

MIT Open Access Articles

Lifting for Simplicity: Concise Descriptions of Convex Sets

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Fawzi, Hamza, Gouveia, Joao, Parrilo, Pablo A, Saunderson, James and Thomas, Rekha R. 2022. "Lifting for Simplicity: Concise Descriptions of Convex Sets." SIAM Review, 64 (4).

Published Version: <https://doi.org/10.1137/20M1324417>

Publisher: Society for Industrial & Applied Mathematics (SIAM)

Permanent Link: <https://hdl.handle.net/1721.1/165366>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Lifting for Simplicity: Concise Descriptions of Convex Sets*

Hamza Fawzi[†]
Joao Gouveia[‡]
Pablo A. Parrilo[§]
James Saunderson[¶]
Rekha R. Thomas^{||}

Abstract. This paper presents a selected tour through the theory and applications of lifts of convex sets. A lift of a convex set is a higher-dimensional convex set that projects onto the original set. Many convex sets have lifts that are dramatically simpler to describe than the original set. Finding such simple lifts has significant algorithmic implications, particularly for optimization problems. We consider both the classical case of polyhedral lifts, described by linear inequalities, as well as that of spectrahedral lifts, defined by linear matrix inequalities, with a focus on recent developments related to spectrahedral lifts.

Given a convex set, ideally we would like to either find a (low-complexity) polyhedral or spectrahedral lift or find an obstruction proving that no such lift is possible. To this end, we explain the connection between the existence of lifts of a convex set and certain structured factorizations of its associated slack operator. Based on this characterization, we describe a uniform approach, via sums of squares, to the construction of spectrahedral lifts of convex sets and illustrate the method on several families of examples. Finally, we discuss two flavors of obstruction to the existence of lifts: one related to facial structure, and the other related to algebraic properties of the set in question.

Rather than being exhaustive, our aim is to illustrate the richness of the area. We touch on a range of different topics related to the existence of lifts and present many examples of lifts from different areas of mathematics and its applications.

Key words. convex sets, polyhedral lifts, positive semidefinite cone, spectrahedral lifts, cone factorization, slack operator

MSC codes. 52-02, 90-02

DOI. 10.1137/20M1324417

*Received by the editors March 11, 2020; accepted for publication (in revised form) November 17, 2021; published electronically November 3, 2022.

<https://doi.org/10.1137/20M1324417>

Funding: The work of the second author was partially supported by the Centre for Mathematics of the University of Coimbra through grant UIDB/00324/2020, funded by the Portuguese Government through FCT/MCTES. The work of the third author was partially supported by the National Science Foundation (NSF) through grant CCF-1565235. The work of the fourth author was partially supported by an Australian Research Council Discovery Early Career Researcher Award (project DE210101056). The work of the fifth author was partially supported by NSF grant DMS-1719538.

[†]Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge, CB3 0WA, United Kingdom (h.fawzi@damtp.cam.ac.uk).

[‡]CMUC, Department of Mathematics, University of Coimbra, 3000-501 Coimbra, Portugal (jgouveia@mat.uc.pt).

[§]Laboratory for Information and Decision Systems (LIDS), Massachusetts Institute of Technology, Cambridge, MA 02139 USA (parrilo@mit.edu).

[¶]Department of Electrical and Computer Systems Engineering, Monash University, VIC 3800, Australia (james.saunderson@monash.edu).

^{||}Department of Mathematics, University of Washington, Seattle, WA 98195-4350 USA (rrthomas@uw.edu).

Contents

I	Introduction	867
1.1	Why Lifts? Two Simple Examples	868
1.2	Lifting to Nonpolyhedral Sets	870
1.3	The Conic Point of View	872
1.4	Lifts in Mathematics	873
1.5	Organization	874
2	Four Examples	875
2.1	Lifting 0/1 Polytopes to Flow Polytopes via Binary Decision Diagrams	875
2.2	The Chain Polytope of a Poset and the Lovász Theta Body of a Graph	877
2.3	Convex Polynomials and Their Epigraphs	882
2.4	The Honeycomb Lift of the Horn Cone	884
3	Slack Operators	887
3.1	Slack Matrices	887
3.2	Nonnegative Factorizations and Yannakakis' Theorem	889
3.3	The Slack Operator	891
3.4	Cone Factorizations of Slack Operators	892
3.5	Bregman Divergence as a Slack Operator	895
4	Constructions of Lifts	896
4.1	Spectrahedral Lifts and Sums of Squares	897
4.2	Hierarchies	899
4.3	Symmetric Convex Bodies	900
4.4	Linear Programming Lifts	901
5	Obstructions and Lower Bounds	902
5.1	Obstructions Based on Facial Structure	903
5.2	Algebraic Obstructions	907
6	Discussion	911
	References	914

I. Introduction. The representation of a convex set in \mathbb{R}^n is a crucial aspect of algorithmic primitives for convex sets, such as testing membership, computing volume, or optimization. The study of representations of *convex* sets is of particular importance in the context of optimization, although it is also of interest well beyond this setting. An established idea for representing a convex set, which has received much attention recently, is to express it as the linear projection of another, “simpler” convex set.

This idea is very powerful—there are by now many examples of convex sets that require highly complicated descriptions in their ambient space, but for which one can find a convex set in a higher-dimensional space, with a concise description, that projects to it. Figure 1 illustrates this idea.

Such concise “lifts” can lead to considerable algorithmic efficiencies for applications such as linear optimization over the original set. They effectively exploit the




Fig. 1 A complicated convex set might admit a concise “lift,” which is a higher-dimensional convex set with a simpler description that projects to it.

nonuniqueness of possible preimages of a projection, in order to find a more computationally convenient description. This paper will take the reader on a selective tour of key ideas, examples, and applications related to representations of convex sets as projections.

1.1. Why Lifts? Two Simple Examples. To motivate our topic, consider *linear programming*, which is the problem of optimizing a linear function over a polytope. The number of linear inequalities needed to represent the polytope can be thought of as a measure of its complexity, as it directly affects the efficiency of certain algorithms for linear programming over the polytope. The obvious linear inequality description of a polytope $P \subset \mathbb{R}^n$ requires one inequality for every facet of P (a codimension-one face). However, the number of facets of P may be very large. In such situations, one might ask whether there is a more compact representation of P . One possible way out is to move to a higher dimension and write P as the linear projection of a polytope $Q \subset \mathbb{R}^\ell$, where $\ell > n$. Such a Q is called a *lift* of P . If, in addition, both ℓ and the number of facets of Q are small in size, say, polynomial in n , then Q is referred to as a *small/compact lift* of P . Such a small lift can be very useful in optimization, since any linear optimization problem over P can be solved via linear optimization over Q . Indeed, if $P = \pi(Q)$, where π is a linear map, then for any cost vector c we have

$$\max_{x \in P} \langle c, x \rangle = \max_{y \in Q} \langle c, \pi(y) \rangle.$$

The right-hand side is a linear optimization problem over Q , which can potentially be much easier to solve.

EXAMPLE 1.1. Consider the ℓ_1 -norm ball in \mathbb{R}^n , which is the n -dimensional cross-polytope defined as the convex hull of the $2n$ vectors $\pm e_i$, $i = 1, \dots, n$, where e_i is the i th standard basis vector in \mathbb{R}^n . This polytope has 2^n facets, and its inequality description is $C_n = \{x \in \mathbb{R}^n : \sum_{i=1}^n \pm x_i \leq 1\}$. The three-dimensional cross-polytope C_3 is shown in Figure 2.

Let π_x denote the projection from \mathbb{R}^{2n} to \mathbb{R}^n sending $(x, y) \mapsto x$. Then $C_n = \pi_x(Q_n)$, where $Q_n = \{(x, y) \in \mathbb{R}^{2n} : \sum_{i=1}^n y_i = 1, -y_i \leq x_i \leq y_i \forall i = 1, \dots, n\}$. Note that Q_n has only $2n$ facets, in contrast to the exponentially many facets of C_n .

Even though the cross-polytope has 2^n facets, it has only $2n$ vertices, so it is




Fig. 2 Cross-polytope C_3 .

perhaps not too surprising that it has a small lift. Next we consider an example of a polytope that has both exponentially many facets and exponentially many vertices.

EXAMPLE 1.2. The permutahedron $\Pi_n \subset \mathbb{R}^n$ is the convex hull of the $n!$ vectors obtained by permuting the coordinates of $(1, 2, \dots, n)$. This is an $(n - 1)$ -dimensional polytope with $2^n - 2$ facets, one corresponding to each proper subset of $[n]$. Figure 3 shows Π_3 and Π_4 .




Fig. 3 On the left is the permutahedron Π_3 shown in the ambient space $x_1 + x_2 + x_3 = 6$. On the right is the permutahedron Π_4 shown in the ambient space $x_1 + x_2 + x_3 + x_4 = 10$.

The Birkhoff polytope B_n , which is the convex hull of all permutation matrices in $\mathbb{R}^{n \times n}$, is a lift of the permutahedron Π_n . The linear map $X \in \mathbb{R}^{n \times n} \mapsto (1, 2, \dots, n)X \in \mathbb{R}^n$ subjects B_n onto Π_n . The Birkhoff-von Neumann theorem [6, 93] states that B_n coincides with the set of doubly stochastic matrices, i.e.,

$$B_n = \left\{ X \in \mathbb{R}^{n \times n} : \sum_{j=1}^n X_{ij} = 1 \forall i \in [n], \sum_{i=1}^n X_{ij} = 1 \forall j \in [n], X_{ij} \geq 0 \forall 1 \leq i, j \leq n \right\}.$$

One can see from this description that B_n has n^2 facets given by the nonnegativity conditions $X_{ij} \geq 0$ and lives in \mathbb{R}^ℓ , where $\ell = n^2$. Hence B_n is a small lift of Π_n .

In general, lifts are far from being unique. For instance, in Figure 4 we show two different polyhedral lifts of a regular octagon. These are nonisomorphic, since the combinatorial structure of the polytopes is different, although both lifts have the minimum possible number of facets (six, in this case).




Fig. 4 Two nonisomorphic polyhedral lifts of a regular octagon, with six facets each.

1.2. Lifting to Nonpolyhedral Sets. The two examples of lifts we saw above were polytopes. From the computational point of view, polytopes are feasible sets of linear programs. A significant portion of this survey will be devoted to linear programming's relative, *semidefinite programming*. *Semidefinite programs* [90] are linear optimization problems where one optimizes a linear function $\sum c_i x_i$ subject to constraints described by a linear matrix inequality (LMI) of the form

$$(1.1) \quad x_1 A_1 + \cdots + x_n A_n \preceq B,$$

where A_1, \dots, A_n, B are symmetric matrices. The notation $A \preceq B$ means that $B - A$ is *positive semidefinite* (psd). The relation $A \preceq B$ (equivalently, $B \succeq A$) is a partial order on the set of symmetric matrices of the same size. LMIs of the form (1.1) describe convex sets that are, in general, nonpolyhedral. Such sets are called *spectrahedra*. We say that a convex set C has a *spectrahedral lift* if it can be expressed as the projection of a spectrahedron. The motivation is clear: if C has a spectrahedral lift, then the maximum of a linear function on C can be computed using semidefinite programming.

EXAMPLE 1.3. Consider the spectrahedron

$$\mathcal{E}_3 = \left\{ (x, y, z) \in \mathbb{R}^3 : \begin{bmatrix} 1 & x & y \\ x & 1 & z \\ y & z & 1 \end{bmatrix} \succeq 0 \right\}.$$

This spectrahedron is known as the *elliptope* and arises in the famous *semidefinite relaxation for the maximum cut problem* [32]. The three-dimensional elliptope is shown in Figure 5 along with various projections.

Why Focus on Convex Sets? The sets we have considered so far are all convex, and this survey will be dedicated to lifts of convex sets. At first glance, focusing exclusively on lifts of convex sets may seem quite restrictive. However, essentially any linearly parameterized family of optimization problems can be rewritten as linear optimization over a convex set. We illustrate this idea through a simple example. Suppose we have the family of optimization problems

$$\max \{ \alpha \sin t + \beta \cos t : t \in [0, 2\pi] \},$$

where $\alpha, \beta \in \mathbb{R}$ parameterize a choice of objective function. Then we can set $x := \sin t$ and $y := \cos t$ and rewrite the family of problems as $\max \{ \alpha x + \beta y : x^2 + y^2 = 1 \}$. Since the objective function is now linear, this problem is equivalent to maximizing the linear function $\alpha x + \beta y$ over the convex hull of the circle, i.e., the unit disc in



Fig. 5 Two different linear images of a spectrahedron. (Note that the linear image on the left is a polytope!)

\mathbb{R}^2 . This last reformulation results in a *convex program*, since the objective function is linear and the feasible region is convex. Figure 6 shows the reformulations for two different objective functions.




Fig. 6 Reformulating the problem of maximizing $\alpha \cos(t) + \beta \sin(t)$ over $t \in [0, 2\pi]$ as a convex optimization problem for $(\alpha, \beta) = (1, 1)$ (solid) and $(\alpha, \beta) = (0, -1)$ (dashed).

The above idea makes sense in much more generality. Suppose we want to reformulate a family of problems of the form $\max\{\sum_{i=1}^k \alpha_i b_i(x) : x \in S\}$, where the $b_i(x)$ are fixed continuous functions from $\mathbb{R}^n \rightarrow \mathbb{R}$, $S \subset \mathbb{R}^n$ is a compact set, and the α_i parameterize the problem instance. This can be rephrased as

$$\max \left\{ \sum_{i=1}^k \alpha_i z_i : z \in b(S) \right\}, \quad \text{where } b(S) := \{(b_1(x), \dots, b_k(x)) : x \in S\}.$$

Since the objective function is linear, we can now optimize over the convex hull of $b(S)$ and obtain the equivalent convex problem $\max\{\sum \alpha_i z_i : z \in \text{conv}(b(S))\}$. The approach is most interesting when a subspace of objective functions is reformulated simultaneously. In the special case where we consider only a single objective function b , this reduces to the (rather tautological) convex optimization problem of finding the largest point in the interval $\text{conv}(b(S))$.

As an example of this general convexification approach, consider applying it to a family of polynomial optimization problems. For a vector $a \in \mathbb{N}^n$, let $|a|$ denote its 1-norm and let x^a denote the monomial $x_1^{a_1} x_2^{a_2} \cdots x_n^{a_n}$. Suppose we wish to maximize

any degree d polynomial $p(x) = \sum_{|a| \leq d} p_a x^a$ over all $x \in S \subset \mathbb{R}^n$. This is equivalent to the convex optimization problem

$$\max \left\{ \sum_{|a| \leq d} p_a y_a : y \in \text{conv}\{(x^a)_{|a| \leq d} : x \in S\} \right\}.$$

Note that the convex sets that arise from these reformulations are, in general, not polyhedra, and so it is not at all obvious how to computationally represent these sets.

In combinatorial optimization we are constantly faced with linear optimization problems over finite sets in \mathbb{R}^n , mostly subsets of vectors in $\{0, 1\}^n$. Any such problem is equivalent to a linear program over the polytope that is the convex hull of the finitely many feasible solutions. We will see examples of lifts of such polytopes in later sections.

1.3. The Conic Point of View. The formal setting we consider in this paper is that of *conic lifts*, which permits an elegant treatment of both polyhedral and spectrahedral lifts in a unified way. This treatment is also motivated from optimization theory, where linear and semidefinite programming are seen as special cases of *conic programming* [74, 4]. To motivate the definition that follows, note that a polyhedron $P = \{x \in \mathbb{R}^n : Ax \leq b\}$, where $A \in \mathbb{R}^{m \times n}$ can be equivalently seen (via a linear bijection) as the intersection $\mathbb{R}_+^m \cap L$, where \mathbb{R}_+^m is the nonnegative orthant in \mathbb{R}^m and $L = b + \text{colspan}(A)$ is an affine space. Similarly, the spectrahedron defined by (1.1) is linearly isomorphic to $\mathcal{S}_+^m \cap L$, where \mathcal{S}_+^m is the convex cone of $m \times m$ symmetric psd matrices and $L = B + \text{span}(A_1, \dots, A_n)$. With these examples in mind, we can state the formal definition of a conic lift that we adopt in this paper.

DEFINITION 1.4. *Let $C \subset \mathbb{R}^n$ be a convex set and $K \subset \mathbb{R}^\ell$ be a closed convex cone. We say that C has a K -lift if $C = \pi(K \cap L)$, where L is an affine space in \mathbb{R}^ℓ and $\pi : \mathbb{R}^\ell \rightarrow \mathbb{R}^n$ is a linear map. The convex set $K \cap L$ is called a K -lift of C .*

If $C = \pi(\mathbb{R}_+^m \cap L)$ (respectively, $C = \pi(\mathcal{S}_+^m \cap L)$), we say that C has a polyhedral (respectively, spectrahedral) lift of *size* m . Going back to Example 1.2, the Birkhoff polytope B_n has $O(n^2)$ inequalities, none of which are redundant, and thus is a polyhedral lift of Π_n of size $O(n^2)$. The smallest m for which $C = \pi(\mathbb{R}_+^m \cap L)$ is called the *linear extension complexity* of C , and the smallest m for which $C = \pi(\mathcal{S}_+^m \cap L)$ is called the *semidefinite extension complexity* of C . It was shown by Goemans [31] that the linear extension complexity of Π_n is $O(n \log n)$. The proof involves constructing an explicit lift of Π_n of that size by using sorting networks.

Another convex cone of special interest in conic programming is the *second-order cone* defined as

$$\mathcal{L}_+^3 = \{(x_0, x) \in \mathbb{R} \times \mathbb{R}^2 : \|x\|_2 \leq x_0\}.$$

If a convex set has a $(\mathcal{L}_+^3)^m$ -lift (for some integer m), then we say it has a *second-order cone lift*. Examples of convex sets with second-order cone lifts include ellipsoids and ℓ_p -norm balls for rational $p \geq 1$, which arise in many practical situations. (See, e.g., [4] for these and many more interesting examples.) Figure 7 summarizes the inclusion relationship between different types of K -lifts, when K is one of \mathbb{R}_+^m , $(\mathcal{L}_+^3)^m$, \mathcal{S}_+^m . A discussion and proof of these relationships, in particular, the strict inclusions, will be given in section 5.

REMARK 1.5 (terminology). *We remark that there is differing terminology in the papers that study lifts of convex sets. Polyhedral lifts are sometimes called extended*




Fig. 7 Inclusion relationship between different types of lifts. The fact that the inclusions are strict will be discussed in more detail in section 5. (The definition of a semialgebraic set is given in section 5.2.1.)

formulations, linear lifts, or LP-lifts. *Spectrahedral lifts* are sometimes called semidefinite lifts or psd-lifts. A convex set that admits a spectrahedral lift is a projected spectrahedron or spectrahedral shadow, and the set is said to have an LMI representation or semidefinite representation. *Linear extension complexity* is also known as nonnegative rank and *semidefinite extension complexity* as psd rank. In this paper we will use the terminology we have introduced in the previous paragraphs.

1.4. Lifts in Mathematics. The idea of going to a higher-dimensional space to seek a “nicer” representation of an object is a common theme in many areas of mathematics. A familiar example is *resolution of singularities* of algebraic varieties, one of the oldest and best-known problems in algebraic geometry. Given a singular algebraic variety, the goal is to write it as a proper birational image of a smooth variety. In 1964 Hironaka [53, 54] proved that such resolutions always exist for varieties over fields of characteristic zero such as \mathbb{C} or \mathbb{R} . In simplified terms, Hironaka’s result says that the set of solutions in k^n , where k is a field of characteristic zero, to a system of polynomial equations can always be rationally parameterized by a smooth manifold. (See, e.g., [49] for a tutorial-style exposition.) Besides the theoretical importance of resolutions, they are also of practical relevance; for instance, they allow us to use our arsenal of techniques for optimization over smooth functions to indirectly optimize over a nonsmooth variety, by passing to the lifted smooth variety.

EXAMPLE 1.6. Consider the deltoid curve (see Figure 8), which is the set of solutions in \mathbb{R}^2 of

$$(x^2 + y^2)^2 - 8(x^3 - 3xy^2) + 18(x^2 + y^2) - 27 = 0.$$

As is evident from Figure 8, this curve has three singularities.




Fig. 8 On the left is the deltoid curve, which has 3 singularities. On the right we see the deltoid as the projection of a smooth curve.

The deltoid is polynomially parameterized by a circle, since it can be written as

$$\{(1 + 2 \cos(t) - 2 \sin(t)^2, 2 \sin(t) - 2 \sin(t) \cos(t)) : t \in [0, 2\pi]\}.$$

This is already very useful, but one can actually write it as a linear projection of a smooth curve by simply adding a $\sin(3t)$ term to the parameterization. Consider the curve

$$Y = \{(1 + 2 \cos(t) - 2 \sin(t)^2, 2 \sin(t) - 2 \sin(t) \cos(t), \sin(3t)) : t \in [0, 2\pi]\} \subseteq \mathbb{R}^3.$$

The curve Y is smooth, and we can see it projecting onto the deltoid.

Here is another instance of the usefulness of “lifting.” Given a probability distribution $\pi_{\mathcal{X}}$ defined on some space \mathcal{X} , a fundamental problem in probability and statistics is to generate a sample from the distribution. One way to do this is to construct an ergodic Markov chain with stationary distribution given by $\pi_{\mathcal{X}}$ and then generate sample paths for the Markov chain. A key measure of the performance of such a sampling method is the *mixing time* of the Markov chain, a measure of how fast the distribution of the samples converges to the stationary distribution.

A *lifting* of a Markov chain is a new Markov chain on a larger state space $\mathcal{X} \times \mathcal{Y}$ so that the marginal on the \mathcal{X} space is the original Markov chain. In certain situations, carefully designed liftings of a Markov chain mix faster than the original chain, leading to faster sampling methods [14]. More generally, many methods for sampling involve introducing auxiliary variables and a probability distribution $\pi_{\mathcal{X} \times \mathcal{Y}}$ on a larger state space, running a Markov chain Monte Carlo (MCMC) method to (approximately) sample from that distribution, and then taking the marginal on the \mathcal{X} variables. Such strategies often tend to mix faster in practice than an MCMC method working directly in the space of the original distribution. These are often called *auxiliary variable samplers* and include approaches such as Hamiltonian Monte Carlo (HMC) [72] and slice sampling [71].

On a philosophical note, recall that in *quantifier elimination*, one is interested in finding a set that has been specified in terms of quantifiers and extra variables. This is a projection operation that is, in general, hard to do. For example, projecting out variables from a system of linear inequalities can be done by *Fourier–Motzkin elimination* [96, Chapter 1], which can take an exponential number of steps. The concept of a lift is inverse to the idea of projection. We start with a highly complicated (convex) set and ask if it can be written as the projection of an (if possible, simple) convex set in some higher dimension.

1.5. Organization. There are a number of applications, drawn from a wide range of areas of mathematics, that illustrate the usefulness of lifted representations of convex sets. In section 2 we give four very different examples, to bring home this point.

Given a convex set $C \subset \mathbb{R}^n$ and a closed convex cone $K \subset \mathbb{R}^\ell$, the first question one can ask is whether C admits a K -lift at all. In section 3 we build the theory that allows us to answer this foundational question. The key ingredient is the *slack operator* of a convex set, which is described in section 3.3. This is a generalization of the *slack matrix* of a polytope that was introduced by Yannakakis [95] to study polyhedral lifts of polytopes. The slack matrix of a polytope is defined in section 3.1, and Yannakakis’ Theorem that decides the existence of an \mathbb{R}_+^m -lift of a polytope $P \subset \mathbb{R}^n$ is described in section 3.2. The main result of section 3 says that a convex set $C \subset \mathbb{R}^n$ has a K -lift for some closed convex cone $K \subset \mathbb{R}^\ell$ if and only if the

slack operator of C factorizes through K and its polar cone K^* . This result is stated and discussed in section 3.4, and we illustrate it with several examples. The slack operator of a convex set can also be seen as a generalization of the notion of Bregman divergence of a convex function. This analogy is discussed in section 3.5.

There are many creative constructions of lifts of convex sets in the literature, and in section 4 we focus on construction methods for spectrahedral lifts of convex sets. We emphasize in particular the connection with *sum-of-squares* certificates of nonnegativity, and semidefinite hierarchies. This viewpoint is illustrated with several examples. We briefly discuss the analogous viewpoint for polyhedral lifts and establish the connection with well-known methods such as the Sherali–Adams hierarchy [84].

Having addressed the existence question and seen some construction methods, a natural next step is to understand obstructions to the existence of a K -lift of C . In section 5 we discuss two flavors of obstruction. The first is combinatorial in nature and depends on the facial structure of both C and K . Such obstructions are already powerful enough to show interesting results such as the result that the psd cone \mathcal{S}_+^3 does not admit a lift into a product of second-order cones [20]. Next we consider algebraic obstructions to the existence of lifts. We prove lower bounds on the semidefinite extension complexity of C based on the degree of its algebraic boundary as well as $\dim(C)$. Finally, we discuss the recent result of Scheiderer [82] that not all convex semialgebraic sets admit a spectrahedral lift and provide concrete examples of such sets (see Figure 7).

The study of lifts of convex sets has many results and theories beyond those we have chosen to present in this paper. We conclude with a brief discussion of related directions and results, as well as pointers to surveys and expository articles on these topics.

2. Four Examples. In the rest of this paper we will study lifts of convex sets as in Definition 1.4. We begin with four examples of lifts, each from a different area of mathematics, to illustrate the ubiquity and usefulness of the idea even beyond the context of optimization. This section can be read independently of the rest of the paper and can be skipped if the reader wishes to get to the mathematical theory behind lifts, which begins in section 3.

2.1. Lifting 0/1 Polytopes to Flow Polytopes via Binary Decision Diagrams.

A *binary decision diagram* (BDD) or *branching program* is a data structure that represents a Boolean function. It is a directed acyclic graph in which exactly one node has in-degree zero (the *source node*), exactly two nodes have out-degree zero (the *sink nodes*), and every nonsink node (known as a *decision node*) has out-degree two (see Figure 9). The sink nodes are labeled 0 and 1. Every decision node is labeled with a variable x_1, x_2, \dots, x_n and has one outgoing arc labeled 0 and one outgoing arc labeled 1 (which are usually depicted as dashed and solid arcs, respectively). Any assignment of Boolean values to the variables, i.e., $x_i = a_i \in \{0, 1\}$ for all i , specifies a path from the source to one of the sinks. The path is defined by taking, for each decision node labeled x_i , the outgoing edge labeled a_i (see Figure 9). Whether the path associated with $a \in \{0, 1\}^n$ ends at the 0-sink or the 1-sink defines a Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$. The diagram on the left of Figure 9 shows a BDD representing the function

$$f_{\text{xor}}(x_1, x_2, x_3, x_4) = x_1 + x_2 + x_3 + x_4 \pmod{2}.$$

It has the additional property that the variables are totally ordered and, along every source-sink path, every variable appears at most once in a way that respects the or-




Fig. 9 On the left is an (ordered) BDD corresponding to the Boolean function f_{xor} that computes the sum of the inputs modulo 2. Solid lines represent arcs labeled by 1, and dashed lines represent arcs labeled by 0. The path corresponding to the assignment $(1, 1, 0, 0)$ to the variables is highlighted in bold. On the right is a zero-suppressed BDD representing the set $\{x \in \{0, 1\}^4 : f_{\text{xor}}(x) = 1\}$.

der. Such a BDD is called an *ordered BDD* (OBDD) or *oblivious read-once branching program* [9]. If a Boolean function has an OBDD representation with a small number of nodes, then this automatically leads to efficient algorithms for counting the number of satisfying assignments, sampling from the uniform distribution on satisfying assignments, and solving the integer linear program $\max_{x \in \{0, 1\}^n : f(x) = 1} \sum_i a_i x_i$. Moreover, natural operations on Boolean functions lead to corresponding operations on OBDDs.

Associated with a Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is the 0/1 polytope $P_f = \text{conv}\{x \in \{0, 1\}^n : f(x) = 1\}$. It turns out that we can construct a polyhedral lift of P_f from an OBDD representation of f . Let A denote the set of arcs of the OBDD. For any subset S of arcs, let $\chi^S \in \{0, 1\}^A$ denote its *characteristic vector* defined as $(\chi^S)_i = 1$ if $i \in S$ and $(\chi^S)_i = 0$ otherwise. Given an OBDD B_f representing f , consider the *flow polytope*

$$F(B_f) = \text{conv} \{ \chi^S \in \{0, 1\}^A : S \subseteq A \text{ is a directed path from the source to the 1-sink in } B_f \}.$$

For $i = 1, 2, \dots, n$, let A_i denote the set of arcs leaving nodes of the OBDD labeled with variable i .

PROPOSITION 2.1. *The polytope $F(B_f)$ is a polyhedral lift of the 0/1 polytope P_f . The linear map defining the lift is $\pi : \mathbb{R}^A \rightarrow \mathbb{R}^n$ defined by $[\pi(y)]_i = \sum_{a \in A_i} y_a$ for $i = 1, 2, \dots, n$.*

Proof. There is a one-to-one correspondence between paths from the source to the 1-sink in the OBDD and vertices of P_f . Moreover, every such path contains at most one arc from each A_i . As such, the vertices of $F(B_f)$ are mapped to those of P_f by π . \square

Since $F(B_f)$ is an example of a network flow polytope, it can be described using

one linear inequality for each arc in the network. Indeed, it is a fundamental result (see, e.g., [83]) that

$$F(B_f) = \{y \in \mathbb{R}^A : y \geq 0, By = b\},$$

where $b \in \mathbb{R}^V$ and $B : \mathbb{R}^A \rightarrow \mathbb{R}^V$ are defined by

$$b_v = \begin{cases} 1 & \text{if } v \text{ is the 1-sink,} \\ -1 & \text{if } v \text{ is the source,} \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad [By]_v = \sum_{a \in v_{\text{in}}} y_a - \sum_{a \in v_{\text{out}}} y_a \quad \forall v \in V,$$

and where $v_{\text{in}} \subseteq A$ and $v_{\text{out}} \subseteq A$ denote the sets of incoming arcs and outgoing arcs from vertex $v \in V$. To summarize, any OBDD representation of a Boolean function leads to a polyhedral lift of the corresponding 0/1 polytope of the same size as the OBDD.

COROLLARY 2.2. *If $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is a Boolean function with an OBDD representation having $|A|$ arcs, then the polytope $P_f = \{x \in \{0, 1\}^n : f(x) = 1\}$ has a polyhedral lift of size $|A|$.*

EXAMPLE 2.3 (odd parity polytope). *The odd parity polytope is the convex hull of binary vectors with an odd number of ones. From Figure 9 it is clear that the Boolean function f_{XOR} has an OBDD with $4(n-1) + 2 = 4n - 2$ arcs. It follows that the odd parity polytope has a polyhedral lift of size $4n - 2$.*

There is a nontrivial upper bound on the size of an OBDD of a general Boolean function f . This gives a corresponding upper bound on the size of a polyhedral lift of an arbitrary 0/1 polytope.

THEOREM 2.4. *Any 0/1 polytope in \mathbb{R}^n has a polyhedral lift of size at most $\frac{6.25}{n} \cdot 2^n$.*

Proof. Any n -variable Boolean function has an OBDD with at most $\frac{3.125}{n} \cdot 2^n$ nodes [66, Theorem 1], and hence at most $\frac{6.25}{n} \cdot 2^n$ arcs. The result then follows from Corollary 2.2. \square

There are many variations and generalizations of OBDDs, a number of which naturally give rise to flow-polytope lifts of 0/1 polytopes. A similar approach can be used to construct polyhedral lifts based on free BDDs (in which every variable appears at most once on every source-sink path, but not necessarily in the same order [68, 94]). The lifted representations of 0/1 polytopes described in this section can be simplified by using zero-suppressed OBDDs (ZBDDs) [69]. These are OBDDs where any arcs leading only to the 0-sink are removed, and the corresponding diagram is reduced by removing any redundancy. A ZBDD is shown on the right in Figure 9. Clearly these arcs do not change any path from the source to the 1-sink, so the polyhedral lifts described above can (possibly) be reduced in size by using a ZBDD representation instead.

2.2. The Chain Polytope of a Poset and the Lovász Theta Body of a Graph.

2.2.1. Chain Polytopes. Let \mathcal{P} be a finite partially ordered set (poset), where the partial order is denoted by \prec ,¹ and let $\mathbb{R}^{\mathcal{P}}$ denote the Euclidean space whose coordinates are indexed by the elements of \mathcal{P} . We start by quickly recalling a series of basic

¹We also use this notation for the specific partial order on symmetric matrices induced by the psd cone. Occasionally both will be used in close proximity, but the meaning should be clear from the context.

concepts from posets that will be used throughout the example. The *characteristic vector* of an element $a \in \mathcal{P}$ is the vector in $\mathbb{R}^{\mathcal{P}}$ with 1 in the coordinate indexed by a and with 0 everywhere else. An *antichain* in \mathcal{P} is a collection of elements of \mathcal{P} such that no two in the collection are comparable in the partial order. A *chain* in \mathcal{P} , on the other hand, is a collection of elements that are totally ordered. The characteristic vector of a collection of elements in \mathcal{P} , for instance, a chain or antichain, is the sum of the characteristic vectors of all elements in that collection. A closely related concept to an antichain is that of a *filter*, i.e., a set I such that $a_1 \in I$ and $a_2 \succeq a_1$ implies $a_2 \in I$. There is a natural bijection between filters I and antichains A , given by

$$I = \{y \in \mathcal{P} : y \succeq x \text{ for some } x \in A\} \leftrightarrow A = \{\text{minimal elements of } I\}.$$

Finally, we say that a_2 *covers* a_1 if $a_1 \prec a_2$ and there is no $x \in \mathcal{P}$ with $a_1 \prec x \prec a_2$. We call this a *cover relation*.

EXAMPLE 2.5. Consider \mathcal{P} to be the poset of the faces of the triangular prism illustrated in Figure 10 ordered by inclusion. We will usually represent such a poset by its Hasse diagram. This is simply a graph where the nodes are the elements of the poset, drawn in such a way that if $a_1 \preceq a_2$, then a_1 is below a_2 and the edges represent the cover relations. In Figure 10 we can also see the Hasse diagram of the poset of faces of the triangular prism.




Fig. 10 Triangular prism and the Hasse diagram of its face poset.

This poset has a single maximal element, the entire prism, and a minimal element, the empty face. An example of a chain is $(\emptyset, A, AB, ABDE, ABCDEF)$ which, in this case, is maximal. An example of a maximal antichain is $\{ABC, AD, BE, CF, DEF\}$. The filter mapping to that antichain is the collection of all the elements that are greater than or equal to any of the elements, namely,

$$\{ABC, AD, BE, CF, DEF, ABDE, ACDF, BCEF, ABCDEF\}.$$

The *chain polytope* $\text{CHAIN}(\mathcal{P})$ of \mathcal{P} , introduced in [87], is the convex hull in $\mathbb{R}^{\mathcal{P}}$ of the characteristic vectors of all antichains of \mathcal{P} . The name comes from its inequality

description:

$$(2.1) \quad \text{CHAIN}(\mathcal{P}) = \{x \in \mathbb{R}_+^{\mathcal{P}} : 0 \leq x_{a_1} + x_{a_2} + \dots + x_{a_k} \leq 1 \\ \forall \text{ chains } a_1 \prec a_2 \prec \dots \prec a_k \text{ in } \mathcal{P}\}.$$

The correctness of this inequality description will become clear at the end of this section. The facets of $\text{CHAIN}(\mathcal{P})$ are given by the nonnegativity constraints $0 \leq x_a$ for all $a \in \mathcal{P}$ and the inequalities $x_{a_1} + x_{a_2} + \dots + x_{a_k} \leq 1$ for all maximal chains $a_1 \prec a_2 \prec \dots \prec a_k$ in \mathcal{P} .

EXAMPLE 2.6. *There are five nonisomorphic posets with three elements. These will have three-dimensional chain polytopes, all of which we can see in Figure 11. Note that two have the same chain polytope up to coordinate labeling.*




Fig. 11 All the three-element posets and their corresponding chain polytopes.

The number of vertices and facets of $\text{CHAIN}(\mathcal{P})$ can be superpolynomial in $|\mathcal{P}|$, as the following example shows.

EXAMPLE 2.7. *Let $\mathcal{B}(n)$ denote the Boolean lattice of the subsets of $[n]$. The number of antichains in $\mathcal{B}(n)$ is the n th Dedekind number, which grows at least as fast as $2^{\frac{N}{\sqrt{\log(N)}}}$, where $N = 2^n$ is the number of elements in $\mathcal{B}(n)$. As antichains correspond to vertices of the chain polytope $\text{CHAIN}(\mathcal{B}(n))$, there is a superpolynomial number of them. On the other hand, since all maximal chains define facets of $\text{CHAIN}(\mathcal{B}(n))$, and there are $n!$ maximal chains in $\mathcal{B}(n)$, the number of facets of $\text{CHAIN}(\mathcal{B}(n))$ is of the order $N^{\log(\log(N))}$, which is also superpolynomial in n .*

The *order polytope* of the poset \mathcal{P} was also introduced in [87]. It is given by

$$(2.2) \quad \text{ORDER}(\mathcal{P}) = \{x \in \mathbb{R}^{\mathcal{P}} : 0 \leq x_{a_1} \leq x_{a_2} \leq 1 \forall a_2 \succeq a_1 \text{ in } \mathcal{P}\}.$$

It is, in fact, sufficient to consider the inequalities corresponding to cover relations, together with $x_a \geq 0$ for all minimal elements $a \in \mathcal{P}$ and $x_a \leq 1$ for all maximal elements $a \in \mathcal{P}$. The vertices of $\text{ORDER}(\mathcal{P})$ are the characteristic vectors of filters, which we have seen to have a natural bijection to antichains, which leads to the following result.

PROPOSITION 2.8 ([87, Theorem 3.2]). *The map $f : \mathbb{R}^{\mathcal{P}} \rightarrow \mathbb{R}^{\mathcal{P}}$ defined by*

$$f(x)_{a_i} = \begin{cases} x_{a_i} & \text{if } a_i \text{ is a minimal element of } \mathcal{P}, \\ \min\{x_{a_i} - x_{a_j} \mid a_i \text{ covers } a_j\} & \text{otherwise} \end{cases}$$

is piecewise linear and maps $\text{ORDER}(\mathcal{P})$ bijectively to $\text{CHAIN}(\mathcal{P})$, preserving the bijection between vertices.

This is not really a lift of the chain polytope, as the map involved is not linear. However, as is current practice in linear optimization, one can introduce slack variables to linearize the minimum function and obtain the following result.

COROLLARY 2.9. *The chain polytope $\text{CHAIN}(\mathcal{P})$ of a poset \mathcal{P} is the projection onto the z coordinates of $(z, x) \in \mathbb{R}^{\mathcal{P}} \times \mathbb{R}^{\mathcal{P}}$ such that $x \in \text{ORDER}(\mathcal{P})$, $0 \leq z_{a_i} \leq x_{a_i} - x_{a_j}$ for all cover relations $a_i \succ a_j$ and $z_{a_i} = x_{a_i}$ for all minimal elements a_i .*

Corollary 2.9 gives an explicit polyhedral lift of $\text{CHAIN}(\mathcal{P})$ with $O(|\mathcal{P}|^2)$ facets, showing that one can use linear programming to efficiently optimize over the chain polytope. Since posets can be thought of as acyclic digraphs, this allows us to efficiently optimize over sets of mutually unreachable vertices in an acyclic digraph. Further discussions on the algorithmic implications of the chain polytope can be found in section 14.5 of [83].

2.2.2. A Spectrahedral Lift of the Chain Polytope. Next we will see that $\text{CHAIN}(\mathcal{P})$ has a spectrahedral lift of size $O(|P|)$. This lift is a special case of a general construction of spectrahedral lifts of *stable set polytopes* of *perfect graphs*. We describe the general construction below.

The *comparability graph* of the poset \mathcal{P} is the graph whose vertices are the elements of \mathcal{P} and whose edges are (a_1, a_2) for all $a_1 \prec a_2$ (not just cover relations). A *stable set* S of a graph $G = (V(G), E(G))$ is a subset of $V(G)$ such that no pair in S lies in $E(G)$. The characteristic vector of a stable set S in G is the vector in $\{0, 1\}^{V(G)}$ with coordinate 1 in positions indexed by $v \in S$ and 0 everywhere else. The *stable set polytope* of G , denoted as $\text{STAB}(G)$, is the convex hull of the characteristic vectors of all stable sets in G . Note that the antichains in a poset \mathcal{P} are precisely the stable sets in its comparability graph. Hence $\text{CHAIN}(\mathcal{P})$ is the stable set polytope of the comparability graph of \mathcal{P} .




Fig. 12 Icosahedral graph with an optimal 4-coloring. Each maximal monochromatic set of nodes is actually a maximal stable set.

Recall that a *perfect graph* is a graph for which the *chromatic number* equals the *clique number* for every induced subgraph. In Figure 12 we can see that the icosahedral graph is not perfect, as its chromatic number is 4 while its maximal clique is a triangle, and hence its clique number is 3. In general, optimizing a linear function over the stable set polytope of a graph is NP-hard. However, if the graph

G is perfect, then $\text{STAB}(G)$ admits a small spectrahedral lift of size $n + 1$, where $n = |V(G)|$. In particular, we get that the largest stable set in a perfect graph can be found in polynomial time via semidefinite programming [44, 45]. This construction marks the beginning of the use of semidefinite programming in discrete optimization. Comparability graphs of posets are perfect graphs (see, for instance, [33, section 5.7]), and hence the chain polytope $\text{CHAIN}(\mathcal{P})$ has a spectrahedral lift of size $|\mathcal{P}| + 1$. We now describe the general construction from which this lift arises.

Let $G = (V(G), E(G))$ be a graph where $V(G) = [n]$. For a stable set $S \subset [n]$, let χ^S denote its characteristic vector in $\{0, 1\}^n$. Setting $x := \chi^S$, notice that

$$X := \begin{bmatrix} \mathbf{1} \\ x \end{bmatrix} \begin{bmatrix} \mathbf{1} & x^\top \end{bmatrix}$$

is a matrix in \mathcal{S}_+^{n+1} with the following properties:

$$\begin{aligned} X_{00} &= 1, & X_{0i} &= X_{i0} = x_i \quad \forall i \in [n], & X_{ii} &= x_i \quad \forall i \in [n], \\ X_{ij} &= 0 \quad \forall ij \in E(G), & \text{rank}(X) &= 1. \end{aligned}$$

Consider the spectrahedron defined by all the above conditions except for the rank condition:

$$L(G) := \left\{ X \in \mathcal{S}_+^{n+1} : \begin{array}{l} X_{00} = 1, \\ X_{0i} = X_{i0} = x_i \quad \forall i \in [n], \\ X_{ii} = x_i \quad \forall i \in [n], \\ X_{ij} = 0 \quad \forall ij \in E(G) \end{array} \right\}.$$

Now consider the linear map π that maps $X \in \mathcal{S}_+^{n+1}$ to $(X_{01}, X_{02}, \dots, X_{0n}) \in \mathbb{R}^n$.

DEFINITION 2.10. *The projection $\text{TH}(G) := \pi(L(G)) \subset \mathbb{R}^n$ is called the theta body of the graph G .*

Theta bodies were introduced in [45]. For the formulation shown above, see [46, Chapter 9]. Since every vertex of $\text{STAB}(G)$ is a χ^S and the matrix X obtained by lifting χ^S as above lies in $L(G)$, it follows from convexity that $\text{STAB}(G) \subseteq \text{TH}(G)$.

THEOREM 2.11 ([45, Corollary 3.11]). *$\text{STAB}(G) = \text{TH}(G)$ if and only if G is perfect.*

Since the comparability graph of a poset \mathcal{P} is perfect, and $\text{CHAIN}(\mathcal{P})$ is its stable set polytope, we get the following corollary.

COROLLARY 2.12. *The chain polytope $\text{CHAIN}(\mathcal{P})$ admits a spectrahedral lift of size $|\mathcal{P}| + 1$. More precisely,*

$$(2.3) \quad \text{CHAIN}(\mathcal{P}) = \left\{ x \in \mathbb{R}^{\mathcal{P}} : \exists Y \in \mathbb{R}^{\mathcal{P} \times \mathcal{P}} \text{ s.t.} \right. \\ \left. \begin{bmatrix} \mathbf{1} & x^\top \\ x & Y \end{bmatrix} \succeq 0, \text{diag}(Y) = x, y_{ab} = 0 \text{ if } a \succ b \text{ or } a \prec b \right\}.$$

Other examples of perfect graphs are chordal graphs and bipartite graphs. In all of these cases, $\text{STAB}(G)$ admits a spectrahedral lift of size $|V(G)| + 1$.

If $G = ([n], E)$ is perfect, then $\text{STAB}(G)$ also has an explicit (although large) polyhedral description, namely,

$$\text{STAB}(G) = \left\{ x \in \mathbb{R}^n : x \geq 0, \sum_{i \in K} x_i \leq 1 \quad \forall \text{ cliques } K \text{ in } G \right\}.$$

This is why the facets of the chain polytope $\text{CHAIN}(P)$ are given by the maximal chain inequalities in the comparability graph of \mathcal{P} and nonnegativities of variables.

We have seen that the chain polytope admits the polyhedral lift $\text{ORDER}(\mathcal{P})$ of size $O(|\mathcal{P}|^2)$ and the spectrahedral lift $L(G)$ of size $|\mathcal{P}| + 1$. We will see later that no spectrahedral lift of $\text{CHAIN}(\mathcal{P})$ can have size smaller than $|\mathcal{P}| + 1$. More generally, any spectrahedral lift of a polytope P has size at least $\dim(P) + 1$.

2.3. Convex Polynomials and Their Epigraphs. In the area of nonlinear programming, it is traditional to express optimization problems in the form

$$(2.4) \quad \min_x f(x) \text{ subject to } g_i(x) \leq 0 \text{ for } i = 1, 2, \dots, m,$$

for functions $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and $g_1, \dots, g_m : \mathbb{R}^n \rightarrow \mathbb{R}$. If these are all convex functions, then such an optimization problem is called a convex optimization problem. Since the advent of interior point methods in the 1990s, it has become common to reformulate such convex optimization problems as

$$\min_{x,t} t \text{ subject to } \begin{cases} (x,t) \in \text{epi}(f), \\ (x,0) \in \text{epi}(g_i) \text{ for } i = 1, 2, \dots, m, \end{cases}$$

where $\text{epi}(f) = \{(x,t) \in \mathbb{R}^n \times \mathbb{R} : f(x) \leq t\}$ is the *epigraph* of f . Here we have introduced an additional variable and expressed the problem as the minimization of a linear functional over a convex set. This motivates the study of lifted representations of the epigraphs of convex functions.

If we restrict ourselves to differentiable convex functions, one approach to constructing lifted representations of their epigraphs is to make use of the first-order characterization of convexity, i.e.,

$$(2.5) \quad f(x) - [f(y) + \langle \nabla f(y), x - y \rangle] \geq 0 \quad \forall (x,y) \in \mathbb{R}^n \times \mathbb{R}^n.$$

This inequality just tells us that the graph of f is above the graph of its linear approximation at y , for any y . The expression on the left-hand side of (2.5) is the *Bregman divergence* associated with f .

2.3.1. Convex and sos-Convex Polynomials. We now restrict ourselves further to convex *polynomials* of degree at most $2d$. (With minor modifications, the discussion extends to convex rational functions $f(x)/g(x)$ with $g(x) > 0$ for all x .) The Bregman divergence of a convex polynomial f of degree at most $2d$ in n variables is a nonnegative polynomial of degree $2d$ in the $2n$ variables (x,y) . Since the Bregman divergence is linear in f , the first-order characterization of convexity (2.5) tells us that the set of convex polynomials of degree at most $2d$ in n variables is a slice of $\text{Pol}_+^{2n,2d}$, the cone of nonnegative polynomials of degree at most $2d$ in $2n$ variables. Denote by $\text{Pol}^{l,2d}$ the set of all polynomials in l variables of degree at most $2d$.

Whenever $n = 1$ or $2d = 2$ or $(n, 2d) = (2, 4)$, it can be shown [1] that the Bregman divergence of a convex polynomial of degree at most $2d$ in n variables can be written as a *sum of squares*, i.e., there exist $f_1, f_2, \dots \in \text{Pol}^{2n,d}$ such that

$$f(x) - [f(y) + \langle \nabla f(y), x - y \rangle] = \sum_k f_k(x,y)^2 \quad \forall (x,y) \in \mathbb{R}^n \times \mathbb{R}^n.$$

Otherwise, the polynomials for which the Bregman divergence is a sum of squares (called *sos-convex polynomials* [51]) form a strict subset of all convex polynomials with a given degree and number of variables.

Let $\Sigma_+^{\ell, 2d}$ denote the cone of polynomials of degree at most $2d$ in ℓ variables that are sums of squares. This cone has a spectrahedral lift given by

$$\Sigma_+^{\ell, 2d} = \{f \in \text{Pol}^{\ell, 2d} : \exists Q \succeq 0 \text{ such that } f(x) = v_d(x)^\top Q v_d(x) \ \forall x \in \mathbb{R}^\ell\}.$$

Here $v_d(x)$ is the length $\binom{\ell+d}{d}$ vector consisting of all monomials of degree at most d in ℓ variables. Using this, we obtain a spectrahedral lift of the cone of sos-convex polynomials.

PROPOSITION 2.13. *The cone of sos-convex polynomials has the spectrahedral lift*

$$(2.6) \quad \left\{ f \in \text{Pol}^{n, 2d} : \exists Q \succeq 0 \text{ s.t.} \right. \\ \left. f(x) - [f(y) + \langle \nabla f(y), x - y \rangle] = v_d(x, y)^\top Q v_d(x, y) \ \forall (x, y) \in \mathbb{R}^n \times \mathbb{R}^n \right\}.$$

In the cases $n = 1$, $2d = 2$, and $(n, 2d) = (2, 4)$, for which all convex polynomials are sos-convex, Proposition 2.13, in fact, gives a spectrahedral lift of the cone of convex polynomials.

2.3.2. Epigraphs of Convex and sos-Convex Polynomials. We can also use the characterization (2.5) to express the epigraph of any fixed convex polynomial as an affine slice of $\text{Pol}_+^{n, 2d}$, and give a spectrahedral lift of the epigraph of any sos-convex polynomial.

PROPOSITION 2.14. *If f is a convex polynomial, then*

$$(2.7) \quad \text{epi}(f) = \{(x, t) \in \mathbb{R}^n \times \mathbb{R} : t - [f(y) + \langle \nabla f(y), x - y \rangle] \in \text{Pol}_+^{n, 2d}\}.$$

If f is an sos-convex polynomial, then

$$(2.8) \quad \text{epi}(f) = \{(x, t) \in \mathbb{R}^n \times \mathbb{R} : \exists Q \succeq 0 \text{ s.t.} \\ t - [f(y) + \langle \nabla f(y), x - y \rangle] = v_d(y)^\top Q v_d(y) \ \forall y \in \mathbb{R}^n\}.$$

Proof. If f is convex (respectively, sos-convex) and $(x, t) \in \text{epi}(f)$, then $f(x) \leq t$ and so

$$t - [f(y) + \langle \nabla f(y), x - y \rangle] = (t - f(x)) + f(x) - [f(y) + \langle \nabla f(y), x - y \rangle]$$

is nonnegative (respectively, a sum of squares). On the other hand, if $t - [f(y) + \langle \nabla f(y), x - y \rangle] \geq 0$ for all $y \in \mathbb{R}^n$, then putting $y = x$ we have that $t \geq f(x)$, and so $(x, t) \in \text{epi}(f)$. \square

Since the epigraphs of sos-convex polynomials have spectrahedral lifts, we can reformulate convex optimization problems of the form (2.4), in which the objective function and the constraints are sos-convex polynomials, as semidefinite programming problems. Figure 13 shows the feasible region of such a problem for two sos-convex polynomials.

EXAMPLE 2.15 (epigraph of a convex quadratic and convex quadratic programming). *If $f(x) = x^\top A x + 2b^\top x + c$ and A is psd, then f is a convex quadratic. Optimization problems of the form (2.4) in which f and the g_i are all convex quadratics*




Fig. 13 The intersection of the zero-sublevel sets of the (sos-)convex polynomials $g_1(x, y, z) = x^4 + y^4 + z^4 - 3$ and $g_2(x, y, z) = (x - 2)^2 + 2y^2 + 3z^2 - 5$.

are known as (convex) quadratic programs. Since any convex quadratic is sos-convex,

$$\begin{aligned} \text{epi}(f) &= \left\{ (x, t) \in \mathbb{R}^n \times \mathbb{R} : \exists Q \succeq 0 \text{ s.t.} \right. \\ &\quad \left. y^\top A y - 2(Ax)^\top y + (t - c - 2b^\top x) = [1 \quad y^\top] Q \begin{bmatrix} 1 \\ y \end{bmatrix} \quad \forall y \in \mathbb{R}^n \right\} \\ &= \left\{ (x, t) \in \mathbb{R}^n \times \mathbb{R} : \begin{bmatrix} t - c - 2b^\top x & -(Ax)^\top \\ -Ax & A \end{bmatrix} \succeq 0 \right\}. \end{aligned}$$

If $f(x) = x^\top x$ and we consider the $t = 1$ slice of $\text{epi}(f)$, we obtain a well-known spectrahedral representation of the sphere:

$$\{x \in \mathbb{R}^n : x^\top x \leq 1\} = \left\{ x \in \mathbb{R}^n : \begin{bmatrix} 1 & -x^\top \\ -x & I \end{bmatrix} \succeq 0 \right\} = \left\{ x \in \mathbb{R}^n : \begin{bmatrix} 1 & x^\top \\ x & I \end{bmatrix} \succeq 0 \right\},$$

where the last description is obtained via a congruence by $\begin{bmatrix} 1 & 0 \\ 0 & -I \end{bmatrix}$, which preserves positive semidefiniteness.

2.4. The Honeycomb Lift of the Horn Cone. Given the eigenvalues of two $n \times n$ Hermitian matrices A and B , a classical question from linear algebra is to characterize the eigenvalues of $A + B$. Equivalently, the problem is to characterize the triplets $(\lambda, \mu, \nu) \in (\mathbb{R}^n)^3$, where λ is the vector of eigenvalues of an $n \times n$ Hermitian matrix A in weakly decreasing order, μ is that of B , and ν is that of C , such that $A + B + C = 0$. Call such a triplet $(\lambda, \mu, \nu) \in (\mathbb{R}^n)^3$ *valid*. The question of characterizing valid triplets was posed in 1912 by Hermann Weyl, who gave a number of necessary conditions, starting with the obvious trace condition that $\sum \lambda_i + \sum \mu_i + \sum \nu_i = 0$.

In 1962 Horn conjectured that the set of valid triplets forms a convex polyhedral cone and proposed a complete set of linear inequalities that describes it via a recursive recipe [55]. Let $\text{HORN}_n \subset \mathbb{R}^{3n}$ be the set of valid triplets (λ, μ, ν) . The inequalities conjectured by Horn have the form

$$\lambda_{i_1} + \cdots + \lambda_{i_r} + \mu_{j_1} + \cdots + \mu_{j_r} + \nu_{k_1} + \cdots + \nu_{k_r} \geq 0,$$

where $1 \leq r < n$ and all triplets of indices $1 \leq i_1 < \cdots < i_r \leq n$, $1 \leq j_1 < \cdots < j_r \leq n$, and $1 \leq k_1 < \cdots < k_r \leq n$ are in a set $T_{n,r}$. Horn's conjecture describes the set

$T_{n,r}$ recursively in terms of sets $T_{r,s}$ for $1 \leq s < r$. This recipe leads to an exponential number of inequalities for HORN_n .

The Horn conjecture was proved by a combination of the works of Klyachko [58] and Knutson and Tao [59, 60]. Remarkably, the work of Knutson and Tao also shows that the convex set HORN_n admits a small polyhedral lift of size $O(n^2)$.

THEOREM 2.16. *The Horn cone for $n \times n$ Hermitian matrices admits a polyhedral lift called HONEY_n with $O(n^2)$ facets.*

In this short section we describe the lift of Knutson and Tao but we do not attempt to explain the proof, which uses sophisticated arguments from algebraic geometry and representation theory (we refer the interested reader to [60] for more details).

Consider a honeycomb structure as shown in Figure 14. The honeycomb is formed of internal hexagons and has $3n$ outgoing edges, each labeled by λ_i , μ_i , or ν_i in the way shown in the figure.




Fig. 14 The honeycomb structure used to describe the lift of HORN_n for the case $n = 5$.

Let E_n be the set of edges in the n -honeycomb. It is not hard to see that $|E_n| = O(n^2)$. Let $\Gamma_n \subset E_n \times E_n$ be the pairs of edges $(a, b) \in E_n \times E_n$ that are in one of the configurations shown in Figure 15. An example of such a configuration is highlighted in Figure 14.




Fig. 15 Configurations of pairs of edges (a, b) in a honeycomb that appear in the lift of HORN_n . Note that these three patterns are 120° rotations of each other.

We are now ready to describe the lift. Consider the polyhedral cone HONEY_n in \mathbb{R}^{E_n} defined by

$$(2.9) \quad \text{HONEY}_n = \{e \in \mathbb{R}^{E_n} : e_a - e_b \geq 0 \quad \forall (a, b) \in \Gamma_n, \\ e_a + e_b + e_c = 0 \quad \forall \text{ edges } a, b, c \text{ meeting at a common vertex}\}.$$

Consider the map ∂ that projects a vector $e \in \mathbb{R}^{E_n}$ onto the $3n$ outgoing edges of the

honeycomb. The result of Knutson and Tao shows that

$$(2.10) \quad \text{HORN}_n = \{(\lambda, \mu, \nu) \in \mathbb{R}^{3n} : \lambda_1 \geq \cdots \geq \lambda_n, \mu_1 \geq \cdots \geq \mu_n, \nu_1 \geq \cdots \geq \nu_n, \\ \text{and } \exists e \in \text{HONEY}_n \text{ s.t. } (\lambda, \mu, \nu) = \partial(e)\}.$$




Fig. 16 A 3-honeycomb.

EXAMPLE 2.17. Take $n = 3$ and suppose $\lambda = (1, 0, -1)$, $\mu = (2, 1, 0)$, and $\nu = (\nu_1, \nu_2, \nu_3)$ are the vectors of eigenvalues of A, B, C such that $C = -A - B$ and all matrices are 3×3 and Hermitian. We would like to use honeycombs to decide the possible values of ν .

A 3-honeycomb with boundary values λ, μ, ν is shown in Figure 16. There are nine internal edges in this honeycomb, of which three are determined by the boundary values. The remaining edges are labeled e_1, \dots, e_6 as shown. Writing the second condition in (2.9) for all vertices of the honeycomb yields the following system of equations:

$$(2.11) \quad \begin{aligned} e_1 + e_2 = 1 &\Rightarrow e_2 = 1 - e_1, \\ e_1 + e_3 = 0 &\Rightarrow e_3 = -e_1, \\ e_3 + e_5 = 1 + \nu_3 &\Rightarrow e_5 = 1 + \nu_3 + e_1, \\ e_5 + e_6 = -\nu_2 &\Rightarrow e_6 = -1 - \nu_2 - \nu_3 - e_1, \\ e_4 + e_6 = \nu_1 &\Rightarrow e_4 = 1 + (\nu_1 + \nu_2 + \nu_3) + e_1, \\ e_2 + e_4 = -1 &\Rightarrow e_2 = -2 - (\nu_1 + \nu_2 + \nu_3) - e_1. \end{aligned}$$

The first and last equations imply the trace equation $\nu_1 + \nu_2 + \nu_3 = -3$. We have eliminated all the e variables except e_1 .

Now we impose the nonnegativity constraints coming from the first condition in (2.9). For example, we can identify a configuration of the leftmost type in Figure 15 that yields $-1 - e_3 \geq 0$. Using the fact that $e_3 = -e_1$, as derived in (2.11), gives

$e_1 \geq 1$. Proceeding like this with all configurations of edges shown in Figure 15 we arrive at the following inequalities:

$$1 \leq e_1 \leq 2, \quad e_1 \geq 2 + \nu_1 + \nu_2, \quad e_1 \geq 1 + \nu_1, \quad e_1 \geq -\nu_2, \quad e_1 \geq 2 + \nu_2, \\ e_1 \leq 3 + \nu_1 + \nu_2, \quad e_1 \leq 2 + \nu_1.$$

In addition, we also have the chamber inequalities: $\nu_1 \geq \nu_2 \geq -3 - \nu_1 - \nu_2$. The possible values of ν are those that allow a value of e_1 to exist given the above inequalities.

For a concrete example, consider the triple of symmetric matrices

$$A = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}, \quad B = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad \text{and } C = -A - B = \begin{bmatrix} -2 & 0 & 0 \\ 0 & -1 & -1 \\ 0 & -1 & 0 \end{bmatrix},$$

which have eigenvalues $\lambda = (1, 0, -1)$, $\mu = (2, 1, 0)$, and $\nu = (\frac{-1+\sqrt{5}}{2}, \frac{-1-\sqrt{5}}{2}, -2)$, respectively. Plugging $\nu_1 = \frac{-1+\sqrt{5}}{2}$ and $\nu_2 = \frac{-1-\sqrt{5}}{2}$ into the inequalities, we should arrive at a feasible value for e_1 and, indeed, we get $\frac{1+\sqrt{5}}{2} \leq e_1 \leq 2$, which is satisfiable.

In this example, HONEY_3 can be thought of as a cone in \mathbb{R}^{6+9} with variables indexed by e_1, \dots, e_6 and the nine variables in λ, μ, ν . It could be in \mathbb{R}^{9+9} , but recall that the variables associated to three of the bounded edges can be eliminated right away and expressed in terms of λ, μ, ν . The Horn cone is in \mathbb{R}^9 . When we fix λ and μ as in this example, we are looking at a slice of HONEY_3 . The method above derives the inequality description of this slice.

3. Slack Operators. A fundamental question concerning lifts is how to characterize when they exist. More formally, given a convex set C and a convex cone K , how can we decide whether or not C has a K -lift? In this section we focus on answering this basic existence question. In order to do this, we introduce a central tool for studying lifts of convex sets, namely, the *slack operator* of the set. We will see that the existence of a K -lift of a convex set C is equivalent to the existence of a *factorization* of the slack operator of C through the cone K and its dual cone K^* . This connection provides us with a useful algebraic handle on the geometric picture of cone lifts of convex sets.

3.1. Slack Matrices. In the case of polyhedral lifts of polytopes, their existence was characterized in [95]. A key component turns out to be the *slack matrix* of a polytope, which we introduce in this section.

There are two traditional ways of representing a polytope P . We can represent P as a \mathcal{V} -polytope, i.e., as the convex hull of a minimal finite set of points, its vertices. We can also represent P as an \mathcal{H} -polytope, i.e., as a bounded intersection of a finite number of closed half-spaces. It is a fundamental property of polytopes that these two representations are equivalent. Thus we would commonly either list all the vertices of P or list a minimal set of affine inequalities that cut out P , or more precisely, one inequality per facet of P .

EXAMPLE 3.1. Consider the polytope P in \mathbb{R}^3 , shown on the left in Figure 17, with the seven vertices $(\pm 1, \pm 1, -1/2)$, $(\pm 1, -1, 1/2)$, and $(0, -1, 3/2)$. This polytope can also be described by the seven affine inequalities that cut it out:

$$1 - x \geq 0, \quad 1 + x \geq 0, \quad 1 + 2z \geq 0, \quad 1 + y \geq 0, \quad 1 - 2y - 2z \geq 0, \\ 1 - x - \frac{1}{2}y - z \geq 0, \quad 1 + x - \frac{1}{2}y - z \geq 0.$$




Fig. 17 Polytope P with seven vertices and seven facets, and its polar P° .

The slack matrix of P combines the \mathcal{V} and \mathcal{H} descriptions of P into a single, less parsimonious, description. To do this we compute the *slack*, $\beta - h^\top v$, of each vertex v in each facet defining inequality $h^\top x \leq \beta$, obtaining a nonnegative matrix with rows indexed by the facets and columns indexed by the vertices of the given polytope P . For instance, the slack matrix of the polytope P from Example 3.1 is the following 7×7 matrix:

$$S_P = \begin{bmatrix} 2 & 2 & 2 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 2 & 2 & 2 & 1 \\ 0 & 0 & 2 & 0 & 0 & 2 & 4 \\ 0 & 2 & 0 & 0 & 2 & 0 & 0 \\ 4 & 0 & 2 & 4 & 0 & 2 & 0 \\ 3 & 2 & 2 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 3 & 2 & 2 & 0 \end{bmatrix}.$$

The rows of S_P are indexed by the facets of P in the order they are listed above. The columns of S_P are indexed by the vertices of P in the following order:

$$\begin{aligned} & \left(-1, -1, -\frac{1}{2}\right), \left(-1, 1, -\frac{1}{2}\right), \left(-1, -1, \frac{1}{2}\right), \left(1, -1, -\frac{1}{2}\right), \\ & \left(1, 1, -\frac{1}{2}\right), \left(1, -1, \frac{1}{2}\right), \left(0, -1, \frac{3}{2}\right). \end{aligned}$$

Since the inequality defining a facet of P is only defined up to scaling by a positive number, the slack matrix S_P is only defined up to positive scaling of rows.

There is another interesting interpretation of the slack matrix, which brings the duality theory of polytopes to the forefront. Recall that, given a polytope $P \subset \mathbb{R}^n$ with the origin in its interior, one can define its *polar* as the set

$$P^\circ = \{y \in \mathbb{R}^n : \langle x, y \rangle \leq 1 \forall x \in P\},$$

which is again a polytope. A fundamental result in the duality of polytopes is that there is a one-to-one correspondence between the vertices of P° and the facets of P that sends vertex y of P° to the facet of P cut out by the inequality $1 - \langle x, y \rangle \geq 0$. In our example, this means that P° is the polytope with vertices $(\pm 1, 0, 0)$, $(0, 0, -2)$, $(0, -1, 0)$, $(0, 2, 2)$, $(1, 1/2, 1)$, and $(-1, 1/2, 1)$, shown on the right in Figure 17. Note that if the facet inequalities of P are scaled so that they are of the form $1 - \langle x, a_i \rangle \geq 0$

(which is possible since the origin lies in the interior of P), then the transposed matrix S_P^\top is a slack matrix of P° . With the language of duality, we can think of the slack matrix S_P as the map $1 - \langle x, y \rangle$ evaluated at the vertices of $P^\circ \times P$.

3.2. Nonnegative Factorizations and Yannakakis' Theorem. In order to characterize the existence of polyhedral lifts for a polytope, we need to introduce the notion of *nonnegative rank* of a matrix. A *nonnegative factorization* of a nonnegative matrix $M \in \mathbb{R}_+^{n \times k}$ is a pair of nonnegative matrices $A \in \mathbb{R}_+^{m \times n}$ and $B \in \mathbb{R}_+^{m \times k}$ such that $M = A^\top B$. We say m , the inner dimension of the product, is the *size* of the factorization, and we define the *nonnegative rank* of M to be the smallest m for which a nonnegative factorization of M exists. Computing nonnegative ranks of matrices is a hard problem. Nonnegative factorizations and nonnegative rank have many interpretations and applications in diverse fields. An early survey can be found in [15], but there is a growing body of literature on this topic. Somewhat surprisingly, it turns out that this hard algebraic problem encodes the geometric problem of interest—the existence of polyhedral lifts.

THEOREM 3.2 (Yannakakis [95]). *The linear extension complexity of a polytope equals the nonnegative rank of its slack matrix.*

Proof. Suppose a polytope P is cut out by the facet inequalities $h_i^\top x \leq \beta_i$ for $i = 1, \dots, f$, and that $S_P = A^\top B$ is a nonnegative factorization of its slack matrix with inner dimension m . If we let a_i , for $i = 1, 2, \dots, f$, denote the columns of A , then we will show that

$$(3.1) \quad P = \{x \in \mathbb{R}^n : \exists y \in \mathbb{R}_+^m \text{ s.t. } a_i^\top y = \beta_i - h_i^\top x \quad \forall i = 1, \dots, f\}.$$

This equality rewrites P as the projection of a polytope with m facets or, equivalently, as the projection of a slice of \mathbb{R}_+^m . Hence P has a size m polyhedral lift.

To see why (3.1) holds, first note that if p is on the right-hand side, we must have $h_i^\top p = \beta_i - a_i^\top y$ for some nonnegative y . Then, since $a_i^\top y$ is nonnegative, it follows that $h_i^\top p \leq \beta_i$ for all i , and so $p \in P$. On the other hand, if v_j is a vertex of P , let $y = b_j$ be the corresponding column of B , which is nonnegative. Then $a_i^\top y = a_i^\top b_j = [S_P]_{ij} = \beta_i - h_i^\top v_j$. Consequently v_j , and indeed all vertices of P , are on the right-hand side set. Since the right-hand side set is convex, all of P must be contained in it, and we have equality between the two sets.

Suppose now that P can be written as the projection of a polytope Q with m facets. Take the slack matrix S_Q and keep only columns that correspond to vertices that project to vertices of P , and call this reduced $m \times v$ matrix \widetilde{S}_Q . Now any facet inequality on P can be pulled back by the projection to a valid inequality on Q , so, by a version of the Farkas' Lemma, it can be written as a nonnegative combination of facet inequalities of Q . Record the coefficients as column vectors and form a matrix $A \in \mathbb{R}_+^{m \times f}$ with those columns. Then $S_P = A^\top \widetilde{S}_Q$ is a nonnegative factorization of S_P with inner dimension m . \square

EXAMPLE 3.3. *Revisiting Example 3.1, one can now ask for the linear extension complexity of the polytope P . We will prove that it is at least 6 in Example 5.2, using basic considerations about how lifts interact with facial structure that will be introduced in section 5. To show that there is a lift with 6 facets, Theorem 3.2 tells us that it is enough to find a nonnegative factorization of the slack matrix S_P of size 6. One such*

factorization is

$$\begin{bmatrix} 2 & 2 & 2 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 2 & 2 & 2 & 1 \\ 0 & 0 & 2 & 0 & 0 & 2 & 4 \\ 0 & 2 & 0 & 0 & 2 & 0 & 0 \\ 4 & 0 & 2 & 4 & 0 & 2 & 0 \\ 3 & 2 & 2 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 3 & 2 & 2 & 0 \end{bmatrix} = \begin{bmatrix} 2 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 2 & 1 \\ 0 & 0 & 0 & 2 & 0 & 4 \\ 0 & 0 & 2 & 0 & 0 & 0 \\ 0 & 4 & 0 & 2 & 0 & 0 \\ 2 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 2 & 0 \end{bmatrix} \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

It is not hard to recover the lift Q from the factorization: We simply start with an inequality description of P and use the factorization to write the lift as explicitly described in (3.1). We then simply use the equality constraints to eliminate some of the variables y .

EXAMPLE 3.4. Recall that in Corollary 2.9 we have shown that the chain polytope $\text{CHAIN}(\mathcal{P})$ of a poset \mathcal{P} has a small polyhedral lift. By Yannakakis' Theorem, the slack matrix of $\text{CHAIN}(\mathcal{P})$ must have small nonnegative factorization, and we will now show that this is indeed the case.

The vertices of $\text{CHAIN}(\mathcal{P})$ are indexed by antichains of \mathcal{P} . The facets, apart from the nonnegativities of the variables, are indexed by maximal chains in \mathcal{P} . The facet inequality corresponding to a maximal chain C is of the form $\sum_{a \in C} x_a \leq 1$. As such, given an antichain A and a maximal chain C , the (C, A) -entry of the slack matrix is 0 if C and A intersect, and 1 otherwise.

In Corollary 2.9 we saw that $\text{CHAIN}(\mathcal{P})$ has a lift of the form $(z, x) \in \mathbb{R}^{\mathcal{P} \times \mathcal{P}}$ given by the inequalities

$$\begin{cases} 0 \leq z_a & \forall a \in \mathcal{P}, \\ 0 \leq x_{a_i} - x_{a_j} - z_{a_i} & \forall \text{ cover relations } a_i \succ a_j, \\ 0 \leq 1 - x_a & \forall a \in \mathcal{P}_{\max}, \end{cases}$$

together with the equalities $z_a = x_a$ for $a \in \mathcal{P}_{\min}$. Here we denote by \mathcal{P}_{\min} and \mathcal{P}_{\max} the sets of minimal and, respectively, maximal, elements of \mathcal{P} . We will denote E as the set of all cover relations. The proof of Theorem 3.2 then gives us a way of extracting a factorization.

For every antichain A , we will assign a vector that is simply the preimage of its indicator vector χ_A through the lift, evaluated at each of the facet inequalities. The preimage of χ_A is (χ_A, χ_F) , where F is the filter over A , i.e., the set of all elements greater than or equal to some element of the antichain. Evaluating this at each of the defining inequalities above gives a vector $(\chi_A, y, w) \in \{0, 1\}^{\mathcal{P} \times \mathcal{P}_{\min} \times E \times \mathcal{P}_{\max}}$ with $y_{a \succ b} = 1$ if and only if a is in $F \setminus A$ and $b \notin F$, and $w_a = 1$ if and only if $a \in \mathcal{P}_{\max} \setminus F$.

For example, in Figure 18 we can see an antichain A and its corresponding filter F in red and blue, respectively. We would, in this case, have $\chi_A = (0, 1, 0, 0, 0, 0, 0, 1)$, since the antichain contains only the nodes 4 and 8, while $\chi_F = (0, 1, 0, 1, 0, 1, 1, 1)$. Setting $z = \chi_A$ and $x = \chi_F$ and looking at the inequalities above, the only interesting ones are those coming from the cover relations that will give us the vector y indicated in the figure. This is the set of edges in the Hasse diagram that go from a node not in the filter to a node that is in the filter but not in the antichain.

As for the facets, the nonnegativities of the variables in $\text{CHAIN}(\mathcal{P})$ lift to nonnegativities of the variables z_a in the lift, so we get the vector $(e_a, 0, 0) \in \{0, 1\}^{\mathcal{P} \times E \times \mathcal{P}_{\max}}$, where e_a is 1 only in coordinate a . The inequality for a maximal chain C of the form $a_0 \prec a_1 \prec \dots \prec a_N$ lifts to $1 - \sum_{i=0}^N z_{a_i} \geq 0$, which is the sum of the inequalities




Fig. 18 Example of an antichain (left), the corresponding filter (center), and the vector y of the factorization (right).

corresponding to the maximal element in C and also all the cover relations. So we obtain a vector $(0, \chi_C, e_{a_N}) \in \{0, 1\}^{\mathcal{P} \times E \times \mathcal{P}_{\max}}$, where χ_C is simply the indicator vector of which cover relations are in the chain.

Since we are following the proof of Theorem 3.2, these must necessarily form a factorization. It is not hard to see this independently. For the nonnegativity inequalities this is trivial, so we just have to see that given an antichain A and a chain C the inner product $\langle (\chi_A, y, z), (0, \chi_C, e_{a_N}) \rangle = \langle y, \chi_C \rangle + z_{a_N}$ is 0 if the chain intersects the antichain, and 1 otherwise. But $\langle y, \chi_C \rangle$ is 1 if and only if the chain enters the filter outside of an antichain element, while z_{a_N} is 1 if and only if the chain never enters the filter at all, so they can never be 1 at the same time. If they are both 0, it means that the chain did enter the filter at an element of the antichain, and so we indeed have a factorization. This is a slightly different factorization from the classical one originally given by Yannakakis [95], but of roughly the same size.

3.3. The Slack Operator. Since we are interested in lifts of convex sets in general and not just polytopes, we next need to generalize the notion of slack matrix to a more continuous object, the slack operator of a convex set. To simplify the exposition we will always consider compact convex sets with the origin in the interior and refer to them as *convex bodies*. Given a convex body C , its polar is

$$C^\circ = \{y \in \mathbb{R}^n : \langle x, y \rangle \leq 1 \forall x \in C\},$$

which is again a convex body. Recall that an *extreme point* of a convex set C is a point in C that cannot be written as a convex combination of two distinct points of C , i.e., it is not in the relative interior of any line segment contained in C . We will denote the set of all extreme points of C as $\text{ext}(C)$. Generalizing the notion of extreme point, a *face* of a convex set C is a convex subset $F \subseteq C$ such that if any point in F is in the relative interior of a line segment contained in C , then the entire line segment is contained in F .

Consider again the function $1 - \langle y, x \rangle$ but now with domain $C^\circ \times C$. In this domain, it is a nonnegative function and induces a one-to-one correspondence between points $\tilde{y} \in C^\circ$ and linear inequalities valid over C , simply given by $1 - \langle \tilde{y}, x \rangle \geq 0$. Since $1 - \langle y, x \rangle$ is affine in x and y (separately), and C and C° are compact, its value at any point is completely determined by its value on $\text{ext}(C^\circ) \times \text{ext}(C)$. We call the resulting map the *slack operator* of C .

DEFINITION 3.5. Given a convex body C , its slack operator is the map $s_C : \text{ext}(C^\circ) \times \text{ext}(C)$ defined as $s_C(y, x) = 1 - \langle x, y \rangle$.

When C is a polytope, both C and C° have a finite number of extreme points, and we recover the slack matrix.




Fig. 19 Convex set C (left) and its polar C° (right) with extreme points highlighted.

EXAMPLE 3.6. Let $C = \{(x, y) \in \mathbb{R}^2 : (1+x)^2(x-1) + y^2 \leq 0, x \geq -1\}$. This is a convex set bounded by a cubic curve. Its boundary is parameterized by

$$(1 - v^2, 2v - v^3)$$

with $v \in [-\sqrt{2}, \sqrt{2}]$, and every boundary point is extreme. The polar of C is the convex hull of the cardioid defined by $4x^4 + 32y^4 + 13x^2y^2 + 18xy^2 - 4x^3 - 27y^2 = 0$, as seen in Figure 19. Not every boundary point of C° is extreme. The extreme boundary points are parameterized by

$$\left(\frac{2 - 3u^2}{2 - u^2 + u^4}, \frac{2u}{2 - u^2 + u^4} \right)$$

with $u \in [-\sqrt{2}, \sqrt{2}]$. The slack operator can be parameterized as $s_C : [-\sqrt{2}, \sqrt{2}]^2 \rightarrow \mathbb{R}_+$, where

$$s_C(u, v) = 1 - \frac{(2 - 3u^2)(1 - v^2) + 2u(2v - v^3)}{2 - u^2 + u^4} = \frac{(2 - v^2)(v - u)^2 + (u^2 - v^2)^2}{u^4 + (2 - u^2)}.$$

It is obvious from the last expression that s_C is nonnegative on its domain.

3.4. Cone Factorizations of Slack Operators. We define these general slack operators in the hope that there might be a geometric interpretation to their factorizations. However, to go from polyhedral lifts to more general conic lifts, we still need to generalize the idea of nonnegative factorizations. Before that, we need to quickly recall the definition of the dual of a cone. Given a cone $K \subseteq \mathbb{R}^\ell$, we define its dual to be $K^* = \{y \in \mathbb{R}^\ell : \langle x, y \rangle \geq 0 \forall x \in K\}$. Here we may think of any inner product, but we will always consider the usual Euclidean one. When working with cones of symmetric matrices this amounts to $\langle X, Y \rangle = \text{tr}(XY)$, where tr is the usual matrix trace, also called the trace inner product.

DEFINITION 3.7. Let K be a closed convex cone and K^* its dual. Furthermore, let C be a convex body and s_C its slack operator. A factorization of s_C through K , or K -factorization of C , is a pair of maps $A : \text{ext}(C) \rightarrow K$ and $B : \text{ext}(C^\circ) \rightarrow K^*$ such that $s_C(y, x) = \langle A(x), B(y) \rangle$ for all $x \in \text{ext}(C)$ and $y \in \text{ext}(C^\circ)$.

This allows us to factorize the slack operator of a nonpolyhedral set C through a cone K .

EXAMPLE 3.8. Recall from Example 3.6 that the slack operator of the set

$$C = \{(x, y) \in \mathbb{R}^2 : (1+x)^2(x-1) + y^2 \leq 0, x \geq -1\}$$

is given by the function $s_C(u, v)$ with domain $[-\sqrt{2}, \sqrt{2}]^2$ and equaling

$$1 - \frac{(2 - 3u^2)(1 - v^2) + 2u(2v - v^3)}{2 - u^2 + u^4} = \frac{2u^2 + u^4 - 4uv + 2v^2 - 3u^2v^2 + 2uv^3}{2 - u^2 + u^4}.$$

Here we are identifying the extreme points of both C and C° by the interval $[-\sqrt{2}, \sqrt{2}]$ since we have an explicit parameterization that we can think of as a labeling of the points. We claim that we can factor this operator through the cone \mathcal{S}_+^3 . To see this, define

$$A(v) = \begin{bmatrix} 1 & 0 & 1 - v^2 \\ 0 & 2 - v^2 & v(2 - v^2) \\ 1 - v^2 & v(2 - v^2) & 1 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \\ 1 - v^2 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ 1 - v^2 \end{bmatrix}^\top + (2 - v^2) \begin{bmatrix} 0 \\ 1 \\ v \end{bmatrix} \begin{bmatrix} 0 \\ 1 \\ v \end{bmatrix}^\top$$

and

$$B(u) = \frac{1}{2 - u^2 + u^4} \begin{bmatrix} (u^2 - 1)^2 & -u(u^2 - 1) & u^2 - 1 \\ -u(u^2 - 1) & u^2 & -u \\ u^2 - 1 & -u & 1 \end{bmatrix} = \frac{1}{2 - u^2 + u^4} \begin{bmatrix} u^2 - 1 \\ -u \\ 1 \end{bmatrix} \begin{bmatrix} u^2 - 1 \\ -u \\ 1 \end{bmatrix}^\top.$$

Since these matrices are psd for $u, v \in [-\sqrt{2}, \sqrt{2}]$, and $s_C(u, v) = \langle A(v), B(u) \rangle$, we have the claimed factorization.

As in the Yannakakis setting of polytopes, we can show that the existence of lifts and the existence of factorizations are equivalent even in this generalized setting of arbitrary convex bodies and closed convex cones. Recall that a convex set C has a K -lift, into a closed convex cone K , if we can write it as $C = \pi(K \cap L)$, where π is a linear map and L an affine subspace. We will also say that C has a *proper* K -lift if L intersects the relative interior of K . Note that if C has a K -lift, it always has a proper lift to some face of K , so the assumption of properness is not a very strong one.

THEOREM 3.9 ([37]). *If a convex body C has a proper K -lift, then its slack operator, s_C , has a K -factorization. Reciprocally, if s_C has a K -factorization, then C has a K -lift.*

Proof. The proof is just a reworking of the proof of Theorem 3.2.

Let A and B form a K -factorization of the slack operator s_C . Then one can check that

$$C = \{x \in \mathbb{R}^n : \exists X \in K \text{ s.t. } \langle X, B(y) \rangle = s_C(y, x) \forall y \in \text{ext}(C^\circ)\},$$

since the equalities force $s_C(y, x) \geq 0$ for all $y \in \text{ext}(C^\circ)$, guaranteeing that the right-hand side is contained in C , while plugging in $X = A(x)$ shows that $\text{ext}(C)$, hence C , is contained on the right-hand side. Although there are potentially infinitely many linear equalities, they necessarily cut out a finite-dimensional space, and we can use them to eliminate the variable x and explicitly get C as a projection of a slice of K .

Suppose now C has a K -lift, i.e., $C = \pi(K \cap L)$ for some linear map π and affine space L . For x in $\text{ext}(C)$, define $A(x)$ to be any element in the preimage $\pi^{-1}(x) \cap (K \cap L)$. For $y \in \text{ext}(C^\circ)$, one can pull back the inequality $1 - \langle x, y \rangle$ by π to a valid inequality on $K \cap L$. The hypothesis that L intersects the interior of K (Slater's condition) can be shown to guarantee that any valid inequality on $K \cap L$, when restricted to L , equals $\langle Y, - \rangle$ for some $Y \in K^*$. Pick such a Y as $B(y)$. Then one can check that $\langle A(x), B(y) \rangle = s_C(y, x)$, as intended. \square

The cones K that we will mostly be interested in are nonnegative orthants and psd cones. In those cases, their special structure allows us to omit the properness condition from the statement of Theorem 3.9 (see [37, Corollary 1]).

EXAMPLE 3.10. *Let us consider again the convex set C given in Example 3.8. Since the slack operator is factorizable through \mathcal{S}_+^3 , Theorem 3.9 implies that C must have an \mathcal{S}_+^3 -lift. In fact,*

$$C = \left\{ (x, y) : \begin{bmatrix} 1 & 0 & x \\ 0 & 1+x & y \\ x & y & 1 \end{bmatrix} \succeq 0 \right\}.$$

Indeed, check that the nonnegativity of the principal minors of the above matrix gives the polynomial inequalities defining C . Moreover, the symmetry of the existence of factorizations and the self-duality of the psd cone implies that the polar C° must also have an \mathcal{S}_+^3 -lift. One can show that

$$C^\circ = \left\{ (z, w) : \exists z_1, z_2, z_3 \in \mathbb{R} \text{ s.t. } \begin{bmatrix} 1 - z_3 - z_1 & z_2 & -(z + z_3)/2 \\ z_2 & z_3 & -w/2 \\ -(z + z_3)/2 & -w/2 & z_1 \end{bmatrix} \succeq 0 \right\}.$$

An important basic property of slack operators is that for any convex body C we have $s_C(y, x) = s_{C^\circ}(x, y)$; in other words, the slack operator of C° is the transpose of the slack operator of C . This follows directly from the definition, since $(C^\circ)^\circ = C$ for convex bodies. Furthermore, it also follows directly from the definition that s_C has a K -factorization if and only if its transpose has a K^* -factorization. An immediate consequence of Theorem 3.9 is that C has a K -lift if and only if C° has a K^* -lift.

Moreover, note that if a convex body C can be written as a linear image of some convex set Q , i.e., $C = \pi(Q)$, then $C^\circ \cong Q^\circ \cap \ker(\pi)^\perp$, which means that we can write C° as a linear slice of the polar of its lift. In the case of polytopes, this has a nice interpretation. We saw that if P has an \mathbb{R}_+^m -lift, then P° also has an \mathbb{R}_+^m -lift. This means that if P is the image of a polytope with m facets, then P° is a slice of a polytope with m vertices.

EXAMPLE 3.11. *We have seen in Figure 4 that the regular octagon is the projection of a polytope with 6 facets. Since the regular octagon is self-polar, we then find that it must also be a section of a polytope with 6 vertices. Figure 20 shows the regular octagon as a slice of an octahedron.*




Fig. 20 The regular octagon as a minimal linear slice of an octahedron.

One should also point out that there is a simple extension of Theorem 3.9 such that it covers approximations of sets, instead of only exact representations. To state it we need to introduce the *generalized slack operator*. Given two convex bodies $C \subseteq D$, we define the generalized slack operator of the pair (C, D) to be $s_{C,D} : \text{ext}(D^\circ) \times \text{ext}(C) \rightarrow \mathbb{R}_+$ given by $s_{C,D}(y, x) = 1 - \langle x, y \rangle$. A look at the proof of Theorem 3.9 shows that it immediately generalizes to the following result.

PROPOSITION 3.12. *Let C and D be a pair of convex bodies with $C \subseteq D$. If there is a convex set B such that B has a proper K -lift and $C \subseteq B \subseteq D$, then $s_{C,D}$ has a K -factorization. Reciprocally, if $s_{C,D}$ has a K -factorization, then there exists a convex set B such that B has a K -lift and $C \subseteq B \subseteq D$.*

By setting $D = (1 + \varepsilon)C$ we obtain a powerful tool to aid the study of approximability of a given set by sets with small lifts.

3.5. Bregman Divergence as a Slack Operator. An important notion that is closely related to that of a slack operator is the *Bregman divergence* of a convex function, which was briefly mentioned in section 2.3 and is widely used in applications. The Bregman divergence can, under some mild assumptions, be seen as the slack operator of the epigraph of the corresponding convex function, while slack operators can be seen as special cases of a nonsmooth version of the Bregman divergence, tying the two notions together.

Following [3], let $f : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ be a closed convex proper function whose domain has nonempty interior. If f is differentiable in the interior of its domain, we define its Bregman divergence as the function $D_f : \mathbb{R}^n \times \text{int}(\text{dom}(f)) \rightarrow [0, +\infty)$ given by

$$D_f(x, y) = f(x) - f(y) - \langle \nabla f(y), x - y \rangle.$$

In other words, $D_f(x, y)$ is the difference between $f(x)$ and the value at x of the linear approximation to f around y . The convexity of f implies that $D_f(x, y) \geq 0$.

EXAMPLE 3.13. *If we take $f(x) = x^\top Qx$ to be a convex quadratic function, then*

$$D_f(x, y) = x^\top Qx - y^\top Qy - 2y^\top Q(x - y) = (x - y)^\top Q(x - y),$$

since Q is a symmetric matrix. In particular, if $f(x) = \|x\|^2$, we have $D_f(x, y) = \|x - y\|^2$.

It turns out that the Bregman divergence is a slack operator of sorts for convex functions, where Fenchel duality replaces polarity of convex sets. Recall that the Fenchel conjugate of f is the function $f^* : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ defined by

$$f^*(y^*) = \sup\{\langle x, y^* \rangle - f(x) \mid x \in \mathbb{R}^n\}.$$

The function f^* is closed and convex and *Fenchel's inequality* says that $f(x) + f^*(y^*) \geq \langle x, y^* \rangle$. We will refer to $F(x, y^*) := f(x) + f^*(y^*) - \langle x, y^* \rangle$ as the *slack* in the Fenchel inequality. Using simple properties of Fenchel conjugation one gets that $D_f(x, y) = F(x, \nabla f(y))$. As such, the Bregman divergence is essentially a slack, but of duality of convex functions rather than convex sets.

In fact, if f is well behaved, namely, if f is a *convex function of Legendre type* (smooth and strictly convex in its domain, and the limit of $\|\nabla f(x)\|$ as x approaches the boundary of the domain through its interior is $+\infty$), then the map $y \rightarrow y^* := \nabla f(y)$ is an isomorphism from the interior of the domain of f to the interior of the domain of f^* . In that case, the Bregman divergence and the slack of the Fenchel inequality are just reparameterizations of each other. For more details, see [79, section 26].

EXAMPLE 3.14. The function $f(x) = e^x$ is a strictly convex smooth function with domain \mathbb{R} . Its conjugate f^* has domain in the nonnegative reals and is given by $f^*(0) = 0$ and $f^*(z) = z \ln(z) - z$ for $z > 0$. The Fenchel inequality becomes $F(x, z) = e^x + z \ln(z) - z - xz \geq 0$. The Bregman divergence is

$$D_f(x, y) = F(x, f'(y)) = F(x, e^y) = e^x - e^y + (y - x)e^y.$$

As above, we can think of the function F as the map $x_0 + y_0 - \langle x, y \rangle$ with domain the product of the graph of f and the graph of f^* . When extended to the epigraphs, this operator creates a one-to-one correspondence between points in the epigraph of a function and affine underestimators of the other function, as illustrated in Figure 21.




Fig. 21 The epigraphs of $f(x) = e^x$ and $f^*(z) = z \ln(z) - z$, with some linear underestimators of f and corresponding points in the epigraph of f^* marked.

For instance, the point $(2, 0)$ in $\text{epi}(f^*)$ gives rise to the inequality $x_0 - 2x \geq 0$, while the point $(1, -1)$, which is an extreme point of $\text{epi}(f^*)$, gives rise to the inequality $x_0 - 1 - x \geq 0$, and so on.

This illustrates that for a smooth convex function of Legendre type, its Bregman divergence is, in a very strong sense, the same as the slack operator of its epigraph. The only difference is that $\text{epi}(f)$ and $\text{epi}(f^*)$ are not actually polar to each other, having instead a different duality relation.

Interestingly, not only can we think of the Bregman divergence as a slack operator, but we can also conversely think of slack operators as straightforward generalizations of the Bregman divergence. For any closed convex set C , consider its *indicator function* $\iota_C(x)$ given by $\iota_C(x) = 0$ if $x \in C$, and $+\infty$ otherwise. This is a proper closed convex function, and its Fenchel conjugate is the *support function* of C given by $h_C(y) = \sup_{x \in C} \langle x, y \rangle$. Since $\iota_C(x)$ is not a smooth convex function, its usual Bregman divergence is not defined. However, there are many generalizations of the Bregman divergence to the nonsmooth case. A simple way of doing this is to think of Bregman divergence as the affine operator $x_0 + y_0 - \langle x, y \rangle$ applied to $\text{epi}(\iota_C) \times \text{epi}(h_C)$. We still have a correspondence between points in the epigraph of h_C and valid inequalities on the epigraph of ι_C . Restricting to the extreme points of $\text{epi}(\iota_C)$, which correspond to the extreme points of C , and a representative of every extreme ray of the cone $\text{epi}(h_C)$, recovers the slack operator of C (up to scaling). An advantage of this viewpoint is that we can remove the technical condition of requiring 0 to be in the interior of C .

4. Constructions of Lifts. Is there a *systematic* way to construct lifts for convex bodies? In this section we will try to answer this question, informed by the factorization theorem from the previous section.

4.1. Spectrahedral Lifts and Sums of Squares. We start by restating the factorization theorem (Theorem 3.9) in the special case of spectrahedral lifts, which will be our main focus in this section.

THEOREM 4.1. *Let $C = \text{conv}(\mathcal{X})$ be a convex body. Then C has a spectrahedral lift of size m if and only if there exists a map $A : \mathcal{X} \rightarrow \mathcal{S}_+^m$ such that the following holds: for any $\ell \in C^\circ$, there exists a $B \in \mathcal{S}_+^m$ such that*

$$(4.1) \quad 1 - \ell(x) = \text{tr}(A(x)B) \quad \forall x \in \mathcal{X}.$$

REMARK 4.2. *Throughout this section we use the notation ℓ for elements of C° to emphasize that these are linear forms defined on the ambient space where C lives. We also write $\ell(x)$ (as opposed to the bracket notation $\langle \ell, x \rangle$) for the evaluation of ℓ at some element x .*

Recall that the slack operator of C is indexed precisely by pairs $(\ell, x) \in \text{ext}(C^\circ) \times \text{ext}(C)$ and the associated entry in the matrix is $1 - \ell(x) \geq 0$. Equation (4.1) is nothing but an \mathcal{S}_+^m factorization of this slack operator.

It is useful to think of (4.1) as a *certificate of nonnegativity* of the affine function $1 - \ell(x)$ on the set \mathcal{X} . Indeed, the right-hand side of (4.1) is “obviously” nonnegative as it is the inner product of two psd matrices. That (4.1) is a certificate of nonnegativity becomes clearer if we factorize $A(x)$ as $A(x) = F(x)F(x)^\top$ and $B = DD^\top$. (Such a factorization is possible since $A(x), B \succeq 0$.) Then we have, for any $x \in \mathcal{X}$,

$$(4.2) \quad 1 - \ell(x) = \text{tr}(A(x)B) = \text{tr}((F(x)^\top D)^\top F(x)^\top D) = \sum_{1 \leq i, j \leq m} M_{ij}(x)^2,$$

where $M(x) = F(x)^\top D$. Equation (4.2) shows that $1 - \ell(x)$ is a *sum of squares* of the entries of $M(x)$. In turn, the entries of $M(x)$ are linear combinations of the entries of $F(x)$. We have thus shown that if $C = \text{conv}(\mathcal{X})$ has a spectrahedral lift of size m , then there is a *subspace* V of functions on \mathcal{X} with $\dim(V) \leq m^2$ such that the following is true:

$$(*\text{-SDP}) \quad \begin{array}{l} \text{For any valid linear inequality } \ell(x) \leq 1 \text{ on } \mathcal{X}, \\ \text{there exist functions } h_k \in V \text{ s.t. } 1 - \ell(x) = \sum_k h_k(x)^2 \quad \forall x \in \mathcal{X}. \end{array}$$

The subspace V here is precisely $V = \text{span}\{F_{ij} : 1 \leq i, j \leq m\}$, where we think of each F_{ij} as a function on \mathcal{X} . (We will look at the nature of these functions later in this section.)

One can also prove a converse of the previous statement. This is the object of the next proposition, which shows how to construct a lift from a subspace satisfying (*-SDP). Note that the statement below is *not an exact* converse, as it gives a spectrahedral lift of size $\dim(V)$ instead of $\sqrt{\dim(V)}$. The reason for this concerns the rank of the map $A(x)$, which is rank-one in the construction (4.3) below, whereas it can be of higher rank in Theorem 4.1.

PROPOSITION 4.3. *Let $C = \text{conv}(\mathcal{X})$ be a convex body. Assume there exists a subspace V of functions on \mathcal{X} such that (*-SDP) is true. Then C has a spectrahedral lift of size $\dim(V)$.*

Proof. We construct a psd factorization of the slack operator of C . Let $m = \dim(V)$, f_1, \dots, f_m , be a basis of V and $F(x) = [f_1(x), \dots, f_m(x)]^\top$. Define

$$(4.3) \quad A(x) = F(x)F(x)^\top \in \mathcal{S}_+^m.$$

Let $\ell \in C^\circ$ so that $1 - \ell(x) \geq 0$ for all $x \in \mathcal{X}$. We want to show that there exists $B \in \mathcal{S}_+^m$ such that $1 - \ell(x) = \text{tr}(A(x)B)$ for all $x \in \mathcal{X}$. We know by (*-SDP) that there exist functions $h_k \in V$ such that $1 - \ell(x) = \sum_k h_k(x)^2$ for $x \in \mathcal{X}$. Since f_1, \dots, f_m form a basis of V , we know there exist vectors $b_k \in \mathbb{R}^m$ such that $h_k(x) = b_k^\top F(x)$. Then

$$1 - \ell(x) = \sum_k h_k(x)^2 = \sum_k (b_k^\top F(x))^2 = \text{tr}(F(x)F(x)^\top B),$$

where $B = \sum_k b_k b_k^\top \in \mathcal{S}_+^m$. This completes the proof. \square

Proposition 4.3 reduces the problem of constructing a spectrahedral lift of C to that of finding a subspace V of functions such that (*-SDP) holds. We now list some examples of lifts constructed in this way.

EXAMPLE 4.4 (elliptope). Consider the convex body $C = [-1, 1]^2$ with extreme points $\mathcal{X} = \{-1, +1\}^2$. Define the subspace $V = \text{span}(1, x, y)$ of affine functions, and consider the valid linear inequality $x \leq 1$ on C . Note that we can write

$$1 - x = \frac{1}{2}(1 - x)^2 \quad \forall (x, y) \in \mathcal{X},$$

where we have used the fact that $x^2 = 1$. Similarly, one can check that the other facet inequalities of C (namely, $1 + x \geq 0$ and $1 \pm y \geq 0$) are sums of squares from functions in V . It thus follows that C has a spectrahedral lift of size $\dim(V) = 3$. Indeed one can verify that

$$(4.4) \quad C = \left\{ (x, y) \in \mathbb{R}^2 : \exists z \in \mathbb{R}, \begin{bmatrix} 1 & x & y \\ x & 1 & z \\ y & z & 1 \end{bmatrix} \succeq 0 \right\}.$$

This is precisely the example of the three-dimensional elliptope seen in the introduction (see Figure 5(left)), which projects onto the two-dimensional square.

EXAMPLE 4.5 (convex hull of rational curves). Consider a plane curve $\mathcal{X} \subset \mathbb{R}^2$ having a rational parameterization $\mathcal{X} = \{(x(t), y(t)) : t \in \mathbb{R}\}$, where $x(t), y(t)$ are rational functions, i.e., $x(t) = a(t)/q(t)$, $y(t) = b(t)/q(t)$, where we assume that $q(t) > 0$ for all $t \in \mathbb{R}$ and $\max(\deg a, \deg b) \leq \deg q$. We can show that $\text{cl conv}(\mathcal{X})$ (where cl denotes the closure) has a spectrahedral lift of size at most $\deg(q)/2 + 1$ [77, 52]; we will show this by exhibiting a subspace of functions V such that property (*-SDP) holds for \mathcal{X} .

We first define a function $t : \mathcal{X} \rightarrow \mathbb{R}$ that gives the “time” coordinate of any point in \mathcal{X} in the parameterization $(x(t), y(t))$; namely, for any $(x_0, y_0) \in \mathcal{X}$ let $t(x_0, y_0) = t_0 \in \mathbb{R}$ be such that $(x(t_0), y(t_0)) = (x_0, y_0)$ (if multiple such t_0 exist, an arbitrary choice is made so that t is single-valued). Consider now the subspace V of bivariate functions of the form

$$(4.5) \quad h(x, y) = \sum_{i=0}^{\deg(q)/2} c_i \frac{t(x, y)^i}{\sqrt{q(t(x, y))}},$$

where the c_i are arbitrary coefficients. Note that $\dim(V) = 1 + \deg(q)/2$.

We claim that any valid linear inequality on $\text{cl conv}(\mathcal{X})$ is a sum of squares from V . Indeed, assume we have a linear form ℓ such that $1 - \ell(x)$ is nonnegative on \mathcal{X} ,

i.e., $1 - \ell(x(t), y(t)) \geq 0$ for all $t \in \mathbb{R}$. Using the fact that $x(t) = a(t)/q(t)$ and $y(t) = b(t)/q(t)$, this gives

$$q(t) - \ell(a(t), b(t)) \geq 0 \quad \forall t \in \mathbb{R}.$$

Note that $q(t) - \ell(a(t), b(t))$ is a (univariate) polynomial of degree at most $\deg q$. It is a well-known simple fact that univariate polynomials are nonnegative in the real line if and only if they are sums of squares (see Chapter 3 of [7] for an overview on sums of squares); therefore $q(t) - \ell(a(t), b(t))$ must be a sum of squares of some polynomials $h_k(t) \in \mathbb{R}[t]$ of degree at most $(\deg q)/2$, i.e.,

$$q(t) - \ell(a(t), b(t)) = \sum_k h_k(t)^2.$$

Evaluating at $t = t(x, y)$ and dividing by q we get

$$1 - \ell(x, y) = \sum_k \left(\frac{h_k(t(x, y))}{\sqrt{q(t(x, y))}} \right)^2 \quad \forall (x, y) \in \mathcal{X}.$$

Since the functions $(x, y) \mapsto h_k(x, y)/\sqrt{q(t(x, y))}$ belong to V , we have shown what we wanted.

4.2. Hierarchies. Consider the case where C is the convex hull of a set defined using polynomial equations. A natural choice of subspace V in this setting is the space $V = V_k$ of *polynomials* of degree at most k , where k is a certain fixed integer. If property (*-SDP) holds with such V_k , then we get a spectrahedral lift of $\text{conv}(\mathcal{X})$ of size $\dim(V_k)$. This spectrahedral lift is precisely the lift produced by the *Lasserre/theta-body method*; see [62, 36].

We note that it is possible for a convex body C to have a spectrahedral lift and yet not verify property (*-SDP) for any subspace of polynomials V (of arbitrary degree). For example, consider the curve defined by the equation $y^2 - x^3 + x^4 = 0$ and shown in Figure 22; this curve is known as the *Piriform curve*. Because of the singularity at the origin, one can show that the Lasserre/theta body hierarchy for this curve is not exact at any finite level k ; see [40, Theorem 9]. However, the convex hull of this curve has an SDP representation since the curve has a rational parameterization; namely, it can be shown that it is parameterized by

$$\begin{cases} x(\theta) = \frac{1+\sin(\theta)}{2}, \\ y(\theta) = \frac{\cos(\theta)(1+\sin(\theta))}{4}, \end{cases}$$

which can be turned into a rational parameterization using $\cos(\theta) = \frac{1-t^2}{1+t^2}$ and $\sin(\theta) = \frac{2t}{1+t^2}$.

EXAMPLE 4.6 (*k-level polytopes*). A polytope P is called *k-level* if every facet-defining linear function of P takes at most k different values on the vertices of the polytope. Put differently, for every affine hyperplane H defining a facet of P , all the vertices of P lie in at most k different translates of H . See Figure 23 for examples of 2- and 3-level polytopes in \mathbb{R}^3 . One important example of a 2-level polytope is the *stable set polytope of perfect graphs* (see section 2.2, Theorem 2.11).

One can prove that if P is a *k-level polytope*, then P admits a spectrahedral lift of size at most $\binom{n+k-1}{k-1} = \dim(V_{k-1})$, where V_{k-1} is the space of polynomials of degree at most $k-1$ in n variables. To do this we will prove that (*-SDP) holds with $V = V_{k-1}$.




Fig. 22 The Piriform curve defined by $y^2 - x^3 + x^4 = 0$. This is a rational curve and so the convex hull of this curve has a spectrahedral lift. However, property (*-SDP) is not true for any subspace V consisting of polynomials. (Note that functions of the form (4.5) are not polynomials.) This example was considered in [35, Example 4.4].




Fig. 23 Left: Examples of 2-level polytopes. Right: A polytope that is 3-level but not 2-level. Figure from [40].

Indeed, let \mathcal{X} be the set of vertices of the polytope and assume that $1 - \ell(x) \geq 0$ for all $x \in \mathcal{X}$. We need to show that there exist $h_i \in V$ such that $1 - \ell(x) = \sum_i h_i(x)^2$ for all $x \in \mathcal{X}$. Note that by Farkas' Lemma, $1 - \ell(x)$ is a nonnegative combination of facet-defining inequalities, so it is enough to show that any facet inequality can be written as a sum of squares. We can therefore assume that $1 - \ell(x) \geq 0$ defines a facet which, since P is k -level, implies that $\{1 - \ell(x) : x \in \mathcal{X}\} = \{a_1, a_2, \dots, a_k\}$ for some nonnegative numbers $a_i \in \mathbb{R}$. Consider a univariate polynomial $r \in \mathbb{R}[t]$ of degree $k-1$ such that $r(a_i) = \sqrt{a_i}$ (e.g., constructed using the Lagrange interpolation formula), and let $h(x) = r(1 - \ell(x))$ be the evaluation of r at $1 - \ell(x)$. Since $\deg r \leq k-1$, we have $h \in V_{k-1}$. Note that for any $x \in \mathcal{X}$ we have $1 - \ell(x) = r(1 - \ell(x))^2 = h(x)^2$, where we used the fact that $1 - \ell(x) \in \{a_1, \dots, a_k\}$ and that $r(a_i)^2 = a_i$ for all $i = 1, \dots, k$. We have thus shown that this choice of subspace works. An illustration of this construction is shown in Figure 24 in the case of the regular hexagon.

4.3. Symmetric Convex Bodies. In many cases, the convex body $C \subset \mathbb{R}^n$ we are interested in has some symmetries. This means that there are certain linear invertible transformations of \mathbb{R}^n that leave C invariant. The symmetry group (or *automorphism group*) of C is precisely the set of such transformations:

$$\text{Aut}(C) = \{T \in GL(\mathbb{R}^n) : T(C) = C\}.$$

For example, the automorphism group of the regular n -gon in the plane is the dihedral group and consists of $2n$ elements. When constructing a lift of the convex body C , one




Fig. 24 Left: The convex hull of the regular hexagon is a 3-level polytope. The levels of the facet inequality $\cos(\pi/6) - x \geq 0$ are $\{0, a_1, a_2\} = \{0, \cos(\pi/6), 2 \cos(\pi/6)\}$. Right: The polynomial $r(t)$ of degree 2 such that $r(0) = 0$, $r(a_1) = \sqrt{a_1}$, $r(a_2) = \sqrt{a_2}$. Any facet inequality of the regular hexagon is a sum of squares of quadratic polynomials. One can show that this gives a spectrahedral lift of the regular hexagon of size 5.

may require that the lift “respects” the symmetry of C . Such lifts are called symmetric or *equivariant* lifts and have been considered since the work of Yannakakis [95]. Having an equivariant spectrahedral lift of a convex body C has a very natural interpretation in terms of the subspace V of functions in $(\ast\text{-SDP})$. Finding an equivariant lift of $C = \text{conv}(\mathcal{X})$ amounts to finding a subspace V of functions on \mathcal{X} that is *invariant* under the action of $\text{Aut}(C)$. Here the action of $\text{Aut}(C)$ on functions $\mathcal{X} \rightarrow \mathbb{R}$ is the natural one defined by $(g \cdot f)(x) = f(g^{-1}x)$ for $g \in \text{Aut}(C)$ and $x \in \mathcal{X}$. A question the reader might ask is whether there is benefit in *breaking symmetry* when constructing lifts of convex bodies. More precisely, given a symmetric convex body, can nonequivariant lifts be much smaller than their equivariant counterparts? It turns out that the answer is yes, and the gap can be very large. Kaibel, Pashkovich, and Theis [57] showed the existence of polytopes (P_n) that have a nonequivariant polyhedral lift of polynomial size in n , and yet any equivariant lift has size at least $n^{\Omega(\log n)}$. Also, it was shown in [24] that the even parity polytope has no polynomial-size equivariant spectrahedral lift, despite having a linear size nonequivariant one (cf. section 2.1).

4.4. Linear Programming Lifts. We have so far only discussed the case of spectrahedral lifts. In this concluding paragraph, we briefly discuss polyhedral lifts. It can be shown via the slack matrix factorization theorem that, for a finite set \mathcal{X} , a polytope $P = \text{conv}(\mathcal{X})$ has a polyhedral lift of size m if and only if there are m nonnegative functions $a_1, \dots, a_m : \mathcal{X} \rightarrow \mathbb{R}_+$ such that the following is true:

$$(\ast\text{-LP}) \quad \begin{array}{l} \text{For any valid linear inequality } \ell(x) \leq 1 \text{ on } \mathcal{X}, \\ \text{there exist } b_1, \dots, b_m \geq 0 \text{ s.t. } 1 - \ell(x) = \sum_{i=1}^m b_i a_i(x) \quad \forall x \in \mathcal{X}. \end{array}$$

As such, whereas the task of constructing a spectrahedral lift corresponds to finding a subspace V such that $(\ast\text{-SDP})$ holds, constructing an LP lift requires finding a finite set of functions $a_1, \dots, a_m : \mathcal{X} \rightarrow \mathbb{R}_+$ such that $(\ast\text{-LP})$ is true. There are some settings in which a natural set of functions can be considered. Assume that $\mathcal{X} \subset \{0, 1\}^n$ is a subset of the Boolean hypercube and is described using polynomial inequalities, i.e.,

$$(4.6) \quad \mathcal{X} = \{x \in \{0, 1\}^n : g_1(x) \geq 0, \dots, g_k(x) \geq 0\},$$

where g_1, \dots, g_k are polynomials in x . Such a setting covers many combinatorial optimization problems. In this case one can consider the family of functions

$$(4.7) \quad a_{T,I,\gamma}(x) = \prod_{i \in T} x_i \prod_{i \in I \setminus T} (1 - x_i) \prod_{i=1}^k g_i(x)^{\gamma_i},$$

where $T \subseteq I \subset \{1, \dots, n\}$ and $(\gamma_1, \dots, \gamma_k) \in \mathbb{N}^k$. The functions (4.7) are clearly nonnegative on \mathcal{X} . Many of the early *lift-and-project methods* in combinatorial optimization (Sherali and Adams [84] or Lovász and Schrijver [67]) can be understood in this framework and correspond to choosing functions a_i of the form (4.7).

5. Obstructions and Lower Bounds. Given a convex body C and a closed convex cone K , how can we show that C does not have a K -lift? In other words, what properties of C are obstructions to C having a K -lift? When K comes from a cone family, such as \mathbb{R}_+^m or \mathcal{S}_+^m , we typically phrase things in more quantitative language, aiming to find *lower bounds on the size* of the corresponding lift.

For polytopes, the study of lower bounds on the sizes of polyhedral and spectrahedral lifts has received considerable attention in recent years. (See section 6 for more information and references.) In this section, we focus on cases where C is not a polytope. In this more general setting, even deciding whether C has a K -lift for some K in a given cone family (such as psd cones of any finite size) can be very challenging.

The characterization of the existence of lifts in terms of factorizations of the slack operator, discussed in section 3, is a key tool in finding obstructions. This connection allows us to conclude the nonexistence of K -lifts by showing that it is impossible to find a K -factorization of the slack operator. Crucially, we can often find such obstructions by considering only very coarse features of the slack operator. For instance, we may consider only the pattern of zeros and nonzeros of the slack operator, capturing only which extreme points lie on which supporting hyperplanes.

As a concrete example, consider the following result of Goemans about the size of polyhedral lifts of polytopes. We use this as a starting point when discussing obstructions to lifts of convex bodies that are not polytopes throughout this section.

PROPOSITION 5.1 ([31]). *Assume C is a polytope with v extreme points. Then any polyhedral lift of C has size at least $\lceil \log_2 v \rceil$.*

The bound of Goemans can be tight (up to a constant multiplicative factor). One family of examples showing this is that of the regular N -gons in the plane. These have N extreme points and are known to have polyhedral lifts of size either $2\lceil \log_2 N \rceil$ or $2\lceil \log_2 N \rceil - 1$ depending on the value of N [89]. These polyhedral lifts are improved versions of a lift of the regular 2^n -gon of size $2n+4$ due to Ben-Tal and Nemirovski [5].

Another family of examples showing that Goemans' bound is tight comes from the permutahedra Π_n . Recall from Example 1.2 that for each n , Π_n is the convex hull of all permutations of $(1, 2, \dots, n)$. As such, Π_n has $n!$ extreme points. Goemans' lower bound therefore tells us that any polyhedral lift of Π_n has size at least $n \log_2(n/e)$. By comparison, the Birkhoff polytope lift of the permutahedron from Example 1.2 has size $O(n^2)$. Goemans constructed a polyhedral lift of the permutahedron of size $O(n \log(n))$ and used Proposition 5.1 to show that this lift cannot be substantially reduced in size [31].

Goemans' lower bound also implies the intuitively obvious fact that a convex body with infinitely many extreme points cannot have a polyhedral lift. In other words, having infinitely many extreme points is an obstruction to the existence of a

polyhedral lift. To see why, we can consider the contrapositive of Proposition 5.1. This tells us that if a convex body has a polyhedral lift with at most m facets, then it has at most 2^m extreme points.

In the rest of this section, we consider generalizations of Proposition 5.1 in two directions. In section 5.1, we view Goemans' bound as relating to obstructions to lifts arising from the complexity of the facial structure (in this case, the number of vertices) of C . In particular, we discuss other obstructions based on facial structure that are applicable even when C is not a polytope. In terms of the slack operator, focusing on facial structure essentially corresponds to considering only its pattern of zeros and nonzeros. In section 5.2, we view Goemans' bound as relating to obstructions to lifts arising from the algebraic complexity of the boundary of C° . In the case of Goemans' bound, the number of vertices of C is the number of facets of C° , which is also the smallest degree of a polynomial vanishing on the boundary of C° . In particular, in section 5.2 we discuss other obstructions of an algebraic nature.

5.1. Obstructions Based on Facial Structure. The faces (see section 3 for the definition) of a convex set are partially ordered by inclusion. Throughout this section we let \mathcal{F}_C denote the poset of faces of C . If $C = \pi(Q)$ is the projection of a convex set Q , then the preimage under π of any face of C is a face of Q . This observation can be used to find lower bounds on the size of lifts.

EXAMPLE 5.2. Consider again the regular octagon for which we see two polyhedral lifts with six facets in Figure 4. In each case, the preimage (under the projection) of each vertex of the octagon is a face of the lift. Moreover, a vertex of the lift can belong to at most one such preimage face. This implies that any polyhedral lift of an octagon must have at least eight vertices. Since any polytope with five or fewer facets has at most six vertices, it follows that any polyhedral lift of an octagon has size at least six.

In general, if Q is a lift of C , then there is an *order embedding* $\phi : \mathcal{F}_C \rightarrow \mathcal{F}_Q$, i.e., a map that satisfies $F \subseteq F'$ if and only if $\phi(F) \subseteq \phi(F')$ for any $F, F' \in \mathcal{F}_C$. Furthermore, if $Q = K \cap L$ for some closed convex cone K and affine space L , then there is a natural order embedding $\psi : \mathcal{F}_Q \rightarrow \mathcal{F}_K$ that sends a face F of Q to the minimal face of K containing F . Overall, then, the composition $\psi \circ \phi : \mathcal{F}_C \rightarrow \mathcal{F}_K$ gives an order embedding from the face poset of C to the face poset of K . For the lift on the right of Figure 4, the embeddings ϕ and ψ , and their composition, are all illustrated in Figure 25.

Since lifts induce embeddings of face posets, we can find obstructions to the existence of K -lifts by finding obstructions to embedding the face poset of C in the face poset of K .

5.1.1. Obstructions from Counting Faces. If the face poset of the cone K is finite, then we can use the order embedding together with a counting argument to obtain obstructions to the existence of K -lifts. This approach makes sense in the case of polyhedral lifts, where $K = \mathbb{R}_+^m$. The face lattice of \mathbb{R}_+^m is the *Boolean lattice of order m* . A counting argument gives a slight generalization of the bound of Goemans.

LEMMA 5.3. *If C has an \mathbb{R}_+^m -lift, then $m \geq \lceil \log_2(|\mathcal{F}_C|) \rceil$.*

Proof. If C has an \mathbb{R}_+^m -lift, then $|\mathcal{F}_C| = |(\psi \circ \phi)(\mathcal{F}_C)| \leq 2^m$ since $\psi \circ \phi$ is injective and the cardinality of the Boolean lattice of order m is 2^m . \square

5.1.2. Obstructions from Chains of Faces. Any order embedding maps a chain of a certain length to a chain of the same length, giving the following basic obstruction to the existence of lifts.




Fig. 25 Poset embeddings implied by the lift $P = \pi(Q) = \pi(\mathbb{R}_+^6 \cap L)$ of the regular octagon P shown on the right in Figure 4. Left: Hasse diagram of \mathcal{F}_P . Center: Hasse diagram of \mathcal{F}_Q with the white nodes showing the embedding of \mathcal{F}_P . Right: Hasse diagram of $\mathcal{F}_{\mathbb{R}_+^6}$ with gray nodes showing the embedding of \mathcal{F}_Q and white nodes showing the embedding of \mathcal{F}_P .

LEMMA 5.4. *If C is a convex body (of dimension at least one) and C has a K -lift for a closed convex cone K , then the longest chain of faces in C is strictly smaller than the longest chain of faces in K .*

Proof. Since C is a convex body, it is compact. Since C has dimension at least one and is compact, it is not a cone. As such, if $C = \pi(K \cap L)$, then $0 \notin L$. Therefore, any chain of faces in C is mapped, by the order embedding ϕ , to a chain of faces of the same length in K that does not include the face $\{0\}$. In particular, a longest chain of faces in C is mapped to a chain of faces in K that is not of maximum length. \square

Since the empty set is always a face of a convex set, we focus the discussion on chains of nonempty faces.

EXAMPLE 5.5 (convex hull of a cardioid). *Consider the convex body*

$$C = \text{conv}\{(x, y) \in \mathbb{R}^2 : 4x^4 + 32y^4 + 13x^2y^2 + 18xy^2 - 4x^3 - 27y^2 = 0\},$$

the convex hull of the cardioid shown on the right in Figure 19. In Example 3.10 we saw that this convex body has an S_+^3 -lift.

The convex hull of a cardioid has a length-three chain of nonempty faces given by

$$\{(-1, \frac{1}{\sqrt{2}})\} \subseteq [(-1, -\frac{1}{\sqrt{2}}), (-1, \frac{1}{\sqrt{2}})] \subseteq C.$$

From Lemma 5.4, we know that if the convex hull of a cardioid is to have a K -lift, then K must have a chain of nonempty faces of length at least four. Since S_+^2 does not have a chain of nonempty faces of length four, we can conclude that the convex hull of a cardioid does not have an S_+^2 -lift. In fact, for any smooth convex cone, the longest chain of nonempty faces has length three (since the only faces are $\{0\}$, the extreme rays, and the cone itself). As such, we can conclude that the convex hull of the cardioid does not have a K -lift where K is a smooth convex cone.

Obstructions based on the lengths of chains of faces can be used to give simple lower bounds on the size of spectrahedral lifts of polytopes.

COROLLARY 5.6 ([37]). *Any spectrahedral lift of a polytope C has size at least $\dim(C) + 1$.*

Proof. For a polytope C , the dimension strictly increases along chains of faces, and any maximal chain of nonempty faces has length $1 + \dim(C)$. On the other hand, the longest chain of nonempty faces in \mathcal{S}_+^m has length at most $m + 1$. This is because the rank (which is constant on the relative interior of faces of \mathcal{S}_+^m) strictly increases along chains of faces, and the rank is zero on the face $\{0\}$. \square

The polytopes C that have spectrahedral lifts of size $\dim(C) + 1$ are known as *psd-minimal polytopes* [39]. These are the polytopes for which the lower bound of Corollary 5.6 is tight. Any 2-level polytope (see section 4.6) is psd-minimal. In particular, the stable set polytopes of perfect graphs from section 2 are examples of psd-minimal polytopes. There is a complete classification of psd-minimal polytopes in dimensions up to four in [38], obtained by a careful study of certain algebraic properties of their slack matrices.

5.1.3. Obstructions Based on Neighborliness. We now consider more elaborate features of the facial structure of convex bodies that, if present, can rule out much larger classes of lifts. These features build on the idea of neighborly polytopes.

DEFINITION 5.7. *A convex polytope C is k -neighborly if, for every subset of k extreme points of C , the convex hull of that set forms a face of C .*

Perhaps the most celebrated examples of neighborly polytopes are the cyclic polytopes [28]. These are obtained by taking the convex hull of v points on the moment curve (t, t^2, \dots, t^n) . Neighborly polytopes are extremal in the sense that among all n -dimensional polytopes with v vertices, the neighborly polytopes have the largest number of faces in each dimension. This property makes them central objects in polyhedral combinatorics. They also provide a geometric interpretation of key results in other contexts, such as compressive sensing [18].

Any face of a polytope is exposed by a linear functional. As such, a convex polytope C is k -neighborly if and only if, for every $I \subseteq \text{ext}(C)$ with $|I| = k$, there is a linear functional $\ell_I \in C^\circ$ that exposes exactly the extreme points in I (i.e., $\ell_I(x) = 1$ for $x \in I$ and $\ell_I(x) < 1$ for $x \in \text{ext}(C) \setminus I$). We can generalize this definition to nonpolyhedral convex bodies and to strict subsets of the extreme points. This property turns out to provide an obstruction to certain lifts of nonpolyhedral convex bodies.

DEFINITION 5.8. *A convex body C is k -neighborly with respect to $V \subseteq \text{ext}(C)$ if, for every $I \subset V$ with $|I| = k$, there exists $\ell_I \in C^\circ$ such that $\ell_I(x) = 1$ for $x \in I$ and $\ell_I(x) < 1$ for all $x \in V \setminus I$.*

One can check that a k -neighborly polytope is k -neighborly with respect to its entire set of extreme points. However, this definition is most interesting beyond the polyhedral setting.

EXAMPLE 5.9. *The set of 3×3 psd matrices with trace one, sometimes called the 3×3 spectraplex, is 2-neighborly with respect to the (countably infinite) set of extreme points*

$$V = \{v_p v_p^\top : p \text{ is an integer}\}, \quad \text{where } v_p = \frac{1}{\sqrt{1+p^2+p^4}} \begin{bmatrix} 1 \\ p \\ p^2 \end{bmatrix}.$$

To see this, we choose any pair $\{i, j\}$ of integers and define the linear functional $\ell_{\{i,j\}}$ in the polar of the spectraplex by

$$\ell_{\{i,j\}}(X) = \text{tr}(X) - w_{ij}^\top X w_{ij}, \quad \text{where} \quad w_{ij} = \begin{bmatrix} ij \\ -(i+j) \\ 1 \end{bmatrix}.$$

Since $w_{ij}^\top v_p = (p-i)(p-j)/\sqrt{1+p^2+p^4}$, it follows that $\ell_{\{i,j\}}(v_p v_p^\top) = 1 - \frac{1}{1+p^2+p^4}(p-i)^2(p-j)^2$, which clearly satisfies Definition 5.8. By a similar argument, the $(k+1) \times (k+1)$ spectraplex is k -neighborly with respect to an infinite set of extreme points.

Convex bodies with neighborliness properties arise very naturally when considering convex reformulations of polynomial optimization problems from the point of view taken in section 1. If $S \subseteq \mathbb{R}^n$ is compact and full-dimensional, we have seen that optimization of polynomials of degree $2d$ over S can be rephrased as linear optimization over the convex body

$$C = \text{cl conv}\{(x^\alpha)_{|\alpha| \leq 2d} : x \in S\}.$$

The extreme points all have the form $(x^\alpha)_{|\alpha| \leq 2d}$ for $x \in S$. The polar, C° , is exactly the set of coefficients of polynomials of degree $2d$ that take value at most one on S . Averkov [2], in a considerable generalization of Example 5.9, showed that for any positive integer $M \geq \binom{n+d}{n} - 1$, there is a subset V_M of extreme points of C with $|V_M| = M$ and such that C is $\binom{n+d}{n} - 1$ -neighborly with respect to V_M . In other words, C is $\binom{n+d}{n} - 1$ -neighborly with respect to arbitrarily large finite subsets of extreme points.

Being k -neighborly with respect to arbitrarily large finite sets of extreme points turns out to be an obstruction to K -lifts, where K is a Cartesian product of “low-complexity” convex cones. The first results in this direction considered the expressive power of second-order cone lifts [20]. Recall that these are K -lifts in which K is a finite Cartesian product of second-order cones, i.e., cones of the form

$$\mathcal{L}_+^{\ell+1} = \{(x_0, x) \in \mathbb{R} \times \mathbb{R}^\ell : \|x\| \leq x_0\}.$$

THEOREM 5.10 ([20]). *Let C be a convex body that is 2-neighborly with respect to arbitrarily large finite sets of extreme points. Then C does not have a second-order cone lift.*

Since the 3×3 spectraplex is 2-neighborly with respect to an infinite set of extreme points, it does not have a second-order cone lift. It follows (by a homogenization argument) that \mathcal{S}_+^3 does not have a second-order cone lift.

The second order cone $\mathcal{L}_+^{\ell+1}$ has an $(\mathcal{S}_+^2)^\ell$ -lift. As such, Theorem 5.10 can also be thought of as giving an obstruction to K -lifts, where K is a finite Cartesian product of 2×2 psd cones. Averkov [2] extended Theorem 5.10 to show that if m is a positive integer and C is k -neighborly with respect to arbitrarily large finite sets of extreme points, then C does not have an $(\mathcal{S}_+^k)^m$ -lift. This allowed him to show that many convex bodies associated with convex approaches to polynomial optimization do not have lifts using small psd blocks.

The property that C is k -neighborly with respect to arbitrarily large finite sets of extreme points is really a property of the face lattice of C . As such, we might expect that such a k -neighborliness property is an obstruction to having a K -lift for a class of cones K defined purely by properties of its faces.

THEOREM 5.11 ([81]). *Let m be a positive integer and let K_1, \dots, K_m be proper convex cones such that, for $i = 1, 2, \dots, m$, the length of the longest chain of nonempty faces of K_i is at most $k + 1$. If C is a convex body that is k -neighborly with respect to arbitrarily large finite sets of extreme points, then C does not have a $K_1 \times K_2 \times \dots \times K_m$ -lift.*

The length of the longest chain of nonempty faces of the $k \times k$ psd cone is $k + 1$. As such, Theorem 5.10 and Averkov's extension are special cases of Theorem 5.11 where we take $K_1 = K_2 = \dots = K_m = S_+^k$. Other specializations of Theorem 5.11 tell us, for instance, that

- the 3×3 psd cone has no K -lift, where K is a finite Cartesian product of smooth convex cones, because any smooth cone only has chains of nonempty faces of length at most three; and
- the 4×4 psd cone has no lifted representation using a finite Cartesian product of exponential cones, $\mathcal{K}_{\text{exp}} = \text{cl}\{(x, y, t) : y > 0, y \exp(x/y) \leq t\}$, since the longest chain of nonempty faces in any three-dimensional cone is four.

This last example is of interest due to recent developments in tractable relaxations of polynomial optimization problems based on geometric programming [12, 19]. The connection arises because geometric programs can, themselves, be reformulated as conic optimization problems over Cartesian products of exponential cones.

5.2. Algebraic Obstructions. We have so far discussed obstructions to the existence of lifts based on combinatorial properties of the face lattice of C . In this subsection we discuss *algebraic obstructions* to the existence of lifts.

5.2.1. Preliminaries. A *semialgebraic* subset of \mathbb{R}^n is a set that is described using Boolean combinations (finite unions, intersections, and complementations) of sets of the form

$$(5.1) \quad \{x \in \mathbb{R}^n : f_i(x) = 0, g_j(x) \geq 0 \ \forall i \in I, j \in J\},$$

where f_i, g_j are polynomials with real coefficients, and I and J are finite index sets. Most subsets of \mathbb{R}^n we have seen in this article are semialgebraic subsets. For example, any *spectrahedron* $S = \{x \in \mathbb{R}^n : A_0 + x_1 A_1 + \dots + x_n A_n \succeq 0\}$ is semialgebraic, since the constraint that a matrix is psd can be expressed using polynomial inequalities in the entries of the matrix. A celebrated result of Tarski shows that the projection of any semialgebraic set is semialgebraic.

THEOREM 5.12 (Tarski). *Let $S \subset \mathbb{R}^\ell$ be a semialgebraic set and $\pi : \mathbb{R}^\ell \rightarrow \mathbb{R}^n$ be a linear map. Then $\pi(S)$ is semialgebraic.*

A consequence of Tarski's theorem is that *if $C \subset \mathbb{R}^n$ has a spectrahedral lift, then C must be semialgebraic*. This already gives us one obstruction to the existence of spectrahedral lift. For example, it tells us that a set like $\{(x, y) \in \mathbb{R}^2 : y \geq \exp(x)\}$ (which is not semialgebraic) cannot have a spectrahedral lift of finite size.

5.2.2. Degree Bounds. Let C be a convex semialgebraic set. Define the (topological) boundary of C as $\partial C = (\text{cl } C) \setminus (\text{int } C)$. Assuming that our convex set C is full-dimensional in \mathbb{R}^n , it can be shown that the boundary is a semialgebraic set of dimension $n - 1$ [86], i.e., it is a hypersurface. This means that there exists a nonzero polynomial $p \in \mathbb{R}[x_1, \dots, x_n]$ such that $\partial C \subset \mathcal{Z}(p)$, where $\mathcal{Z}(p) = \{x \in \mathbb{R}^n : p(x) = 0\}$ is the zero set of p . The set $\mathcal{Z}(p)$ is known as the *Zariski closure* of ∂C . The *degree* of ∂C is the smallest degree of a polynomial p such that $\partial C \subset \mathcal{Z}(p)$, and the *algebraic*

boundary of C is defined as $\mathcal{Z}(p)$. Below we give several examples of convex sets for which we describe the algebraic boundary and the corresponding degree.

EXAMPLE 5.13 (disk). Consider the convex set $C = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq 1\}$, which has the 2×2 spectrahedral representation

$$C = \left\{ (x, y) \in \mathbb{R}^2 : \begin{bmatrix} 1-x & y \\ y & 1+x \end{bmatrix} \succeq 0 \right\}.$$

The topological boundary of this convex set is the circle $\{(x, y) \in \mathbb{R}^2 : x^2 + y^2 = 1\}$, which is described by a polynomial of degree two. Thus the degree of ∂C is equal to 2.

EXAMPLE 5.14 (oval). Consider the two-dimensional convex set defined by

$$(5.2) \quad C = \left\{ (x, y) \in \mathbb{R}^2 : \begin{bmatrix} x & 0 & y \\ 0 & 1 & -x \\ y & -x & 1 \end{bmatrix} \succeq 0 \right\},$$

which corresponds to the blue oval depicted in Figure 26. The smallest degree of a polynomial that vanishes on the boundary of C is 3; an example of such a polynomial is $p(x, y) = x - x^3 - y^2 = 0$. The zero set of $p(x, y)$ is shown with a thick black line in Figure 26. We see that it has two components. In this example the topological boundary and algebraic boundary are distinct.




Fig. 26 In blue is the convex set of (5.2). The solid black line represents the algebraic boundary of C , which has two components. The degree of the algebraic boundary of C is 3.

EXAMPLE 5.15 (k -ellipse). Given points $p_1, \dots, p_k \in \mathbb{R}^2$ in general position, consider the convex set

$$(5.3) \quad C = \{x \in \mathbb{R}^2 : \|x - p_1\|_2 + \dots + \|x - p_k\|_2 \leq 1\}.$$

For $k = 2$, the set C is an ellipse with focal points p_1 and p_2 , and the degree of the boundary of C is equal to 2. For higher k , the degree of the boundary of C was computed in [75] and was shown to be 2^k if k is odd, and $2^k - \binom{k}{k/2}$ if k is even. Figure 27 shows an example where $k = 3$.

We saw in Theorem 5.12 that the projection of any semialgebraic set is semialgebraic. More explicit formulations of Tarski's result give bounds on the degrees of the polynomials that define the semialgebraic set $\pi(S)$, in terms of the polynomials that




Fig. 27 A 3-ellipse (see (5.3)) with the focal points shown in red. The smallest degree of a polynomial that vanishes on the boundary of C is 8. The figure shows the zero set of this degree 8 polynomial.

define S ; see, e.g., [78]. These results can be used to show, for example, that if S is a spectrahedron defined using matrices of size $m \times m$, then the degree of the boundary of $\pi(S)$ is at most $m^{O(m^2n)}$ [37], where $n = \dim(\pi(S))$. It follows immediately that if C is a convex body of dimension n having a spectrahedral lift of size m , then necessarily $m^{O(m^2n)} \geq d$, where d is the degree of the boundary of C . Inverting this inequality allows us to get a lower bound on m in terms of d and n which reads

$$(5.4) \quad m \geq \Omega \left(\sqrt{\frac{\log d}{n \log \log(d/n)}} \right),$$

where Ω hides a constant.

The previous lower bound on m used very little information about the structure of the spectrahedron S living in higher dimensions and projecting onto C . A more refined analysis allows us to get the better lower bound [23]

$$(5.5) \quad m \geq \sqrt{\log d}.$$

This bound relies on the specific structure of the lift and uses the Karush–Kuhn–Tucker system of optimality conditions from convex optimization, combined with the Bézout bound to count zeros of systems of polynomial equations.

Observe that the bound (5.5) only depends on the degree d of the boundary of C and does not depend, say, on the dimension of C . One can easily show that no convex body can have a spectrahedral lift of size smaller than $\sqrt{\dim(C)}$. Indeed, if $C = \pi(S)$ and S is defined using an LMI of size $m \times m$, then we must have

$$(5.6) \quad \dim(C) \leq \dim(S) \leq \dim S^m = \binom{m+1}{2} \leq m^2,$$

which gives $m \geq \sqrt{\dim(C)}$.

We note that the two bounds presented above, based on degree and dimension, are not in general comparable:

- Consider the regular N -gon in the plane. The degree of the boundary is N and so the lower bound (5.5) gives $m \geq \sqrt{\log N}$. On the other hand, the bound (5.6) gives $m \geq \sqrt{2}$.
- Now consider the unit Euclidean ball in \mathbb{R}^n . The degree of the boundary is equal to two, since the boundary is given by the degree two polynomial equation $x_1^2 + \cdots + x_n^2 - 1 = 0$. Thus the lower bound (5.5) gives $m \geq \sqrt{2}$. On the other hand, the dimension lower bound gives $m \geq \sqrt{n}$.

One interesting open question is whether we can find a bound combining both degree d and dimension $\dim(C)$ which improves on (5.4).

5.2.3. Scheiderer's Result. In the previous section we saw, using Tarski's theorem, that in order for a convex set C to have a spectrahedral lift, it must be semialgebraic. Is this condition sufficient, i.e., does every convex semialgebraic set have a spectrahedral lift? Nemirovski raised this question in his 2006 ICM survey [73]. In [50, 51] Helton and Nie proved a series of results showing that every convex semialgebraic set whose boundary satisfies certain smoothness conditions admits a spectrahedral lift. They also conjectured that *any* convex semialgebraic set has a spectrahedral lift. In a breakthrough result, Scheiderer [82] disproved this conjecture and showed the existence of convex semialgebraic sets with no spectrahedral lift.

In this section we give a simple proof of Scheiderer's result based on the arguments in [21]. For concreteness we will focus on the convex set

$$C_{n,2d} = \text{cl conv} \{ (x^\alpha)_{|\alpha| \leq 2d} : x \in \mathbb{R}^n \},$$

where n and d are two integers. Note that $C_{n,2d}$ is a convex semialgebraic set since it is the convex hull of a semialgebraic set. We will show the following result.

THEOREM 5.16. *Assume n and d are two integers such that there is a homogeneous polynomial of degree $2d$ in n variables that is not a sum of squares. Then $C_{n,2d}$ has no spectrahedral lift.*

The theorem above shows, in particular, that $C_{3,6}$ and $C_{4,4}$ have no spectrahedral lift since it is known by a result of Hilbert that in the cases $(n, 2d) = (3, 6)$ and $(n, 2d) = (4, 4)$ there exist nonnegative polynomials that cannot be written as a sum of squares; see, e.g., [7, Chapter 4].

We remark that the assumption in the theorem above can be used to prove that $C_{n-1,2d}$ has no spectrahedral lift (which readily implies that $C_{n,2d}$ has no spectrahedral lift, since $C_{n-1,2d}$ is a projection of $C_{n,2d}$), but the argument is more technical and uses additional properties about semialgebraic functions; we refer the reader to [21] for more details.

Sketch of proof of Theorem 5.16. Let n and d be as in the statement of the theorem and assume for contradiction that $C_{n,2d}$ has a spectrahedral lift. We know from section 4 (see (*-SDP)) that there is a finite-dimensional subspace of functions that can be used to certify all valid linear inequalities on $C_{n,2d}$ using sums of squares. Linear inequalities valid on $C_{n,2d}$ are nothing more than nonnegative polynomials of degree $2d$ in n variables. Thus, the existence of a spectrahedral lift for $C_{n,2d}$ means that there is a finite number of functions $f_1, \dots, f_m : \mathbb{R}^n \rightarrow \mathbb{R}$ such that any nonnegative polynomial in $\mathbb{R}[x_1, \dots, x_n]_{\leq 2d}$ is a sum of squares from $\text{span}(f_1, \dots, f_m)$.

A key observation here is that the functions f_i can be taken to be *semialgebraic*. A semialgebraic function is one whose graph $\{(x, f(x)) : x \in \mathbb{R}^n\} \subset \mathbb{R}^{n+1}$ is a semialgebraic set. The observation about the f_i can be shown by going back to the proof of the factorization theorem (Theorem 3.9) and using some standard results from

semialgebraic geometry, e.g., the cylindrical algebraic decomposition of semialgebraic sets; see [21] for details. Semialgebraic functions are tame and possess nice qualitative properties. The following fact will be crucial for our argument (see, e.g., [48, Theorem 1.7] for a proof).

THEOREM 5.17. *A semialgebraic function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is smooth (C^∞) almost everywhere.*

We now have all the ingredients to prove our result. Since the f_i ($i = 1, \dots, m$) are smooth almost everywhere, there is at least one point $a \in \mathbb{R}^n$ such that all the f_i are smooth at the point a .

Let p be a homogeneous polynomial in n variables of degree $2d$ that is nonnegative but not a sum of squares. Consider the shifted polynomial $p(x + a)$. Since $p(x + a)$ is in $\mathbb{R}[x_1, \dots, x_n]_{\leq 2d}$ and nonnegative, it must be a sum of squares from $\text{span}(f_1, \dots, f_m)$. Shifting back to the origin, this means that $p(x)$ is a sum of squares from $\text{span}(\tilde{f}_1, \dots, \tilde{f}_m)$, where $\tilde{f}_i(x) = f_i(x - a)$. Note that all the functions in $\text{span}(\tilde{f}_1, \dots, \tilde{f}_m)$ are C^∞ at the origin. We now use the following proposition.

PROPOSITION 5.18. *Let p be a homogeneous polynomial of degree $2d$ in n variables and assume that $p = \sum_j h_j^2$, where the h_j are arbitrary functions that are smooth (C^∞) at the origin. Then p is a sum of squares of polynomials.*

Proof. We can write a Taylor expansion of each h_j of order d : $h_j(x) = q_j(x) + o(\|x\|^d)$, where $\deg q_j = d$. Let $\alpha_j = \text{mindeg } q_j \leq d$ (where $\text{mindeg } q_j$ is the smallest degree of a monomial in q_j) and note that $h_j(x)^2 = q_j(x)^2 + o(\|x\|^{d+\alpha_j})$. Then we get

$$p(x) = \sum_j q_j(x)^2 + o(\|x\|^{d+\alpha_j}).$$

Note that since p is homogeneous we must have $\alpha_j = d$ for all j , since otherwise there is a term of degree $2\alpha_j$ in $q_j(x)^2$ that cannot be canceled by the other terms. This means that the q_j are homogeneous of degree d . But then this means that we must have exact equality $p(x) = \sum_j q_j(x)^2$ and the remainder term $o(\|x\|^{2d})$ is identically zero. This shows that $p(x)$ is a sum of squares of polynomials. \square

End of proof of Theorem 5.16. The proposition above implies that p is a sum of squares of polynomials, which is a contradiction. We have thus shown that $C_{n,2d}$ has no spectrahedral lift. \square

6. Discussion. In this article we have focused on understanding exact lifts of convex sets. We saw several scenarios in which lifted representations greatly simplify the description of a convex set with applications to optimization and other fields. We also discussed in detail the key tool of slack matrices and operators that both determine the existence of lifts into specific cones and provide a method to construct them. Various construction methods for lifts, as well as obstructions to their existence, were also discussed. The main focus in these last two sections was on spectrahedral lifts. The purpose of this article was to show the broad connections of lifts to mathematics at large with a focus on spectrahedral lifts in the second half. A short and focused exposition on just spectrahedral lifts can be found in [88].

There is much more that one can say and study in and around the theory of lifts, and in this final section we collect several of these directions with pointers to the literature. We highlight the surveys and expository articles in each area. These aspects have not been discussed in this article but knit together a more complete view

of this rich subject. They also showcase the wide ranging connections between the theory of lifts and other disciplines.

Integer Programming/Combinatorial Optimization. In this article we have emphasized recent developments in spectrahedral lifts of convex semialgebraic sets, at the expense of more classical topics like polyhedral lifts of polytopes, usually called *extended formulations*. The study of such extended formulations has a long history, motivated by their importance in combinatorial optimization. A broad exposition of extended formulations of polytopes that arise in combinatorial optimization can be found in [16]. For a survey of mixed integer linear programming reformulations, see [92].

Lower Bounds for Polyhedral Lifts. The main result of the seminal paper of Yannakakis [95], mentioned in section 3, is that the *perfect matching polytope* of a complete graph does not admit a small *symmetric* polyhedral lift. Symmetric lifts were discussed in section 4. As mentioned there, Yannakakis asked whether symmetry imposes restrictions on the linear extension complexity of a polytope. In [57] Kaibel, Pashkovich, and Theis showed that, indeed, symmetry can force the size of a lift to be higher than necessary. Their examples are certain slices of the perfect matching polytope of a complete graph. A friendly introduction to polyhedral lifts including the symmetry questions can be found in [56]. A second breakthrough result on polyhedral lifts came in [27], where the authors showed that *cut polytopes* and *traveling salesman polytopes* do not admit small polyhedral lifts. The proof relied on showing that a submatrix of the slack matrix has exponentially high nonnegative rank. The underlying combinatorial problems in these examples, the *max cut problem* and the *traveling salesman problem*, are NP-hard, and so these results were expected although not known to be true explicitly. On the other hand, there is a polynomial time algorithm to find the maximum size matching in a graph which makes it plausible that the matching polytope has a small polyhedral lift. Rothvoss showed in [80] that the matching polytope does not in fact have a small extended formulation. Several results on combinatorial lower bounds on the extension complexity of polytopes and the techniques for establishing such bounds can be found in [26].

Lower Bounds for Spectrahedral Lifts. Lower bounds on the semidefinite extension complexity of polytopes have also received attention. In [65] the authors show that the cut and traveling salesman polytopes of graphs have super polynomial lower bounds on their semidefinite extension complexity. As before, such a result is expected but was not known explicitly. It remains open whether matching polytopes also have high semidefinite extension complexity. Also, no family of polytopes is known with a significant gap between linear extension complexity and semidefinite extension complexity; see [25] for an explicit example with a nontrivial gap, based on special cyclic polytopes. Semidefinite extension complexity can be studied for convex semialgebraic sets beyond polytopes. Not much is known in this direction beyond the results we have presented in section 5 of this paper.

Other Cones. Lifts of convex sets into cone families beyond nonnegative orthants, second-order cones, and psd cones have also been studied. For instance, a striking result in this direction is that every polytope obtained as the convex hull of a collection of 0/1 vectors in \mathbb{R}^n admits a lift into the *completely positive cone* of size $n + 1$ [10]. A symmetric matrix M is *copositive* if $x^\top M x \geq 0$ for all $x \geq 0$ and the *copositive cone* of size k is the set of all $k \times k$ copositive matrices. Its dual cone is the *completely positive cone* and consists of all symmetric matrices of the form BB^\top ,

where B is a nonnegative matrix. While the lifts mentioned above are small in terms of the size of the matrices involved, linear programming over copositive and completely positive cones does not admit efficient algorithms unless $P = NP$ [17, 70]. There are other families of convex cones for which the associated conic program enjoys efficient algorithms. Examples include *hyperbolic programming* [47], *geometric programming* [8], and *relative entropy programming* [13]. Studying lifted representations of convex sets using hyperbolicity cones or relative entropy cones gives a systematic approach to understanding the expressive power of these families of optimization problems.

Approximate Lifts and Hierarchies. Since finding an exact lift of a convex set might be hard, it is useful to look for an approximate lift in the sense of a set in a higher dimension whose projection contains the original convex set. There is a vast body of literature on such approximate lifts, as well as construction methods for producing a hierarchy of lifts in increasing dimensions whose projections create a nested sequence of convex approximations of the original set. The hierarchies most closely related to the discussion in section 4 of this paper are those of Lasserre [61] based on moments of probability measures and its dual due to Parrilo [76] based on sums of squares of polynomials. These methods are connected to deep results in real algebraic geometry, the theory of moments, and optimization, which give tools to address questions of convergence, construction, geometry, and more. See [64] for a survey on the Lasserre and Parrilo hierarchies. For a discussion of some of the central hierarchies for convex approximations of a set and their relationships, see [63]. See [7] for broad connections to convex algebraic geometry and related topics. The chapter [40] in *Handbook of Semidefinite, Conic and Polynomial Optimization* is about *theta bodies* of algebraic varieties, which is a specific hierarchy of spectrahedral lifts that successively approximate the convex hull of an algebraic variety. Theta bodies are also the subject of the chapter titled “Convex Hulls of Algebraic Sets” in [7]. The tight connection between conic lifts of convex semialgebraic sets and cone factorizations was inspired by the work on theta bodies and the paper of Yannakakis [95] that connects polyhedral lifts to nonnegative factorizations.

Approximation Algorithms. Hierarchies of convex relaxations have been studied extensively in theoretical computer science from the point of view of approximation algorithms. The most celebrated result in this context is perhaps the rounding algorithm for the max cut problem due to Goemans and Williamson, which yields an α -approximation with $\alpha \approx 0.878$ [32]. This is based on a simple spectrahedral relaxation of the cut polytope. Recently, hierarchies have been used to produce lower bounds on the approximation power of arbitrary polynomial-size linear (semidefinite) programming relaxations of hard combinatorial problems; see, e.g., [11, 65] and the references therein.

Computational Complexity. A natural question is to ask, what is the complexity of deciding whether a convex set C admits a K -lift?² To make the question precise, one needs to specify how the input convex set C is provided, e.g., as a set of extreme points, a set of valid inequalities, membership/separation oracle, etc., and similarly for the cone K . Existing results on the hardness of computing the nonnegative and psd rank of matrices [91, 85] can be used to show some NP-hardness results on the existence of lifts. For example, one can show that if $P \subset Q$ are two polytopes, where P is given by its set of vertices and Q by its set of facets, the question of deciding

²We thank one of the reviewers for raising this question.

whether there is a polytope S with an \mathbb{R}_+^r -lift such that $P \subset S \subset Q$ is NP-hard (here $r \in \mathbb{N}$ is part of the input). The same problem is also NP-hard if we replace \mathbb{R}_+^r by \mathcal{S}_+^r lift and is a consequence of the NP-hardness of computing the psd rank of matrices [85]. These hardness results follow from the fact that any nonnegative matrix M can be interpreted as the *generalized slack matrix/operator* (see Proposition 3.12) of two such polytopes P, Q by a simple rank factorization of M ; see, e.g., [30]. As far as we know, the same question where $P = Q$ is given by a \mathcal{V} or \mathcal{H} -representation is open, as not every nonnegative matrix is the slack matrix of a polytope [34].

Other Applications of Factorizations. Factorizations of nonnegative matrices have many uses beyond the theory of lifts. Nonnegative factorization of matrices is an essential tool for dimensionality reduction. In this case, one is often interested in approximate factorizations, rather than exact ones; see the chapter [29] for applications of nonnegative factorizations to data analysis. Nonnegative factorizations also have interpretations in probabilistic modeling, information theory, and communication complexity; see [15] and [27]. Positive semidefinite factorizations of nonnegative matrices in turn have applications in quantum information theory and quantum communication complexity; these are explained in [22] and [27]. The survey [22] concerns psd factorizations and psd ranks of general nonnegative matrices. It does not focus on slack matrices, which are the nonnegative matrices needed for lifts.

Symmetric factorizations also have applications. Recall that a matrix M is completely positive if there exist vectors $\{b_i\} \subset \mathbb{R}_+^k$ such that $M_{ij} = \langle b_i, b_j \rangle$. The *completely positive rank* of M is the smallest k for which such a symmetric factorization is possible. Analogously, a square matrix $M \in \mathbb{R}_+^{p \times p}$ is said to be *psd-completely-positive* if there exist psd matrices A_1, \dots, A_p (of arbitrary size) such that $M_{ij} = \text{tr}(A_i A_j)$ for all $i, j = 1, \dots, p$. These factorizations have found many applications in quantum information, in the context of nonlocal games; see, e.g., [42]. Tools have recently been developed to bound both the completely positive and the *psd-completely-positive ranks* of matrices [41, 43].

REFERENCES

- [1] A. A. AHMADI AND P. A. PARRILO, *A complete characterization of the gap between convexity and SOS-convexity*, SIAM J. Optim., 23 (2013), pp. 811–833, <https://doi.org/10.1137/110856010>. (Cited on p. 882)
- [2] G. AVERKOV, *Optimal size of linear matrix inequalities in semidefinite approaches to polynomial optimization*, SIAM J. Appl. Algebra Geom., 3 (2019), pp. 128–151, <https://doi.org/10.1137/18M1201342>. (Cited on p. 906)
- [3] H. H. BAUSCHKE AND J. M. BORWEIN, *Legendre functions and the method of random Bregman projections*, J. Convex Anal., 4 (1997), pp. 27–67. (Cited on p. 895)
- [4] A. BEN-TAL AND A. NEMIROVSKI, *Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications*, MPS/SIAM Ser. Optim. 2, SIAM, 2001, <https://doi.org/10.1137/1.9780898718829>. (Cited on p. 872)
- [5] A. BEN-TAL AND A. NEMIROVSKI, *On polyhedral approximations of the second-order cone*, Math. Oper. Res., 26 (2001), pp. 193–205, <https://doi.org/10.1287/moor.26.2.193.10561>. (Cited on p. 902)
- [6] G. BIRKHOFF, *Tres observaciones sobre el algebra lineal*, Univ. Nac. Tucumán Rev. Ser. A, 5 (1946), pp. 147–151. (Cited on p. 869)
- [7] G. BLEKHERMAN, P. A. PARRILO, AND R. R. THOMAS, EDS., *Semidefinite Optimization and Convex Algebraic Geometry*, MOS-SIAM Ser. Optim. 13, SIAM, Mathematical Optimization Society, 2013, <https://doi.org/10.1137/1.9781611972290>. (Cited on pp. 899, 910, 913)
- [8] S. BOYD, S.-J. KIM, L. VANDENBERGHE, AND A. HASSIBI, *A tutorial on geometric programming*, Optim. Eng., 8 (2007), pp. 67–127, <https://doi.org/10.1007/s11081-007-9001-7>. (Cited on p. 913)

- [9] R. E. BRYANT, *Graph-based algorithms for Boolean function manipulation*, IEEE Trans. Comput., 100 (1986), pp. 677–691, <https://doi.org/10.1109/TC.1986.1676819>. (Cited on p. 876)
- [10] S. BURER, *On the copositive representation of binary and continuous nonconvex quadratic programs*, Math. Program., 120 (2009), pp. 479–495, <https://doi.org/10.1007/s10107-008-0223-z>. (Cited on p. 912)
- [11] S. O. CHAN, J. R. LEE, P. RAGHAVENDRA, AND D. STEURER, *Approximate constraint satisfaction requires large LP relaxations*, J. ACM, 63 (2016), art. 34, <https://doi.org/10.1145/2811255>. (Cited on p. 913)
- [12] V. CHANDRASEKARAN AND P. SHAH, *Relative entropy relaxations for signomial optimization*, SIAM J. Optim., 26 (2016), pp. 1147–1173, <https://doi.org/10.1137/140988978>. (Cited on p. 907)
- [13] V. CHANDRASEKARAN AND P. SHAH, *Relative entropy optimization and its applications*, Math. Program., 161 (2017), pp. 1–32, <https://doi.org/10.1007/s10107-016-0998-2>. (Cited on p. 913)
- [14] F. CHEN, L. LOVÁSZ, AND I. PAK, *Lifting Markov chains to speed up mixing*, in Proceedings of the 31st ACM Symposium on Theory of Computing, Atlanta, GA, ACM, 1999, pp. 275–281, <https://doi.org/10.1145/301250.301315>. (Cited on p. 874)
- [15] J. E. COHEN AND U. G. ROTHBLUM, *Nonnegative ranks, decompositions, and factorizations of nonnegative matrices*, Linear Algebra Appl., 190 (1993), pp. 149–168, [https://doi.org/10.1016/0024-3795\(93\)90224-C](https://doi.org/10.1016/0024-3795(93)90224-C). (Cited on pp. 889, 914)
- [16] M. CONFORTI, G. CORNUÉJOLS, AND G. ZAMBELLI, *Extended formulations in combinatorial optimization*, Ann. Oper. Res., 204 (2013), pp. 97–143, <https://doi.org/10.1007/s10479-012-1269-0>. (Cited on p. 912)
- [17] P. J. C. DICKINSON AND L. GJIBEN, *On the computational complexity of membership problems for the completely positive cone and its dual*, Comput. Optim. Appl., 57 (2014), pp. 403–415, <https://doi.org/10.1007/s10589-013-9594-z>. (Cited on p. 913)
- [18] D. L. DONOHO AND J. TANNER, *Sparse nonnegative solution of underdetermined linear equations by linear programming*, Proc. Natl. Acad. Sci. USA, 102 (2005), pp. 9446–9451, <https://doi.org/10.1073/pnas.0502269102>. (Cited on p. 905)
- [19] M. DRESSLER, S. İLIMAN, AND T. DE WOLFF, *A Positivstellensatz for sums of nonnegative circuit polynomials*, SIAM J. Appl. Algebra Geom., 1 (2017), pp. 536–555, <https://doi.org/10.1137/16M1086303>. (Cited on p. 907)
- [20] H. FAWZI, *On representing the positive semidefinite cone using the second-order cone*, Math. Program., 175 (2019), pp. 109–118, <https://doi.org/10.1007/s10107-018-1233-0>. (Cited on pp. 875, 906)
- [21] H. FAWZI, *The set of separable states has no finite semidefinite representation except in dimension 3×2* , Comm. Math. Phys., 386 (2021), pp. 1319–1335, <https://doi.org/10.1007/s00220-021-04163-2>. (Cited on pp. 910, 911)
- [22] H. FAWZI, J. GOUVEIA, P. A. PARRILO, R. Z. ROBINSON, AND R. R. THOMAS, *Positive semidefinite rank*, Math. Program., 153 (2015), pp. 133–177, <https://doi.org/10.1007/s10107-015-0922-1>. (Cited on p. 914)
- [23] H. FAWZI AND M. SAFEY EL DIN, *A lower bound on the positive semidefinite rank of convex bodies*, SIAM J. Appl. Algebra Geom., 2 (2018), pp. 126–139, <https://doi.org/10.1137/17M1142570>. (Cited on p. 909)
- [24] H. FAWZI, J. SAUNDERSON, AND P. A. PARRILO, *Equivariant semidefinite lifts and sum-of-squares hierarchies*, SIAM J. Optim., 25 (2015), pp. 2212–2243, <https://doi.org/10.1137/140966265>. (Cited on p. 901)
- [25] H. FAWZI, J. SAUNDERSON, AND P. A. PARRILO, *Sparse sums of squares on finite abelian groups and improved semidefinite lifts*, Math. Program., 160 (2016), pp. 149–191, <https://doi.org/10.1007/s10107-015-0977-z>. (Cited on p. 912)
- [26] S. FIORINI, V. KAIBEL, K. PASHKOVICH, AND D. O. THEIS, *Combinatorial bounds on nonnegative rank and extended formulations*, Discrete Math., 313 (2013), pp. 67–83, <https://doi.org/10.1016/j.disc.2012.09.015>. (Cited on p. 912)
- [27] S. FIORINI, S. MASSAR, S. POKUTTA, H. R. TIWARY, AND R. DE WOLF, *Exponential lower bounds for polytopes in combinatorial optimization*, J. ACM, 62 (2015), art. 17, <https://doi.org/10.1145/2716307>. (Cited on pp. 912, 914)
- [28] D. GALE, *Neighborly and cyclic polytopes*, in Proc. Sympos. Pure Math., Vol. 7, AMS, 1963, pp. 225–232, <https://doi.org/10.1090/pspum/007/0152944>. (Cited on p. 905)
- [29] N. GILLIS, *The why and how of nonnegative matrix factorization*, in Regularization, Optimization, Kernels, and Support Vector Machines, Chapman & Hall/CRC Mach. Learn. Pattern Recogn. Ser., CRC Press, Boca Raton, FL, 2015, pp. 257–291, <https://doi.org/10.1201/b17558>. (Cited on p. 914)

- [30] N. GILLIS AND F. GLINEUR, *On the geometric interpretation of the nonnegative rank*, Linear Algebra Appl., 437 (2012), pp. 2685–2712, <https://doi.org/10.1016/j.laa.2012.06.038>. (Cited on p. 914)
- [31] M. X. GOEMANS, *Smallest compact formulation for the permutahedron*, Math. Program., 153 (2015), pp. 5–11, <https://doi.org/10.1007/s10107-014-0757-1>. (Cited on pp. 872, 902)
- [32] M. X. GOEMANS AND D. P. WILLIAMSON, *Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming*, J. ACM, 42 (1995), pp. 1115–1145, <https://doi.org/10.1145/227683.227684>. (Cited on pp. 870, 913)
- [33] M. C. GOLUBIC, *Algorithmic Graph Theory and Perfect Graphs*, Ann. Discrete Math. 57, North-Holland, 2004. (Cited on p. 881)
- [34] J. GOUVEIA, R. GRAPPE, V. KAIBEL, K. PASHKOVICH, R. Z. ROBINSON, AND R. R. THOMAS, *Which nonnegative matrices are slack matrices?*, Linear Algebra Appl., 439 (2013), pp. 2921–2933, <https://doi.org/10.1016/j.laa.2013.08.009>. (Cited on p. 914)
- [35] J. GOUVEIA AND T. NETZER, *Positive polynomials and projections of spectrahedra*, SIAM J. Optim., 21 (2011), pp. 960–976, <https://doi.org/10.1137/100801913>. (Cited on p. 900)
- [36] J. GOUVEIA, P. A. PARRILO, AND R. R. THOMAS, *Theta bodies for polynomial ideals*, SIAM J. Optim., 20 (2010), pp. 2097–2118, <https://doi.org/10.1137/090746525>. (Cited on p. 899)
- [37] J. GOUVEIA, P. A. PARRILO, AND R. R. THOMAS, *Lifts of convex sets and cone factorizations*, Math. Oper. Res., 38 (2013), pp. 248–264, <https://doi.org/10.1287/moor.1120.0575>. (Cited on pp. 893, 894, 905, 909)
- [38] J. GOUVEIA, K. PASHKOVICH, R. Z. ROBINSON, AND R. R. THOMAS, *Four-dimensional polytopes of minimum positive semidefinite rank*, J. Combin. Theory Ser. A, 145 (2017), pp. 184–226, <https://doi.org/10.1016/j.jcta.2016.08.002>. (Cited on p. 905)
- [39] J. GOUVEIA, R. Z. ROBINSON, AND R. R. THOMAS, *Polytopes of minimum positive semidefinite rank*, Discrete Comput. Geom., 50 (2013), pp. 679–699, <https://doi.org/10.1007/s00454-013-9533-x>. (Cited on p. 905)
- [40] J. GOUVEIA AND R. R. THOMAS, *Convex hulls of algebraic sets*, in Handbook on Semidefinite, Conic and Polynomial Optimization, Springer, 2012, pp. 113–138, https://doi.org/10.1007/978-1-4614-0769-0_5. (Cited on pp. 899, 900, 913)
- [41] S. GRIBLING, D. DE LAAT, AND M. LAURENT, *Matrices with high completely positive semidefinite rank*, Linear Algebra Appl., 513 (2017), pp. 122–148, <https://doi.org/10.1016/j.laa.2016.10.015>. (Cited on p. 914)
- [42] S. GRIBLING, D. DE LAAT, AND M. LAURENT, *Bounds on entanglement dimensions and quantum graph parameters via noncommutative polynomial optimization*, Math. Program., 170 (2018), pp. 5–42, <https://doi.org/10.1007/s10107-018-1287-z>. (Cited on p. 914)
- [43] S. GRIBLING, D. DE LAAT, AND M. LAURENT, *Lower bounds on matrix factorization ranks via noncommutative polynomial optimization*, Found. Comput. Math., 19 (2019), pp. 1013–1070, <https://doi.org/10.1007/s10208-018-09410-y>. (Cited on p. 914)
- [44] M. GRÖTSCHHEL, L. LOVÁSZ, AND A. SCHRIJVER, *Polynomial algorithms for perfect graphs*, in Topics on Perfect Graphs, North-Holland Math. Stud. 88, North-Holland, Amsterdam, 1984, pp. 325–356, [https://doi.org/10.1016/S0304-0208\(08\)72943-8](https://doi.org/10.1016/S0304-0208(08)72943-8). (Cited on p. 881)
- [45] M. GRÖTSCHHEL, L. LOVÁSZ, AND A. SCHRIJVER, *Relaxations of vertex packing*, J. Combin. Theory Ser. B, 40 (1986), pp. 330–343, [https://doi.org/10.1016/0095-8956\(86\)90087-0](https://doi.org/10.1016/0095-8956(86)90087-0). (Cited on p. 881)
- [46] M. GRÖTSCHHEL, L. LOVÁSZ, AND A. SCHRIJVER, *Geometric Algorithms and Combinatorial Optimization*, 2nd ed., Algorithms Combin. 2, Springer-Verlag, Berlin, 1993, <https://doi.org/10.1007/978-3-642-78240-4>. (Cited on p. 881)
- [47] O. GÜLER, *Hyperbolic polynomials and interior point methods for convex programming*, Math. Oper. Res., 22 (1997), pp. 350–377, <https://doi.org/10.1287/moor.22.2.350>. (Cited on p. 913)
- [48] H.-V. HÀ AND T.-S. PHAM, *Genericity in Polynomial Optimization*, Vol. 3, World Scientific, 2016. (Cited on p. 911)
- [49] H. HAUSER, *The Hironaka theorem on resolution of singularities (or: A proof we always wanted to understand)*, Bull. Amer. Math. Soc., 40 (2003), pp. 323–403, <https://doi.org/10.1090/S0273-0979-03-00982-0>. (Cited on p. 873)
- [50] J. W. HELTON AND J. NIE, *Sufficient and necessary conditions for semidefinite representability of convex hulls and sets*, SIAM J. Optim., 20 (2009), pp. 759–791, <https://doi.org/10.1137/07070526X>. (Cited on p. 910)
- [51] J. W. HELTON AND J. NIE, *Semidefinite representation of convex sets*, Math. Program., 122 (2010), pp. 21–64, <https://doi.org/10.1007/s10107-008-0240-y>. (Cited on pp. 882, 910)
- [52] D. HENRION, *Semidefinite representation of convex hulls of rational varieties*, Acta Appl. Math., 115 (2011), art. 319, <https://doi.org/10.1007/s10440-011-9623-9>. (Cited on p. 898)

- [53] H. HIRONAKA, *Resolution of singularities of an algebraic variety over a field of characteristic zero: I*, Ann. of Math. (2), 79 (1964), pp. 109–203, <https://doi.org/10.2307/1970486>. (Cited on p. 873)
- [54] H. HIRONAKA, *Resolution of singularities of an algebraic variety over a field of characteristic zero: II*, Ann. of Math. (2), 79 (1964), pp. 205–326, <https://doi.org/10.2307/1970547>. (Cited on p. 873)
- [55] A. HORN, *Eigenvalues of sums of Hermitian matrices*, Pacific J. Math., 12 (1962), pp. 225–241, <https://doi.org/10.2140/pjm.1962.12.225>. (Cited on p. 884)
- [56] V. KAIBEL, *Extended formulations in combinatorial optimization*, Optima, 85 (2011), pp. 2–7, <http://www.mathopt.org/Optima-Issues/optima85.pdf>. (Cited on p. 912)
- [57] V. KAIBEL, K. PASHKOVICH, AND D. O. THEIS, *Symmetry matters for sizes of extended formulations*, SIAM J. Discrete Math., 26 (2012), pp. 1361–1382, <https://doi.org/10.1137/110839813>. (Cited on pp. 901, 912)
- [58] A. A. KLYACHKO, *Stable bundles, representation theory and Hermitian operators*, Selecta Math. (N.S.), 4 (1998), pp. 419–445, <https://doi.org/10.1007/s000290050037>. (Cited on p. 885)
- [59] A. KNUTSON AND T. TAO, *The honeycomb model of $GL_n(\mathbb{C})$ tensor products. I. Proof of the saturation conjecture*, J. Amer. Math. Soc., 12 (1999), pp. 1055–1090, <https://doi.org/10.1090/S0894-0347-99-00299-4>. (Cited on p. 885)
- [60] A. KNUTSON AND T. TAO, *Honeycombs and sums of Hermitian matrices*, Notices Amer. Math. Soc., 48 (2001), pp. 175–186, <http://www.ams.org/notices/200102/fea-knutson.pdf>. (Cited on p. 885)
- [61] J. B. LASSERRE, *Global optimization with polynomials and the problem of moments*, SIAM J. Optim., 11 (2001), pp. 796–817, <https://doi.org/10.1137/S1052623400366802>. (Cited on p. 913)
- [62] J. B. LASSERRE, *Convex sets with semidefinite representation*, Math. Program., 120 (2009), pp. 457–477, <https://doi.org/10.1007/s10107-008-0222-0>. (Cited on p. 899)
- [63] M. LAURENT, *A comparison of the Sherali-Adams, Lovász-Schrijver, and Lasserre relaxations for 0-1 programming*, Math. Oper. Res., 28 (2003), pp. 470–496, <https://doi.org/10.1287/moor.28.3.470.16391>. (Cited on p. 913)
- [64] M. LAURENT, *Sums of squares, moment matrices and optimization over polynomials*, in Emerging Applications of Algebraic Geometry, IMA Vol. Math. Appl. 149, Springer, New York, 2009, pp. 157–270, https://doi.org/10.1007/978-0-387-09686-5_7. (Cited on p. 913)
- [65] J. R. LEE, P. RAGHAVENDRA, AND D. STEURER, *Lower bounds on the size of semidefinite programming relaxations*, in Proceedings of the 47th 2015 ACM Symposium on Theory of Computing, Portland, OR, ACM, New York, 2015, pp. 567–576, <https://doi.org/10.1145/2746539.2746599>. (Cited on pp. 912, 913)
- [66] H.-T. LIAW AND C.-S. LIN, *On the OBDD-representation of general Boolean functions*, IEEE Trans. Comput., 41 (1992), pp. 661–664, <https://doi.org/10.1109/12.144618>. (Cited on p. 877)
- [67] L. LOVÁSZ AND A. SCHRIJVER, *Cones of matrices and set-functions and 0-1 optimization*, SIAM J. Optim., 1 (1991), pp. 166–190, <https://doi.org/10.1137/0801013>. (Cited on p. 902)
- [68] W. J. MASEK, *A Fast Algorithm for the String Editing Problem and Decision Graph Complexity*, Master’s thesis, Massachusetts Institute of Technology, 1976. (Cited on p. 877)
- [69] S. MINATO, *Zero-suppressed BDDs for set manipulation in combinatorial problems*, in Proceedings of the 30th ACM/IEEE Design Automation Conference, IEEE, 1993, pp. 272–277, <https://doi.org/10.1145/157485.164890>. (Cited on p. 877)
- [70] K. G. MURTY AND S. N. KABADI, *Some NP-complete problems in quadratic and nonlinear programming*, Math. Program., 39 (1987), pp. 117–129, <https://doi.org/10.1007/BF02592948>. (Cited on p. 913)
- [71] R. M. NEAL, *Slice sampling*, Ann. Statist., 31 (2003), pp. 705–767, <https://doi.org/10.1214/aos/1056562461>. (Cited on p. 874)
- [72] R. M. NEAL, *MCMC using Hamiltonian dynamics*, in Handbook of Markov Chain Monte Carlo, Chapman and Hall/CRC, 2011, pp. 139–188. (Cited on p. 874)
- [73] A. NEMIROVSKI, *Advances in convex optimization: Conic programming*, in Proceedings of the International Congress of Mathematicians, Vol. 1, 2007, pp. 413–444. (Cited on p. 910)
- [74] Y. NESTEROV AND A. NEMIROVSKII, *Interior-Point Polynomial Algorithms in Convex Programming*, SIAM, 1994, <https://doi.org/10.1137/1.9781611970791>. (Cited on p. 872)
- [75] J. NIE, P. A. PARRILO, AND B. STURMFELS, *Semidefinite representation of the k -ellipse*, in Algorithms in Algebraic Geometry, Springer, 2008, pp. 117–132, https://doi.org/10.1007/978-0-387-75155-9_7. (Cited on p. 908)
- [76] P. A. PARRILO, *Semidefinite programming relaxations for semialgebraic problems*, Math. Program., 96 (2003), pp. 293–320, <https://doi.org/10.1007/s10107-003-0387-5>. (Cited on p. 913)

- [77] P. A. PARRILO, *Exact semidefinite representation for genus zero curves*, in Talk at the Banff Workshop “Positive Polynomials and Optimization,” Banff, Canada, 2006, pp. 8–12. (Cited on p. 898)
- [78] J. RENEGAR, *On the computational complexity and geometry of the first-order theory of the reals. I. Introduction. Preliminaries. The geometry of semi-algebraic sets. The decision problem for the existential theory of the reals*, J. Symbolic Comput., 13 (1992), pp. 255–299, [https://doi.org/10.1016/S0747-7171\(10\)80003-3](https://doi.org/10.1016/S0747-7171(10)80003-3). (Cited on p. 909)
- [79] R. T. ROCKAFELLAR, *Convex Analysis*, Princeton University Press, 1970. (Cited on p. 895)
- [80] T. ROTHVOSS, *The matching polytope has exponential extension complexity*, J. ACM, 64 (2017), art. 41, <https://doi.org/10.1145/3127497>. (Cited on p. 912)
- [81] J. SAUNDERSON, *Limitations on the expressive power of convex cones without long chains of faces*, SIAM J. Optim., 30 (2020), pp. 1033–1047, <https://doi.org/10.1137/19M1245670>. (Cited on p. 907)
- [82] C. SCHEIDERER, *Spectrahedral shadows*, SIAM J. Appl. Algebra Geom., 2 (2018), pp. 26–44, <https://doi.org/10.1137/17M1118981>. (Cited on pp. 875, 910)
- [83] A. SCHRIJVER, *Combinatorial Optimization: Polyhedra and Efficiency*, Algorithms Combin. 24, Springer Science & Business Media, 2003. (Cited on pp. 877, 880)
- [84] H. D. SHERALI AND W. P. ADAMS, *A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems*, SIAM J. Discrete Math., 3 (1990), pp. 411–430, <https://doi.org/10.1137/0403036>. (Cited on pp. 875, 902)
- [85] Y. SHITOV, *The complexity of positive semidefinite matrix factorization*, SIAM J. Optim., 27 (2017), pp. 1898–1909, <https://doi.org/10.1137/16M1080616>. (Cited on pp. 913, 914)
- [86] R. SINN, *Algebraic boundaries of convex semi-algebraic sets*, Res. Math. Sci., 2 (2015), art. 3, <https://doi.org/10.1186/s40687-015-0022-0>. (Cited on p. 907)
- [87] R. P. STANLEY, *Two poset polytopes*, Discrete Comput. Geom., 1 (1986), pp. 9–23, <https://doi.org/10.1007/BF02187680>. (Cited on pp. 878, 879)
- [88] R. R. THOMAS, *Spectrahedral lifts of convex sets*, in Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. IV. Invited Lectures, World Scientific, Hackensack, NJ, 2018, pp. 3819–3842. (Cited on p. 911)
- [89] A. VANDAELE, N. GILLIS, AND F. GLINEUR, *On the linear extension complexity of regular n -gons*, Linear Algebra Appl., 521 (2017), pp. 217–239, <https://doi.org/10.1016/j.laa.2016.12.023>. (Cited on p. 902)
- [90] L. VANDENBERGHE AND S. BOYD, *Semidefinite programming*, SIAM Rev., 38 (1996), pp. 49–95, <https://doi.org/10.1137/1038003>. (Cited on p. 870)
- [91] S. A. VAVASIS, *On the complexity of nonnegative matrix factorization*, SIAM J. Optim., 20 (2009), pp. 1364–1377, <https://doi.org/10.1137/070709967>. (Cited on p. 913)
- [92] J. P. VIELMA, *Mixed integer linear programming formulation techniques*, SIAM Rev., 57 (2015), pp. 3–57, <https://doi.org/10.1137/130915303>. (Cited on p. 912)
- [93] J. VON NEUMANN, *A certain zero-sum two-person game equivalent to the optimal assignment problem*, in Contributions to the Theory of Games, Vol. 2, H. W. Kuhn and A. W. Tucker, eds., Princeton University Press, 1953, pp. 5–12, <https://doi.org/10.1515/9781400881970-002>. (Cited on p. 869)
- [94] I. WEGENER, *Branching Programs and Binary Decision Diagrams: Theory and Applications*, Discrete Math. Appl. 4, SIAM, 2000, <https://doi.org/10.1137/1.9780898719789>. (Cited on p. 877)
- [95] M. YANNAKAKIS, *Expressing combinatorial optimization problems by linear programs*, J. Comput. System Sci., 43 (1991), pp. 441–466, [https://doi.org/10.1016/0022-0000\(91\)90024-Y](https://doi.org/10.1016/0022-0000(91)90024-Y). (Cited on pp. 874, 887, 889, 891, 901, 912, 913)
- [96] G. M. ZIEGLER, *Lectures on 0/1-polytopes*, in Polytopes—Combinatorics and Computation, Springer, 2000, pp. 1–41, https://doi.org/10.1007/978-3-0348-8438-9_1. (Cited on p. 874)