Studies Integrating Geometry, Probability, and Optimization under Convexity

by

Alexandre Belloni Nogueira

Submitted to the Sloan School of Management
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Operations Research

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2006

Author .....
Sloan School of Management
May 18, 2006

Certified by .....
Robert M. Freund
Theresa Seley Professor of Operations Research,
Sloan School of Management
Thesis Supervisor

Accepted by .....
James B. Orlin
Edward Pennell Brooks Professor of Operations Research,
Sloan School of Management

Studies Integrating Geometry, Probability, and Optimization under Convexity

by

Alexandre Belloni Nogueira

Submitted to the Sloan School of Management
on May 18, 2006, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Operations Research

Abstract

Convexity has played a major role in a variety of fields over the past decades. Nevertheless, the convexity assumption continues to reveal new theoretical paradigms and applications. This dissertation explores convexity in the intersection of three fields, namely, geometry, probability, and optimization.

We study in depth a variety of geometric quantities. These quantities are used to describe the behavior of different algorithms. In addition, we investigate how to algorithmically manipulate these geometric quantities. This leads to algorithms capable of transforming ill-behaved instances into well-behaved ones. In particular, we provide probabilistic methods that carry out such task efficiently by exploiting the geometry of the problem.

More specific contributions of this dissertation are as follows. (i) We conduct a broad exploration of the symmetry function of convex sets and propose efficient methods for its computation in the polyhedral case. (ii) We also relate the symmetry function with the computational complexity of an interior-point method to solve a homogeneous conic system. (iii) Moreover, we develop a family of pre-conditioners based on the symmetry function and projective transformations for such interior-point method. The implementation of the pre-conditioners relies on geometric random walks. (iv) We developed the analysis of the re-scaled perceptron algorithm for a linear conic system. In this method a sequence of linear transformations is used to increase a condition measure associated with the problem. (v) Finally, we establish properties relating a probability density induced by an arbitrary norm and the geometry of its support. This is used to construct an efficient simulating annealing algorithm to test whether a convex set is bounded, where the set is represented only by a membership oracle.

Thesis Supervisor: Robert M. Freund
Title: Theresa Seley Professor of Operations Research,
Sloan School of Management

Acknowledgments

Best words cannot express my gratitude to my thesis advisor Professor Robert M. Freund. I am indebted to him for his unconditional support as a mentor and as a friend during my PhD program. His high research and ethical standards will have a neverending influence on my career.

I am also very grateful to Victor Chernozhukov and Santosh Vempala. Their professional passion and talent deeply changed my perspective on research. Also, I found in Dimitris J. Bertsimas a great source of motivation and encouragement. As for Pablo Parrilo, I wish he had joined MIT earlier since he is a professional I have so much to learn from.

I greatly benefitted from discussions with many faculty members in conferences or at MIT. I would like to thank Don Goldfarb, Kurt M. Anstreicher, Clovis Gonzaga, Mike Todd, Yurii Nesterov, Arkadi Nemirovskii, Stephen Boyd, and Ward Whitt to name a few outside MIT, and Duncan Simester, Jim Orlins, Andreas Schultz, Stephen Graves among many others from MIT. I also want to add to this list my colleagues from IBM Thomas J. Watson Research Center. In particular, I am extremely grateful to Katya Scheinberg, Oktay Günlük, and Retsef Levi for their support and many advices.

The ORC was a second home for me. First of all, I would like to thank Jim Orlin and John Tsitsiklis who served as co-directors during my time at the ORC. The ORC students were great friends. I can only thank Fernando Ordonez, Peng Sun, Melvyn Sim, Nico, Victor, Felipe Caro, Agustin, Mike, Tim, Pranava, Guillaume, Dan, Katy, Shobit, Rags, Lincoln, among many others for enriching my experience at MIT. Finally, I also thank the staff of the ORC, Paulette P. Mosley, Laura A. Rose, Andrew Carvalho, Veronica Mignott, and Shiba Nemat-Nasser.

Still in the realm of academia, I also need to thank the ones who incentivated me to pursue a PhD in the United States, in particular, Claudia Sagastizabal, Michael Solodov, Abílio Lucena, Alfredo Iusem, Wilfredo Maldonado, Humberto Moreira, Juan Pablo Torres-Martinéz, and Aloisio Araújo.

I also would like to thank many friends for the encouragement during these four years. Primo, Fred, Willy (and Laura), Rafael, Coelho, Garcia, Maia, Caldinhas, Casa, Mipe, Leo, Leite, D., Ash, Jim, Craig, Travassos, Albuka, Alvaro, Novello, Varela, Henrique, João Pedro, Petrucio, José Koiller, Palmeira, Fuka, Pedro, Roberto and Bianca are just a few among many other names. Special thanks for Gustavo and Flavio for uncountable discussions in person or over the phone during my time at MIT and IMPA.

Above all, I am extremelly greateful to my parents, to my sister, and to my wife Rafaela. I owe to them more than I can say here. I cannot imagine my life without anyone of them. To them this dissertation is dedicated.

Finally, I would like to acknowledge that this research was supported in part by a MIT Presidential Fellowship and IBM PhD Fellowship.

To the Masschusetts Institute of Technology,
a home for the creative intellect.
(Norbert Wiener, "*Invention*")

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivations and philosophy

Convexity has played a major role in a variety of fields over the past decades. Nevertheless, the convexity assumption continues to reveal new theoretical paradigms and applications. This dissertation explores convexity in the intersection of three fields, namely, geometry, probability, and optimization.

In all likelihood, the fundamental result in convex analysis is the Hahn-Banach Theorem. Most of its versions pertain very general vector spaces, including infinite-dimensional ones. Here we focus on a geometric version of the Hahn-Banach Theorem for finite dimensional spaces.

Theorem 1 *Assume S is a nonempty, closed, convex subset of $\mathbb{R}^n$. If $x \in \mathbb{R}^n$ but $x \notin S$, then there exists $s \in \mathbb{R}^n$ such that*

$$\langle s, x \rangle < \langle s, y \rangle \quad \textit{for all } y \in S.$$

We say that the hyperplane defined by $(s, \langle s, x \rangle)$ separates x from S.

The importance of this result cannot be overstated. Theorem 1 is the driving force underlying many fundamental results including convex duality theory for optimization.

Given the geometric nature of this result, we expect many connections between convex analysis and convex geometry, where the latter studies the geometry of convex sets and related geometric inequalities in finite dimensional Euclidean spaces. Herein we will be dealing with both bounded and unbounded convex sets. Among unbounded convex sets, convex cones play an important role. In particular, convex cones are the convex analysis the counterpart of subspaces in analysis. Although most of our results hold for arbitrary convex cones, we are particularly interested on the non-negative orthant, on the second order cone, and on the cone of positive semi-definite matrices. These three cones form the backbone of modern convex optimization theory and methods.

A remarkable property of convexity is its close relation with computational complexity both in theory and in practice. The subject of computational complexity has its roots in the effort in the 1930's to classify problems into decidable or undecidable. Over the next decades, the focus changed to further classify decidable problems into computationally tractable or untractable. In the 1970's these notions were formalized by many NP-completeness results mainly for discrete problems. Roughly speaking, a method is said to be efficient for a problem if its computational time to solve any instance is bounded by a fixed polynomial on the size of the instance. We point out that the size of the instance is usually measured in bits due to the discrete nature of the problems.

In real number algorithms, like Newton Method, several additional issues need to be addressed. First one needs to accept real numbers as input and the bit model no longer applies. Moreover, many problems cannot be solved exactly even assuming exact arithmetic. In such cases we need to introduce an additional parameter $\varepsilon < 1$ which captures the error between the desired and the computed solution. Finally, there are instances, called ill-posed, where many algorithms simply fail. To deal with that one defines a condition measure μ for each instance; the larger μ is, the more ill-posed the instance is. In this context, a method is said to be efficient if its complexity is bounded by a function that grows at most polynomially in the size of the instance, and logarithmically on the approximation error and on the condition measure.

Probabilistic methods, such as simulated annealing, represent a different paradigm with respect to computational complexity. In fact, these methods are capable of efficiently solving a broader class of problems by allowing a positive probability of failure $\delta < 1$ which can be made arbitrarily small. In addition to the previous requirement, the dependence of the computational complexity on $1/\delta$ should be logarithmic for a method to be considered efficient. A celebrated example of that is the computation of the volume of a convex set, see [59] and the references therein.

As mentioned before, convexity still plays a central role in defining the frontier of polynomial time algorithms. Interior point methods for conic programming, ellipsoid method for convex feasibility, and sampling random points according to log-concave densities are examples where convexity plays a crucial role to establish efficiency. In other contexts, we are interested in the computation of certificates whose existence is equivalent to a desired property being true. Important examples are certificates of infeasibility, Farkas Lemma, S-Lemma, Nullstellensatz, and Positivstellensatz. Surprisingly, convexity is fundamental for the existence of such certificates in many important cases. With that in mind - and since convexity is the unifying theme herein - the development of efficient algorithms to compute different quantities is also an object of interest in this dissertation.

In the chapters to come we will propose several methods combining deterministic and probabilistic algorithms. These methods will be applied to solve the homogeneous linear conic feasibility system

$$\begin{aligned} Ax &\in K \\ x &\neq 0 \end{aligned} \tag{1.1}$$

and its alternative system

$$\begin{aligned} A^*y &= 0 \\ y &\in K^*. \end{aligned} \tag{1.2}$$

That sets the stage for this thesis. We propose several methods to solve (1.1) and (1.2) and study their computational complexity. The condition measures associated with these algorithms are geometric objects whose properties are also investigated here. As opposed to most of the literature which has a passive approach with respect

to the condition measure, we show how to algorithmically improve these condition measures to obtain a better behaved instance. Remarkably, probabilistic methods play a central role in the implementation of such algorithms.

1.2 Organization and contributions

1.2.1 Symmetry of convex sets

We develop properties and connections between the *symmetry function* of a convex set $S \subset \mathbb{R}^n$ $\mathbf{sym}(x, S)$, where $\mathbf{sym}(S) = \max_x \mathbf{sym}(x, S)$, and other arenas of convexity including convex functions, convex geometry, probability theory on convex sets, and computational complexity.

There are many important properties of symmetric convex sets; herein we explore how these properties extend as a function of $\mathbf{sym}(S)$ and/or $\mathbf{sym}(x, S)$. By accounting for the role of the symmetry function, we reduce the dependence of many mathematical results on the strong assumption that S is symmetric, and we are able to quantify many of the ways in which the symmetry function influences properties of convex sets and functions.

The results include functional properties of $\mathbf{sym}(x, S)$, relations with several convex geometry quantities such as volume, distance, and cross-ratio distance, as well as set approximation results, including a refinement of the Löwner-John rounding theorems, and applications of symmetry to probability theory on convex sets. We provide a characterization of symmetry points x^* for general convex sets. Finally, in the polyhedral case, we show how to efficiently compute $\mathbf{sym}(S)$ and a symmetry point x^*.

1.2.2 Projective pre-conditioners for IPM

The traditional strategy for speeding up the performance of modern interior-point methods for solving convex optimization problems has been to focus on pre-conditioners to try to improve the numerical performance of the equation systems solved at each

iteration. In contrast, we present a methodology for mapping the entire problem to one that is better behaved (geometrically and in other respects) in order to actually reduce the number of interior-point iterations that are required to solve the problem. This is achieved by improving a condition measure associated with the problem.

We present a general theory for transforming a normalized homogeneous conic system

$$F : \begin{cases} Ax &= 0 \\ \bar{s}^T x &= 1 \\ x &\in C \end{cases}$$

to an equivalent system via projective transformation induced by the choice of a point $\hat{v}$ in the convex set $H_{\bar{s}}^{\circ} = \{v : \bar{s} - A^T v \in C^*\}$ associated with the system F. Such a projective transformation serves to pre-condition the conic system into a system that has both geometric and computational properties with certain guarantees. We characterize both the geometric behavior and the computational behavior of the transformed system as a function of the symmetry of $\hat{v}$ in $H_{\bar{s}}^{\circ}$ as well as the complexity parameter ϑ of the barrier for C. Under the assumption that F has an interior solution, $H_{\bar{s}}^{\circ}$ must contain a point v whose symmetry is at least $1/m$; if we can find a point whose symmetry is $\Omega(1/m)$ then we can projectively transform the conic system to one whose geometric properties and computational complexity will be strongly-polynomial-time in m and ϑ. We present a method for generating such a point $\hat{v}$ based on sampling and on a geometric random walk on $H_{\bar{s}}^{\circ}$ with associated complexity and probabilistic analysis.

1.2.3 Efficiency of a re-scaled perceptron algorithm for conic systems

The classical perceptron algorithm is a simple and intuitive row-action/relaxation algorithm for solving a homogeneous linear inequality system $Ax \geq 0$, $x \neq 0$. A condition measure associated with this algorithm is the Euclidean width of the cone of feasible solutions. In fact, the complexity of the perceptron algorithm is bounded

by $1/\tau^2$ which is inefficient in general since we would like a logarithmic dependence on the condition measure.

Recently, Dunagan and Vempala have developed a re-scaling version of the perceptron algorithm with an improved complexity of $O(n \ln(1/\tau))$ iterations (with high probability), which is notable in that it is theoretically efficient, i.e., polynomial-time in the bit-length model. We explore extensions of the concepts of these perceptron methods to the general homogeneous conic system $Ax \in C,\ x \neq 0$. We provide a conic extension of the re-scaling perceptron algorithm based on the notion of a deep separation oracle of a cone, which essentially computes a certificate of strong separation for cones. We propose an implementation of this oracle which makes use of new properties developed here. In the case of linear, second-order, semi-definite cones, as well as certain other cones, we show that the resulting re-scaling algorithm is efficient.

1.2.4 Norm-induced densities and testing the boundedness of a convex set

The geometry of convex sets has been extensively studied during the past century. More recently, the interplay between convex geometry and probability theory has been investigated. Among many possible research directions, log-concave probability measures provide an interesting framework that generalizes uniform densities on convex sets but preserves many interesting properties. Here we will focus on a subset of logconcave functions called norm-induced densities.

In this chapter we have two goals. The first one is to explore relations between a variety of geometric quantities associated with convex sets and norm-induced densities. The second is to develop an efficient algorithm to test if a given convex set $K \subset \mathbb{R}^n$ is bounded or unbounded.

Chapter 2

On the symmetry function of a convex set

We attempt a broad exploration of properties and connections between the *symmetry function* of a convex set $S \subset \mathbb{R}^n$ and other areas of convexity including convex functions, convex geometry, probability theory on convex sets, and computational complexity. Our starting point is the definition of our object of interest. Given a closed convex set S and a point $x \in S$, define the symmetry of S about x as follows:

$$\mathbf{sym}(x, S) := \max\{\alpha \geq 0 : x + \alpha(x - y) \in S \text{ for every } y \in S\} , \tag{2.1}$$

which intuitively states that $\mathbf{sym}(x, S)$ is the largest scalar α such that every point $y \in S$ can be reflected through x by the factor α and still lie in S. The symmetry value of S then is:

$$\mathbf{sym}(S) := \max_{x \in S} \mathbf{sym}(x, S) , \tag{2.2}$$

and x^* is a *symmetry point* of S if x^* achieves the above maximum (also called a "critical point" in [25], [27] and [44]). S is *symmetric* if $\mathbf{sym}(S) = 1$. There are a variety of other measures of symmetry (or asymmetry) for a convex set that have been studied over the years, see Grünbaum [25] for a survey; the symmetry measure based on (2.2) is due to Minkowski [44], which in all likelihood was the first and most

useful such symmetry measure.

We explore fundamental properties of $\mathbf{sym}(x, S)$, and we present new results in other areas of convexity theory that are connected to the symmetry function. In Section 2.1 we examine analytical properties of the function $\mathbf{sym}(x, S)$. We show that $\mathbf{sym}(x, S)$ is a quasiconcave function, and more importantly, that $\mathbf{sym}(x, S)$ is a log-concave function and therefore inherits some of the strong results of logconcave functions related to sampling on convex sets (Theorem 2). We also show that $\mathbf{sym}(x, S)$ is the infimum of linear fractional functions related to the supporting hyperplanes of S (Proposition 1). In Proposition 3 we explore the behavior of $\mathbf{sym}(x, S)$ under basic set operations such as intersection, Minkowski sums, polarity, Cartesian product, and affine transformation. Furthermore, we completely characterize $\mathbf{sym}(x, S)$ when S is symmetric in Proposition 2.

In Section 2.2 we focus on connections between $\mathbf{sym}(x, S)$ and a wide variety of geometric properties of convex bodies, including volume ratios, distance metrics, set-approximation and rounding results, and probability theory on convex sets. It is well-known that any half-space whose bounding hyperplane passes through the center of mass z_S of S will cut off at least $1/\mathbf{e}$ and at most $1 - 1/\mathbf{e}$ of the volume of S, see Grünbaum [24]. In a similar vein, in Section 2.2.1 we present lower and upper bounds on ratios of volumes of S to the intersection of S with a half-space whose bounding hyperplane passes through x, as a function of $\mathbf{sym}(x, S)$ (Theorem 3), as well as lower bounds on the $(n-1)$-dimensional volume ratios of slices of S defined by the intersection of S with a hyperplane passing through x, as a function of $\mathbf{sym}(x, S)$ (Theorem 4).

If S is a symmetric convex body, then it is a straightforward exercise to show that the symmetry point of S is unique. Furthermore, if S is nearly symmetric, intuition suggests that two points in S with high symmetry values cannot be too far apart. This intuition is quantified Section 2.2.2, where we present upper bounds on the relative distance (in any norm) between two points $x, y \in S$ as a function of $\mathbf{sym}(x, S)$ and $\mathbf{sym}(y, S)$ (Theorem 5) and upper bounds on the "cross-ratio distance" in Theorem 6.

Section 2.2.3 examines the approximation of the convex set S by another convex set P. We say that P is a β-approximation of S if there exists a point $x \in S$ such that $\beta P \subset S - x \subset P$. In the case when P is an ellipsoid centered at the origin, then the statement "P is a β-approximation of S" is equivalent to "βP provides a $\frac{1}{\beta}$-rounding of S." We examine the interrelationship between the symmetry function and bounds on β-approximations for S. We show that for any $x \in S$ there exists a $\sqrt{n}/\mathbf{sym}(x, S)$-rounding of S centered at x (Theorem 8). A classical example of β-approximation is given by the Löwner-John theorem [30], which guarantees a $1/\sqrt{n}$-approximation for a symmetric convex body and a $1/n$-approximation for general convex body using ellipsoids. Unfortunately, the theorem does not provide more precise bounds for case when S is nearly symmetric, i.e., $\mathbf{sym}(S) = 1-\varepsilon$ for ε small. This is partially rectified herein, where we prove a slightly stronger rounding results using $\mathbf{sym}(x, S)$ (Theorem 10). We also show that if two convex sets are nearly the same, then their symmetry must be nearly the same (Theorem 9). Moreover, we show how to construct a norm based on $\mathbf{sym}(S)$ that yields the optimal β-approximation of S among all symmetric convex bodies (Lemma 1).

Subsection 2.2.4 is concerned with connections between symmetry and probability theory on convex sets. Let X be a random vector uniformly distributed on S. We show that the expected value of $\mathbf{sym}(X, S)$ is nicely bounded from below (by $\mathbf{sym}(S)/(2(n+1))$) and we present lower bounds on the probability that $\mathbf{sym}(X, S)$ is within a constant M of $\mathbf{sym}(S)$. Furthermore, in the case when S is symmetric, these quantities have closed-form expressions independent of the specific set S (Theorem 11). We also present an extension of Anderson's Lemma [1] concerning the the integral of a nonnegative logconcave even function on S, to the case of non-symmetric convex sets (Theorem 12), which has many statistical applications.

Since symmetry points enjoy many interesting properties, it is natural to explore methods for computing a symmetry point and for computing $\mathbf{sym}(S)$, which is the subject of Section 2.4. As expected, the representation of S plays a major role in any computational scheme. While the problem of simply evaluating the $\mathbf{sym}(x, S)$ for a given $x \in S$ is a hard problem in general, it turns out that for polyhedra, whose most

common representations are as the convex hull of points and as the intersection of half-spaces, computing a symmetry point can be accomplished via linear programming. When S is given as the convex hull of m points, we show that determining a symmetry point can be computed by solving a linear program in m^2 nonnegative variables, or as non-differentiable concave maximization problem where supergradients can be computed by solving m decoupled linear programming subproblems with only m nonnegative variables each. The more interesting case is when S is given as the intersection of m half-spaces. Then a symmetry point and $\mathbf{sym}(S)$ can be computed by solving $m+1$ linear programs with m nonnegative variables. We present an interior-point algorithm that, given an approximate analytic center x^a of S, will compute an approximation of $\mathbf{sym}(S)$ to any given relative tolerance ε in no more than

$$\left\lceil 10m^{1.5} \ln\left(\frac{10m}{\epsilon}\right)\right\rceil$$

iterations of Newton's method.

This work also contains a variety of discussions of open questions as well as unproved conjectures regarding the symmetry function and its connection to other areas of convexity theory.

2.0.5 Notation

Let $S \subset \mathbb{R}^n$ denote a convex set and let $\langle \cdot, \cdot \rangle$ denote the conventional inner product in the appropriate Euclidean space. **int** S denotes the interior of S. Using traditional convex analysis notation, we let aff(S) be the minimal affine subspace that contains S and let $S^{\perp}$ be its orthogonal subspace complement. The polar of S is defined as $S^{\circ} = \{y \in \mathbb{R}^n : \langle x, y \rangle \leq 1 \text{ for all } x \in S\}$. Given a convex function $f(\cdot)$, for $x \in \text{dom} f(\cdot)$ the subdifferential of $f(\cdot)$ is defined as $\partial f(x) := \{s \in \mathbb{R}^n : f(y) \geq f(x) + \langle s, y - x \rangle \text{ for all } y \in \text{dom} f(\cdot)\}$. Let $e = (1, \ldots, 1)^T$ denote the vector of ones whose dimension is dictated by context, let $\mathbf{e}$ denote the base of the natural logarithm, and let $\text{dist}(x, T) := \min_{y \in T} \|y - x\|$ be the distance from x to the set T in the norm $\|\cdot\|$ dictated by context.

2.1 Analytical properties of $\mathbf{sym}(x, S)$

We make the following assumption:

Assumption A: S is a convex body, i.e., S is a nonempty closed bounded convex set with a nonempty interior.

When S is a convex set but is either not closed or is unbounded, then certain properties of $\mathbf{sym}(S)$ break down; we refer the interested reader to Appendix B.1 for a discussion of these general cases. We assume that S has an interior as a matter of convenience, as one can always work with the affine hull of S or its subspace translation with no loss of generality, but at considerable notational and expositional expense.

There are other definitions of $\mathbf{sym}(x, S)$ equivalent to (2.1). In [52], $\mathbf{sym}(x, S)$ is defined by considering the set $\mathcal{L}(x, S)$ of all chords of S that pass through x. For $L \in \mathcal{L}(x, S)$, let $r(L)$ denote the ratio of the length of the smaller to the larger of the two intervals in $L \cap (S \setminus \{x\})$, and define

$$\mathbf{sym}(x, S) = \inf_{L \in \mathcal{L}(x,S)} r(L) \ . \tag{2.3}$$

Herein it will be convenient to also use the following set-containment definition of $\mathbf{sym}(x, S)$:

$$\mathbf{sym}(x, S) = \max \{\alpha \geq 0 : \alpha(x - S) \subseteq (S - x)\} \ . \tag{2.4}$$

It turns out that this definition is particularly useful to motivate and prove many of our results.

Intuition suggests that $\mathbf{sym}(x, S)$ inherits many nice properties from the convexity of S, as our first result shows:

Theorem 2 *Under Assumption A,*
(i) $\mathbf{sym}(\cdot, S) : S \to [0, 1]$ *is a continuous quasiconcave function,*
(ii) $h(x, S) := \dfrac{\mathbf{sym}(x, S)}{1 + \mathbf{sym}(x, S)}$ *is a concave function on* S*, and*
(iii) $\mathbf{sym}(\cdot, S)$ *is a logconcave function on* S*.*

Regarding part (iii) of the theorem, note that logconcave functions play a central

role in the theory of probability and sampling on convex bodies, see [41]. The proof of this theorem will use the following proposition, which will also be useful in the development of an algorithm for computing $\mathbf{sym}(S)$ in Section 2.4.

Proposition 1 *Let S be a convex body, and consider the representation of S as the intersection of halfspaces: $S = \{x \in \mathbb{R}^n : a_i^T x \leq b_i \ , i \in I\}$ for some (possibly unbounded) index set I, and let $\delta_i^* := \max_{x \in S}\{-a_i^T x\}$ for $i \in I$. Then for all $x \in S$,*

$$\mathbf{sym}(x, S) = \inf_{i \in I} \left\{ \frac{b_i - a_i^T x}{\delta_i^* + a_i^T x} \right\} .$$

Proof Let $\alpha = \mathbf{sym}(x, S)$ and $\gamma := \min\limits_{i \in I} \left\{ \dfrac{b_i - a_i^T x}{\delta_i^* + a_i^T x} \right\}$. Then for all $y \in S$, $x + \alpha(x - y) \in S$, so

$$a_i^T x + \alpha a_i^T x + \alpha(-a_i^T y) \leq b_i \ , i \in I \ .$$

This implies that

$$a_i^T x + \alpha a_i^T x + \alpha \delta_i^* \leq b_i \ , i \in I \ ,$$

whereby $\alpha \leq \gamma$. On the other hand, for all $y \in S$ we have:

$$b_i - a_i^T x \geq \gamma(\delta_i^* + a_i^T x) \geq \gamma(-a_i^T y + a_i^T x) \ .$$

Thus $a_i^T x + \gamma a_i^T x + \gamma(-a_i^T y) \leq b_i$, and therefore $a_i^T(x + \gamma(x - y)) \leq b_i$ which implies that $\alpha \geq \gamma$. Thus $\alpha = \gamma$.□

Proof of Theorem 2.

Proof We first prove (ii). It follows from Proposition 1 that

$$h(x, S) = \frac{\min_{i \in I} \left\{ \frac{b_i - a_i^T x}{\delta_i^* + a_i^T x} \right\}}{1 + \min_{i \in I} \left\{ \frac{b_i - a_i^T x}{\delta_i^* + a_i^T x} \right\}} = \min_{i \in I} \left\{ \frac{\frac{b_i - a_i^T x}{\delta_i^* + a_i^T x}}{1 + \frac{b_i - a_i^T x}{\delta_i^* + a_i^T x}} \right\} = \min_{i \in I} \left\{ \frac{b_i - a_i^T x}{b_i + \delta_i^*} \right\} \ ,$$

which is the minimum of linear functions and so is concave.

To prove (i), first observe that $\mathbf{sym}(x, S)$ is monotone in the concave function $h(x, S)$, and so is quasiconcave. To prove the continuity of $\mathbf{sym}(x, S)$ it suffices to

prove the continuity of $h(x,S)$. It follows from concavity that $h(x,S)$ is continuous on int$S$. For $\bar{x} \in \partial S$ it follows from (2.1) that $\mathbf{sym}(x,S) = 0$ and hence $h(x,S) = 0$. Because S is a convex body there exists a ball of radius $r > 0$ that is contained in S. Now suppose that $x^j \to \bar{x}$, whereby $\text{dist}(x^j, \partial S) \to 0$. It follows from (2.4) that $\mathbf{sym}(x^j,S) \cdot r \leq \text{dist}(x^j, \partial S)$, whereby $\mathbf{sym}(x^j,S) \to 0 = \mathbf{sym}(\bar{x},S)$, showing continuity of $h(x,S)$ and hence of $\mathbf{sym}(x,S)$ on S.

To prove (iii) define the following functions:

$$f(t) = \frac{t}{1+t} \quad \text{and} \quad g(t) = \ln\left(\frac{t}{1-t}\right).$$

For these functions, we have the following properties:
(i) f is monotone, concave and $f(\mathbf{sym}(x,S)) \in [0,1/2]$ for any $x \in S$;
(ii) g is monotone for $t \in (0,1)$ and concave for $t \in (0,1/2]$;
(iii) $g(f(t)) = \ln t$.

Now, for any $\alpha \in [0,1]$, $x, y \in S$,

$$\begin{aligned}
\ln\left(\mathbf{sym}(\alpha x + (1-\alpha)y, S)\right) &= g\left(f\left(\mathbf{sym}(\alpha x + (1-\alpha)y, S)\right)\right) \\
&\geq g\left(\alpha f\left(\mathbf{sym}(x,S)\right) + (1-\alpha) f\left(\mathbf{sym}(y,S)\right)\right) \\
&\geq \alpha g\left(f\left(\mathbf{sym}(x,S)\right)\right) + (1-\alpha) g\left(f\left(\mathbf{sym}(y,S)\right)\right) \\
&= \alpha \ln \mathbf{sym}(x,S) + (1-\alpha) \ln \mathbf{sym}(y,S)\ ,
\end{aligned}$$

where the first inequality follows from the concavity of $h(\cdot,S) = f(\mathbf{sym}(\cdot,S))$ and the monotonicity of g, and the second inequality follows from the concavity of g on $[0,1/2]$. □

It is curious that $\mathbf{sym}(\cdot,S)$ is not a concave function. To see this, consider $S = [0,1] \subset \mathbb{R}$; then a trivial computation yields $\mathbf{sym}(x,S) = \min\left\{\frac{x}{(1-x)}; \frac{(1-x)}{x}\right\}$, which is not concave on S and is not differentiable at $x = \frac{1}{2}$. Part (ii) of Theorem 2 shows that a simple nonlinear transformation of the symmetry function is concave.

For a symmetric convex body S, i.e., $\mathbf{sym}(S) = 1$, it is possible to prove a stronger statement and completely characterize the symmetry function using the norm induced

by S. Suppose S is a symmetric convex set centered at the origin. Let $\|\cdot\|_S$ denote the norm induced by S, namely $\|x\|_S := \min\{\gamma : x \in \gamma S\}$.

Proposition 2 *Under Assumption A, let S be symmetric and centered at the origin. Then for every $x \in S$,*

$$\mathbf{sym}(x, S) = \frac{1 - \|x\|_S}{1 + \|x\|_S} .$$

Proof We start by observing that for any $y \in S$, $\|y\|_S \leq 1$. For any $x \in S$, consider any chord of S that intersects x, and let p, q be the endpoints of this chord. Notice that $\|p\|_S = \|q\|_S = 1$ and using the triangle inequality,

$$\|p - x\|_S \leq \|x\|_S + \|p\|_S \quad \text{and} \quad \|q\|_S \leq \|q - x\|_S + \|x\|_S$$

Thus,

$$\frac{\|q - x\|_S}{\|p - x\|_S} \geq \frac{\|q\|_S - \|x\|_S}{\|x\|_S + \|p\|_S} = \frac{1 - \|x\|_S}{1 + \|x\|_S}.$$

Finally, the lower bound is achieved by the chord that passes through x and the origin. □

The next proposition presents properties of the symmetry function under basic set operations on S.

Proposition 3 *Let S, $T \subset \mathbb{R}^n$ be convex bodies, and let $x \in S$ and $y \in T$. Then:*

1. *(Superminimality under intersection) If $x \in S \cap T$,*

$$\mathbf{sym}(x, S \cap T) \geq \min\{\mathbf{sym}(x, S), \mathbf{sym}(x, T)\} \tag{2.5}$$

2. *(Superminimality under Minkowski sums)*

$$\mathbf{sym}(x + y, S + T) \geq \min\{\mathbf{sym}(x, S), \mathbf{sym}(y, T)\} \tag{2.6}$$

3. *(Invariance under polarity)*

$$\mathbf{sym}(0, S - x) = \mathbf{sym}(0, (S - x)^\circ) \tag{2.7}$$

4. *(Minimality under Cartesian product)*

$$\mathbf{sym}((x,y), S \times T) = \min\{\mathbf{sym}(x,S), \mathbf{sym}(y,T)\} \tag{2.8}$$

5. *(Lower bound under affine transformation) Let $A(\cdot)$ be an affine transformation. Then*

$$\mathbf{sym}(A(x), A(S)) \geq \mathbf{sym}(x,S) \tag{2.9}$$

with equality if $A(\cdot)$ is invertible.

Proof To prove 2.5, without loss of generality, we can translate the sets and suppose that $x = 0$. Let $\alpha = \min\{\mathbf{sym}(0,S), \mathbf{sym}(0,T)\}$. Then $-\alpha S \subset S$, $-\alpha T \subset T$ which implies

$$-\alpha(S \cap T) = -\alpha S \cap -\alpha T \subset S \cap T,$$

and (2.5) is proved.

To prove (2.6), again, without loss of generality, we can translate both sets and suppose that $x = y = 0$, and define $\alpha = \mathbf{sym}(0,S)$ and $\beta = \mathbf{sym}(0,T)$. By definition, $-\alpha S \subset S$ and $-\beta T \subset T$. Then it follows trivially that

$$-\alpha S - \beta T \subset (S+T)$$

Replacing α and β by the minimum between them, the result follows.

In order to prove (2.7), we can assume $x = 0$, then

$$\mathbf{sym}(0,S) = \alpha \Rightarrow -\alpha S \subseteq S.$$

Assuming $\mathbf{sym}(0,S^\circ) < \alpha$, there exist $\bar{y} \in S^\circ$ such that $-\alpha\bar{y} \notin S^\circ$.
Thus, there exists $x \in S$, $-\alpha\bar{y}^T x > 1$. However, since $-\alpha x \in -\alpha S \subseteq S$, then

$$-\alpha\bar{y}^T x = \bar{y}^T(-\alpha x) \leq 1 \text{ , since } \bar{y} \in S^\circ,$$

which is a contradiction. Thus

$$\mathbf{sym}(0, S) \leq \mathbf{sym}(0, S^{\circ}) \leq \mathbf{sym}(0, S^{\circ\circ}) = \mathbf{sym}(0, S).$$

Equality (2.8) is left as a simple exercise.

To prove inequality (2.9), we can assume that $A(\cdot)$ is a linear operator and that $x = 0$ (since $\mathbf{sym}(x, S)$ is invariant under translation), and suppose that $\alpha < \mathbf{sym}(x, S)$. Then, $-\alpha S \subseteq S$ which implies that $A(-\alpha S) \subseteq A(S)$. Since $A(\cdot)$ is a linear operator, $A(-\alpha S) = -\alpha A(S) \subseteq A(S)$. It is straightforward to show that equality holds in (2.9) when $A(\cdot)$ is invertible. □

Remark 1 *Unlike the case of affine transformation,* $\mathbf{sym}(x, S)$ *is not invariant under projective transformation. For instance, let* $S = [-1, 1] \times [-1, 1]$ *be unit cube, for which* $\mathbf{sym}(S) = 1$*, and consider the projective transformation that maps* $x \in \mathbb{R}^2$ *to* $x/(1 + x_2/3) \in \mathbb{R}^2$*. Then, the symmetric set* S *will be mapped to the trapezoid*

$$T := \mathbf{conv}\left\{\left(\frac{3}{4}, \frac{3}{4}\right), \left(-\frac{3}{4}, \frac{3}{4}\right), \left(-\frac{3}{2}, -\frac{3}{2}\right), \left(\frac{3}{2}, -\frac{3}{2}\right)\right\},$$

for which $\mathbf{sym}(T) < 1$*. This lack of invariance is used in Charter 3 in the development of a methodology designed to improve the symmetry of a point in a set using a projective transformation.*

2.2 Geometric properties

Whereas there always exists an n-rounding of a convex body $S \subset \mathbb{R}^n$, a symmetric convex body S possesses some even more powerful geometric properties, for example there exists a $\sqrt{n}$-rounding of S when S is symmetric, see [30]. The geometric flavor of the definition of the symmetry function in (2.4) suggests that $\mathbf{sym}(\cdot, S)$ is connected to extensions of these geometric properties and gives rise to new properties as well; these properties are explored and developed in this section. We examine volumes of intersections of S with halfspaces and halfplanes that cut through $x \in S$ in Section

2.2.1, notions of distance and symmetry in Section 2.2.2, set approximation results in Section 2.2.3, and results on probability and symmetry in Section 2.2.4.

2.2.1 Volumes and symmetry

We start with two theorems that connect $\mathbf{sym}(x, S)$ to bounds on the n-dimensional volume of the intersection of S with a halfspace cut through x, and with the $(n-1)$-dimensional volume of the intersection of S with a hyperplane passing through S. Similar results have been extensively used in the literature. For example, if S is symmetric around some point x^*, it is clear that the intersection of S with a halfspace cut through x^* contains exactly one half of the volume of S. Moreover, it is well known that a halfspace cut through the center of mass generates a set with at least $1/\mathbf{e}$ of the original volume, and this fact has been utilized in [6] to develop theoretically efficient probabilistic methods for solving convex optimization problems.

Let $v \in \mathbb{R}^n$, $v \neq 0$ be given, and for all $x \in S$ define $H(x) := \{z \in S : v^T z = v^T x\}$ and $H^+(x) := \{z \in S : v^T z \leq v^T x\}$. Also let $\mathbf{Vol}_n(\cdot)$ denotes the volume measure on $\mathbb{R}^n$. We have:

Theorem 3 *Under Assumption A, if $x \in S$, then*

$$\frac{\mathbf{sym}(x,S)^n}{1+\mathbf{sym}(x,S)^n} \leq \frac{\mathbf{Vol}_n(H^+(x))}{\mathbf{Vol}_n(S)} \leq \frac{1}{1+\mathbf{sym}(x,S)^n} \ . \tag{2.10}$$

Proof Without loss of generality, assume that x is the origin and $\alpha = \mathbf{sym}(x, S)$. Define $K_1 = H^+(x)$ and $K_2 = S \backslash K_1$. Clearly, $\mathbf{Vol}_n(K_1) + \mathbf{Vol}_n(K_2) = \mathbf{Vol}_n(S)$. Notice that $-\alpha K_2 \subset K_1$ and $-\alpha K_1 \subset K_2$. Therefore

$$\mathbf{Vol}_n(S) \geq \mathbf{Vol}_n(K_1) + \mathbf{Vol}_n(-\alpha K_1) = \mathbf{Vol}_n(K_1)(1+\alpha^n)$$

which proves the second inequality. The first inequality follows easily from

$$\mathbf{Vol}_n(S) = \mathbf{Vol}_n(K_1) + \mathbf{Vol}_n(K_2) \leq \mathbf{Vol}_n(K_1) + \frac{\mathbf{Vol}_n(K_1)}{\alpha^n} \ .$$

□

For the next theorem, define the function $f(x) = \mathbf{Vol}_{n-1}(H(x))^{1/(n-1)}$ for all $x \in S$.

Theorem 4 *Under Assumption A, for every point $x \in S$,*

$$\frac{f(x)}{\max_{y \in S} f(y)} \geq \frac{2\mathbf{sym}(x,S)}{1+\mathbf{sym}(x,S)} \ . \tag{2.11}$$

Proof Let $\alpha = \mathbf{sym}(x,S)$ and let y^* satisfy $y^* \in \arg\max_y f(y)$. Note that

$$x + \alpha(x - H(y^*)) \subset S \ ,$$

and the set on the left in this inclusion passes through $x+\alpha(x-y^*)$, and so $x+\alpha(x-H(y^*)) \subset H(x+\alpha(x-y^*))$. Next, recall that the $(n-1)$-dimensional volume of a set S is invariant under translations and $\mathbf{Vol}_{n-1}(aS) = a^{n-1}\mathbf{Vol}_{n-1}(S)$ for any set S and positive scalar a. Therefore

$$\begin{aligned} \alpha f(y^*) &= (\mathbf{Vol}_{n-1}(x+\alpha(x-H(y^*))))^{1/(n-1)} \\ &\leq (\mathbf{Vol}_{n-1}(H(x+\alpha(x-y^*))))^{1/(n-1)} \\ &= f(x+\alpha(x-y^*)) \ . \end{aligned} \tag{2.12}$$

Note that we can write

$$x = \frac{\alpha}{1+\alpha}y^* + \frac{1}{1+\alpha}(x+\alpha(x-y^*)) \ .$$

where $x+\alpha(x-y^*) \in S$.

Noting that $f(\cdot)$ is concave (this follows from the Brunn-Minkowski inequality [20]), we have:

$$\begin{aligned} f(x) &\geq \frac{\alpha}{1+\alpha}f(y^*) + \frac{1}{1+\alpha}f(x+\alpha(x-y^*)) \\ &\geq \frac{\alpha}{1+\alpha}f(y^*) + \frac{\alpha}{1+\alpha}f(y^*) \\ &= \frac{2\alpha}{1+\alpha}f(y^*) \ , \end{aligned}$$

where the second inequality follows from (2.12). □

Remark 2 *We conjecture that any symmetry point x^* satisfies*

$$\frac{f(x^*)}{\max_{y \in S} f(y)} \geq \frac{2}{3} .$$

2.2.2 Distance and symmetry

If S is a symmetric convex body, then it is a straightforward exercise to show that the symmetry point of S is unique. Furthermore, if S is nearly symmetric, intuition suggests that two points in S with high symmetry values cannot be too far apart. The two theorems in this subsection quantify this intuition. Given $x, y \in S$ with $x \neq y$, let $p(x,y), q(x,y)$ be the pair of endpoints of the chord in S passing through x and y, namely:

$$\begin{aligned} p(x,y) &= x + s(x-y) \in \partial S \quad \text{where } s \text{ is a maximal scalar} \\ q(x,y) &= y + t(y-x) \in \partial S \quad \text{where } t \text{ is a maximal scalar.} \end{aligned} \tag{2.13}$$

Theorem 5 *Under Assumption A, let $\|\cdot\|$ be any norm on $\mathbb{R}^n$. For any $x, y \in S$ satisfying $x \neq y$, let $\alpha = \mathbf{sym}(x, S)$ and $\beta = \mathbf{sym}(y, S)$. Then:*

$$\|x - y\| \leq \left(\frac{1 - \alpha\beta}{1 + \alpha + \beta + \alpha\beta} \right) \|p(x,y) - q(x,y)\| .$$

Proof For convenience let us denote the quantities $p(x,y), q(x,y)$ by p, q, and note that the chord from p to q contains, in order, the points p, x, y, and q. It follows from the symmetry values of x, y that

$$\|p-x\| \geq \alpha\|q-x\| = \alpha(\|y-x\|+\|q-y\|) \text{ and } \|q-y\| \geq \beta\|p-y\| = \beta(\|y-x\|+\|p-x\|) .$$

Multiplying the first inequality by $1 + \beta$, the second inequality by $1 + \alpha$, adding the

result and rearranging yields:

$$(1+\alpha+\beta+\alpha\beta)\|x-y\| \leq (1-\alpha\beta)(\|p-x\|+\|x-y\|+\|q-y\|) = (1-\alpha\beta)\|p-q\| ,$$

which yields the desired result. □

Another relative measure of distance is the "cross-ratio distance" with respect to S. Let $x, y \in S$, $x \neq y$, be given and let s, t be as defined in (2.13); the cross-ratio distance is given by:

$$d_S(x,y) := \frac{(1+t+s)}{ts} .$$

Theorem 6 *Under Assumption A, for any $x, y \in S$, $x \neq y$, let s, t be as defined in (2.13). Then*

$$d_S(x,y) \leq \frac{1}{\mathbf{sym}(x,S) \cdot \mathbf{sym}(y,S)} - 1 .$$

Proof Let $\alpha = \mathbf{sym}(x,S)$ and $\beta = \mathbf{sym}(y,S)$. By definition of symmetry, $t \geq \beta(1+s)$ and $s \geq \alpha(1+t)$. Then

$$\begin{aligned} d_S(x,y) &= \frac{(1+t+s)}{ts} \leq \frac{(1+t+s)}{\alpha(1+s)\beta(1+t)} \\ &= \frac{1}{\alpha\beta}\frac{(1+t+s)}{(1+s+t+st)} = \left(\frac{1}{\alpha\beta}\right)\frac{1}{1+\frac{1}{d_S(x,y)}}. \end{aligned} \tag{2.14}$$

Thus $d_S(x,y) \leq \frac{1}{\alpha\beta} - 1$. □

We end this subsection with a comment on a question posed by Hammer in [27]: what is the upper bound on the difference between $\mathbf{sym}(S)$ and $\mathbf{sym}(x^c, S)$, where x^c is the centroid (center of mass) of S? It is well known that $\mathbf{sym}(x^c, S) \geq 1/n$, see [27], and it follows trivially from the Löwner-John theorem that $\mathbf{sym}(S) \geq 1/n$ as well. Now let S be the Euclidean half-ball: $S := \{x \in \mathbb{R}^n : \langle x, x\rangle \leq 1, x_1 \geq 0\}$. It is an easy exercise to show that the unique symmetry point of S is $x^* = (\sqrt{2}-1)e^1$ and that $\mathbf{sym}(S) = \frac{1}{\sqrt{2}}$, and so in this case $\mathbf{sym}(S)$ is a constant independent of the dimension n. On the other hand, $\mathbf{sym}(x^c, S) = \Omega\left(\frac{1}{\sqrt{n}}\right)$ (see [3]), and so for this class of instances the symmetry of the centroid is substantially less than the symmetry of

the set for large n. For an arbitrary convex body S, note that in the extreme cases where $\mathbf{sym}(S) = 1$ or $\mathbf{sym}(S) = 1/n$ the difference between $\mathbf{sym}(S)$ and $\mathbf{sym}(x^c, S)$ is zero; we conjecture that tight bounds on this difference are only small when $\mathbf{sym}(S)$ is either very close to 1 or very close to $1/n$.

2.2.3 Set-approximation and symmetry

In this subsection we examine the approximation of the convex set S by another convex set P. We say that P is a β-approximation of S if there exists a point $x \in S$ such that $\beta P \subset S - x \subset P$. In the case when P is an ellipsoid centered at the origin, then the statement "P is a β-approximation of S" is equivalent to "βP provides a $\frac{1}{\beta}$-rounding of S." We examine the interrelationship between the symmetry function and bounds on β-approximations for S in the following three theorems.

A classical example of β-approximation is given by the Löwner-John theorem [30], which guarantees a $1/\sqrt{n}$-approximation for a symmetric convex body and a $1/n$-approximation for general convex body using ellipsoids. Unfortunately, the theorem does not provide more precise bounds for case when S is nearly symmetric, i.e., $\mathbf{sym}(S) = 1 - \varepsilon$ for ε small. This is partially rectified in the fourth result of this subsection, Theorem 10.

Theorem 7 *Under Assumption A, let P be a convex body that is a β-approximation of S, and suppose that $\mathbf{sym}(0, P) = \alpha$. Then, $\mathbf{sym}(S) \geq \beta\alpha$.*

Proof By definition we have $\beta P \subset S - x \subset P$ for some $x \in S$. Since $\mathbf{sym}(\cdot, \cdot)$ is invariant under translations, we can assume that $x = 0$. Since $\mathbf{sym}(0, P)$ is invariant under nonzero scalings of P, we have

$$-\alpha\beta S \subset -\alpha\beta P \subset \beta P \subset S \ .$$

□

Theorem 8 *Under Assumption A, suppose that $x \in \mathbf{int}\ S$. Then there exists an ellipsoid E centered at 0 such that*

$$E \subset S - x \subset \left(\frac{\sqrt{n}}{\mathbf{sym}(x, S)} \right) E \ . \tag{2.15}$$

Proof Suppose without loss of generality that $x = 0$ (otherwise we translate S), and let $\alpha = \mathbf{sym}(0, S)$. Clearly, $-\alpha S \subset S$, and $\alpha S \subset S$. Consider a $\sqrt{n}$-rounding E of $S \cap (-S)$. Then $\alpha S \subset S \cap (-S) \subset \sqrt{n} E \subset \sqrt{n} S$. $\square$

Theorem 9 *Let $\| \cdot \|$ be any norm on $\mathbb{R}^n$, and let $B(x, r)$ denote the ball centered at x with radius r. Under Assumption A, suppose that*

$$B(x, r) \subset S \subset P \subset S + B(0, \delta) \tag{2.16}$$

for some r and δ with $0 < \delta < r$. Then

$$\left(1 - \frac{\delta}{r}\right) \leq \frac{\mathbf{sym}(x, S)}{\mathbf{sym}(x, P)} \leq \left(\frac{1}{1 - \delta/r}\right).$$

Proof Let $\alpha = \mathbf{sym}(x, P)$. Consider any chord of P that passes through x, dividing the chord into two segments. Assume that the length of one segment is Δ, then the length of the other segment must be at most Δ/α. It then follows that the length of the first segment of this chord in S must be at least $\Delta - \delta$, while the length of the second segment of this chord in S must be at most Δ/α. Since these inequalities hold for any chord, it follows that

$$\mathbf{sym}(x, S) \geq \frac{\Delta - \delta}{\Delta/\alpha} = \alpha \left(1 - \frac{\delta}{\Delta}\right) \geq \alpha \left(1 - \frac{\delta}{r}\right) \tag{2.17}$$

where the last inequality follows since $\Delta \geq r$, thereby showing that $\mathbf{sym}(x, S) \geq \mathbf{sym}(x, P)\left(1 - \frac{\delta}{r}\right)$. Note also that:

$$B(x, r) \subset P \subset S + B(0, \delta) \subset P + B(0, \delta) \ .$$

Letting P play the role of S in (2.16) and $S + B(0, \delta)$ play the role of P in (2.16), it

also follows from (2.17) that

$$\mathbf{sym}(x, P) \geq \mathbf{sym}(x, S + B(0, \delta)) \left(1 - \frac{\delta}{r}\right) .$$

However, using the superminimality of of $\mathbf{sym}(\cdot, \cdot)$ under Minkowski sums (2.6) of Theorem 3, we have

$$\mathbf{sym}(x, S + B(0, \delta)) \geq \min\{\mathbf{sym}(x, S), \mathbf{sym}(0, B(0, \delta))\} = \mathbf{sym}(x, S) ,$$

which when combined with the previous inequality completes the proof. □

The center x^L of the minimum-volume ellipsoid E containing S is called the Löwner-John center of S, and John showed in [30] that E provides a $\sqrt{n}$-rounding of S in the case when S is symmetric and an n-rounding of S when S is not symmetric. The following theorem provides a sharpening of this result:

Theorem 10 *Under Assumption A, let E be the minimum volume ellipsoid containing S, and let x^L be the Löwner-John center of S. Then E provides a $\sqrt{\frac{n}{\mathbf{sym}(x^L, S)}}$-rounding of S.*

Remark 3 *It follows from Theorem 10 that*

$$\mathbf{sym}(x^L, S) \geq \sqrt{\frac{\mathbf{sym}(x^L, S)}{n}}$$

and hence $\mathbf{sym}(x^L, S) \geq 1/n$. This in turn yields the Löwner-John result [23] that the rounding in the theorem is an n-rounding, and hence $\mathbf{sym}(S) \geq \mathbf{sym}(x^L, S) \geq 1/n$. Noting that when S is symmetric the Löwner-John center must also be the symmetry point of S, it also follows from Theorem 10 that S admits a $\sqrt{n}$-rounding when $\mathbf{sym}(S) = 1$.

Remark 4 *Theorem 8 is valid for every point in S and Theorem 10 focuses on the Löwner-John center. We conjecture that Theorem 10 can be strengthened to prove the existence of a $\left(\sqrt{\frac{n}{\mathbf{sym}(S)}}\right)$-rounding of S.*

The proof of Theorem 10 is based in part on ideas communicated by Earl Barnes [4] in 1998. We start with the following two elementary propositions:

Proposition 4 *Let $w_1, \ldots, w_k$ be scalars and define $w_{\min}, w_{\max}$ to be the smallest and largest values among these scalars. For any $p \in \mathbb{R}^k$ satisfying $p \geq 0$ and $e^T p = 1$ define $\mu = p^T w$ and $\sigma^2 = \sum_{i=1}^k p_i (w_i - \mu)^2$. Then $(w_{\max} - \mu)(\mu - w_{\min}) \geq \sigma^2$.*

Proof Clearly, $\sum_{i=1}^k p_i (w_{\max} - w_i)(w_i - w_{\min}) \geq 0$. Therefore $\mu w_{\max} + \mu w_{\min} - \sum_{i=1}^k p_i w_i^2 - w_{\min} w_{\max} \geq 0$. It then follows that $(w_{\max} - \mu)(\mu - w_{\min}) = \mu w_{\max} + \mu w_{\min} - \mu^2 - w_{\min} w_{\max} \geq \sum_{i=1}^k p_i w_i^2 - \mu^2 = \sigma^2$. $\square$

Proposition 5 *Let $y^1, \ldots y^k \in \mathbb{R}^n$ be given, let $p \in \mathbb{R}^k$ satisfy $p \geq 0$ and $e^T p = 1$, and suppose that $\sum_{i=1}^n p_i y^i = 0$ and $\sum_{i=1}^k p_i y^i (y^i)^T = \frac{1}{n} I$. Then for any $\bar{b} \in \mathbb{R}^n$ with $\|\bar{b}\|_2 = 1$ it holds that*

$$\max_{i=1,\ldots,k} \bar{b}^T y^i \geq \sqrt{\frac{\mathbf{sym}(0, \mathbf{conv}(\{y^i\}_{i=1}^k))}{n}} \, .$$

Proof Let $\bar{b} \in \mathbb{R}^n$ satisfying $\|\bar{b}\|_2 = 1$ be given, and define $w_i = \bar{b}^T y^i$. Then

$$\mu = \sum_{i=1}^k p_i w_i = \sum_{i=1}^k p_i \bar{b}^T y^i = \bar{b}^T \left(\sum_{i=1}^k p_i y^i \right) = 0$$

and

$$\sigma^2 = \sum_{i=1}^k p_i (w_i - \mu)^2 = \sum_{i=1}^k p_i w_i^2 = \sum_{i=1}^k p_i \bar{b}^T y^i (y^i)^T \bar{b} = \frac{1}{n} \bar{b}^T I \bar{b} = \frac{1}{n} \, .$$

It then follows from Proposition 4 that $(\max_i w_i)(-\min_i w_i) = (\max_i w_i - \mu)(\mu - \min_i w_i) \geq \sigma^2 = \frac{1}{n}$. Let $\alpha := \mathbf{sym}(0, \mathbf{conv}(\{y^i\}_{i=1}^k))$, and notice that $-\min_i \bar{b}^T y^i \leq \frac{1}{\alpha} \max_i \bar{b}^T y^i$. Therefore

$$\frac{1}{n} \leq \max_i w_i \left(-\min_i w_i \right) \leq \max_i w_i \left(\frac{1}{\alpha} \max_i w_i \right) = \frac{(\max_i w_i)^2}{\alpha}$$

from which the result readily follows. $\square$

Proof *of Theorem 10:* We first suppose that S is the convex hull of finitely many points, and we write $S = \mathbf{conv}(\{v^i\}_{i=1}^k)$. The minimum volume ellipsoid containing

S is obtained using the solution of the following optimization problem:

$$\begin{array}{ll} \min_{Q,c} & -\ln\det Q \\ \text{s.t.} & (v^i-c)^TQ(v^i-c)\le 1, \quad i=1,\dots,k \\ & Q\succeq 0\ . \end{array} \tag{2.18}$$

If Q, c solves (2.18) then $E^O := \{x \in \mathbb{R}^n : (x-c)^TQ(x-c) \le 1\}$ is the minimum volume ellipsoid containing S and c is the Löwner-John center. Letting $E^I := \{x \in \mathbb{R}^n : (x-c)^TQ(x-c) \le \frac{\alpha}{n}\}$ where $\alpha := \mathbf{sym}(c, S)$, we need to show that $E^I \subset S$. Equivalently, for every $b \in \mathbb{R}^n$ we need to show that

$$\max\{b^Tx : x \in E^I\} \le \max\{b^Tx : x \in S\}\ .$$

The KKT conditions for (2.18) are necessary and sufficient, see John [30], and can be written as:

$$\begin{array}{l} -Q^{-1}+\sum_{i=1}^k \lambda_i(v^i-c)(v^i-c)^T = 0 \\ \sum_{i=1}^k \lambda_i Q(v^i-c) = 0 \\ \lambda_i \ge 0, \quad i=1,\dots,k \\ (v^i-c)^TQ(v^i-c)\le 1, \quad i=1,\dots,k \\ \lambda_i(v^i-c)^TQ(v^i-c) = \lambda_i, \quad i=1,\dots,k \\ Q\succ 0\ . \end{array}$$

Defining $y^i = Q^{1/2}(v^i-c)$ and $p_i = \frac{\lambda_i}{n}$ we have $p \ge 0$, and using the KKT conditions we obtain:

$$\begin{array}{rcl} n = \text{trace}(I) & = & \text{trace}(Q^{1/2}Q^{-1}Q^{1/2}) \\ & = & \sum_{i=1}^k \lambda_i \text{trace}\left(Q^{1/2}(v^i-c)(v^i-c)^TQ^{1/2}\right) \\ & = & \sum_{i=1}^k \lambda_i \text{trace}\left((v^i-c)^TQ^{1/2}Q^{1/2}(v^i-c)\right) \\ & = & \sum_{i=1}^k \lambda_i (v^i-c)^TQ^{1/2}Q^{1/2}(v^i-c) \\ & = & \sum_{i=1}^k \lambda_i = e^T\lambda\ , \end{array}$$

and it follows that $e^T p = \frac{e^T \lambda}{n} = 1$. Furthermore,

$$\sum_{i=1}^{k} p_i y^i = \frac{1}{n} Q^{-1/2} \sum_{i=1}^{k} \lambda_i Q(v^i - c) = 0$$

and

$$\sum_{i=1}^{k} p_i y^i (y^i)^T = \sum_{i=1}^{k} \frac{\lambda_i}{n} Q^{1/2}(v_i - c)(v_i - c)^T Q^{1/2} = \frac{1}{n} I \ .$$

For any $b \in \mathbb{R}^n$, $b \neq 0$, define $\bar{b} := \frac{Q^{-1/2} b}{\sqrt{b^T Q^{-1} b}}$ and note that $\|\bar{b}\|_2 = 1$. Then $p, y^1, \ldots, y^k$, and $\bar{b}$ satisfy the hypotheses of of Proposition 5, and so

$$\begin{aligned} \max\{b^T x : x \in S\} &= \max_i b^T v^i \\ &= b^T c + \sqrt{b^T Q^{-1} b}(\max_i \bar{b}^T Q^{1/2}(v^i - c)) \\ &= b^T c + \sqrt{b^T Q^{-1} b}(\max_i \bar{b}^T y^i) \\ &\geq b^T c + \sqrt{\tfrac{\alpha}{n}} \sqrt{b^T Q^{-1} b} \ , \end{aligned}$$

where the inequality is from Proposition 5, and we use the fact that

$$\mathbf{sym}(0, \mathbf{conv}(\{v^i\}_{i=1}^k)) = \mathbf{sym}(0, \mathbf{conv}(\{y^i\}_{i=1}^k))$$

which follows from the invariance of $\mathbf{sym}(\cdot, \cdot)$ under invertible affine transformation, see (5.) of Theorem 2. On the other hand we have:

$$\max\left\{b^T x : x \in E^I\right\} = \max\left\{b^T x : (x - c)^T Q (x - c) \leq \frac{\alpha}{n}\right\} = b^T c + \sqrt{\frac{\alpha}{n}} \sqrt{b^T Q^{-1} b} \ ,$$

which then yields $\max\{b^T x : x \in E^I\} \leq \max\{b^T x : x \in S\}$, proving the result under the hypothesis that S is the convex hull of finitely many points.

Finally, suppose S is not the convex hull of finitely many points. For any $\delta > 0$ there is a polytope P_δ that approximates S in the sense that $S \subset P_\delta \subset S + B(0, \delta)$, where $B(0, \delta)$ is the ball of radius δ centered at 0. Limiting arguments can then be used to show the result by taking a limiting sequence of polytopes P_δ as $\delta \to 0$ and noticing from Theorem 9 that $\lim_{\delta \to 0} \mathbf{sym}(0, P_\delta) = \mathbf{sym}(0, S)$. $\square$

We close this subsection by discussing a norm closely related to the symmetry function that was also used in [13]. Without loss of generality, assume that $x^* = 0$ is a symmetry point of S and define the following norm associated with S:

$$\|x\|_S = \min_t \{t : x \in t(S \cap -S)\} , \tag{2.19}$$

and let $B_S(c, r)$ denote the ball of radius r centered at c using the norm defined in (2.19).

Lemma 1 *Under Assumption A, suppose that $x^* = 0$ is a symmetry point of S. Then*

$$B_S(0, 1) \subset S \subset B_S(0, 1/\mathbf{sym}(S)).$$

Proof By construction, $B_S(0, 1) = S \cap -S \subset S$. For the second inclusion, observe that $-\mathbf{sym}(S)S \subset S$, which then implies that $S \subset -\frac{1}{\mathbf{sym}(S)}S$. Therefore $S \subset \frac{1}{\mathbf{sym}(S)}(S \cap -S)$. □

Remark 5 *The norm defined by (2.19) induces the best β-approximation among all norms in $\mathbb{R}^n$. That is, $(0, 1, 1/\mathbf{sym}(S), \|\cdot\|_S)$ solves the following optimization problem*

$$\min_{x,r,R,\|\cdot\|} \left\{ \frac{R}{r} : B_{\|\cdot\|}(x, r) \subset S \subset B_{\|\cdot\|}(x, R) \right\}$$

Proof Suppose that there exists a norm $\|\cdot\|$, r, R, and x, such that

$$B_{\|\cdot\|}(x, r) \subset S \subset B_{\|\cdot\|}(x, R)$$

and $\frac{R}{r} < \frac{1}{\mathbf{sym}(S)}$. Using Theorem 7, we have $\mathbf{sym}(x, S) \geq \frac{r}{R} > \mathbf{sym}(S)$, a contradiction. □

2.2.4 Probability and symmetry

This subsection contains two results related to symmetry and probability. To set the stage for the first result, suppose that X is a random vector uniformly distributed on

the given convex body $S \subset \mathbb{R}^n$. Theorem 11 gives lower bounds on the expected value of $\mathbf{sym}(X, S)$ and on the probability that $\mathbf{sym}(X, S)$ will be larger than a constant fraction $1/M$ of $\mathbf{sym}(S)$. Roughly speaking, Theorem 11 states that it is likely that $\mathbf{sym}(X, S)$ is relatively large. The second result, Theorem 12, is an extension of Anderson's Lemma [1] concerning the integral of a nonnegative logconcave even function on S, and has many statistical applications.

Theorem 11 *Under Assumption A, let X be a random vector uniformly distributed on S. Then*

(i) $\mathsf{E}[\mathbf{sym}(X, S)] \geq \dfrac{\mathbf{sym}(S)}{2n+1}$

(ii) For any $M \geq 1$, $\mathsf{P}\left(\mathbf{sym}(X, S) \geq \dfrac{\mathbf{sym}(S)}{M}\right) \geq \left(1 - \dfrac{2}{M+1}\right)^n$

(iii) Among symmetric sets S, $\mathsf{E}[\mathbf{sym}(X, S)]$ *and* $\mathsf{P}\left(\mathbf{sym}(X, S) \geq \frac{1}{M}\right)$ *are functions only of the dimension n and are independent of the specific set S, and satisfy:*

(iii.a) $\mathsf{E}[\mathbf{sym}(X, S)] \leq \dfrac{1}{2(n+1)} + \dfrac{1}{(n+1)(n+2)}$

(iii.b) For any $M \geq 1$, $\mathsf{P}\left(\mathbf{sym}(X, S) \geq \dfrac{1}{M}\right) = \left(1 - \dfrac{2}{M+1}\right)^n$.

Proof Without loss of generality we assume for convenience that $x^* = 0$ is a symmetry point of S. Let $t \in [0, 1]$. For any $x \in tS$, consider any chord of S that intersects x, and let p, q be the endpoints of this chord. Note that $\|p\|_S \leq 1/\mathbf{sym}(S)$ and $\|x\|_S \leq t/\mathbf{sym}(S)$, where $\|\cdot\|_S$ is the norm defined in (2.19). Also, it follows from basic convexity that $tS + (1-t)B_S(0,1) \subset tS + (1-t)S \subset S$, where $B_S(0,1)$ is the unit ball centered at the origin for the norm $\|\cdot\|_S$. This then implies that if $x \in tS$ and $q \in \partial S$ then $\|q - x\|_S \geq 1 - t$. Therefore

$$\frac{\|q - x\|_S}{\|p - x\|_S} \geq \frac{1-t}{\|p\|_S + \|x\|_S} \geq \frac{1-t}{1/\mathbf{sym}(S) + t/\mathbf{sym}(S)},$$

which implies that

$$\mathbf{sym}(x, S) \geq \mathbf{sym}(S)\frac{1-t}{1+t} \ . \tag{2.20}$$

Now suppose that X is a random vector uniformly distributed on S, and consider the random variable $t(X)$ defined uniquely by the inclusion $X \in \partial(tS)$. Then

$$P(t(X) \leq t) = P(X \in tS) = \frac{\mathbf{Vol}(tS)}{\mathbf{Vol}(S)} = t^n \ ,$$

which implies that the density of $t(X)$ is given simply by $f(t) = nt^{n-1}$. Therefore using (2.20) we have:

$$\begin{aligned}
\mathsf{E}[\mathbf{sym}(X, S)] &\geq \mathsf{E}\left[\mathbf{sym}(S)\frac{1-t(X)}{1+t(X)}\right] \\
&= \int_0^1 \mathbf{sym}(S)\frac{1-t}{1+t} nt^{n-1} dt \\
&= n\mathbf{sym}(S)\int_0^1 \frac{1-t}{1+t} t^{n-1} dt \\
&\geq n\mathbf{sym}(S)\int_0^1 t^{n-1}\left(1-\sqrt{t}\right) dt \\
&= \tfrac{\mathbf{sym}(S)}{2n+1} \ ,
\end{aligned}$$

where the second inequality follows from the observation that $\frac{1-t}{1+t} \geq 1 - \sqrt{t}$ for $t \in [0, 1]$. This proves (i).

To prove (ii), let $M \geq 1$ be given and define $\bar{t} := 1 - \frac{2}{M+1}$ and note the relationship

$$\frac{1-\bar{t}}{1+\bar{t}} = \frac{1}{M} \ .$$

Since $\{x \in S : x \in \bar{t}S\} \subset \left\{x \in S : \mathbf{sym}(x, S) \geq \frac{\mathbf{sym}(S)}{M}\right\}$ from (2.20), we have:

$$\mathsf{P}\left(\mathbf{sym}(X, S) \geq \frac{\mathbf{sym}(S)}{M}\right) \geq \mathsf{P}(X \in \bar{t}S) = (\bar{t})^n \ ,$$

which establishes (ii). To prove (iii) notice from Proposition 2 that (2.20) holds with equality in this case, whereby the above derivations yield $\mathsf{E}[\mathbf{sym}(X, S)] = n\int_0^1 \frac{1-t}{1+t} t^{n-1} dt$ and $\mathsf{P}\left(\mathbf{sym}(X, S) \geq \frac{1}{M}\right) = (\bar{t})^n$, which are functions of n and are

independent of S, thus showing (iii) and $(iii.b)$. Noting that $\frac{1-t}{1+t} \leq 1 - \frac{3}{2}t + \frac{1}{2}t^2$ for $t \in [0,1]$, we obtain in the symmetric case that

$$\mathsf{E}[\mathbf{sym}(X,S)] \leq n \int_0^1 \left(1 - \frac{3}{2}t + \frac{1}{2}t^2\right) t^{n-1} dt = \frac{1}{2(n+1)} + \frac{1}{(n+1)(n+2)} ,$$

which shows $(iii.a)$. □

Corollary 1 *Let X be a random vector uniformly distributed on $S \subset \mathbb{R}^n$ for $n \geq 2$. Then*

$$\mathsf{P}\left(\mathbf{sym}(X,S) \geq \frac{\mathbf{sym}(S)}{n}\right) \geq \left(1 - \frac{2}{n+1}\right)^n \geq 1/9 ,$$

and the lower bound goes to $1/(\mathbf{e})^2$ as $n \to \infty$.

The following is an extension of Anderson's Lemma [1], whose proof relies on the Brunn-Minkowski inequality in the symmetric case.

Theorem 12 *Let $S \subset \mathbb{R}^n$ be a compact convex set which contains the origin in its interior, and let $\alpha = \mathbf{sym}(0,S)$. Let $f(\cdot)$ be a nonnegative quasiconcave even function that is Lebesgue integrable. Then for $0 \leq \beta \leq 1$ and any $y \in \mathbb{R}^n$,*

$$\int_S f(x+\beta y)dx \geq \alpha^n \int_S f\left(x + \frac{y}{\alpha}\right) dx . \tag{2.21}$$

Proof We refer to [10] for a proof in the symmetric case $\alpha = 1$. Suppose that $f(\cdot)$ is an indicator function of a set K. This implies that K is convex and $\mathbf{sym}(0,K) = 1$. Therefore:

$$\begin{aligned} \textstyle\int_S f(x+\beta y)dx &\geq \textstyle\int_{S\cap -S} f(x+\beta y)dx \\ &\geq \textstyle\int_{S\cap -S} f(x+y)dx \\ &= \mathbf{Vol}_n((S\cap -S)\cap (K-y)) = \mathbf{Vol}_n((S\cap -S)\cap \alpha(\tfrac{K-y}{\alpha})) \\ &\geq \mathbf{Vol}_n(\alpha S \cap \alpha(K - \tfrac{y}{\alpha})) = \alpha^n \mathbf{Vol}_n(S \cap (K - \tfrac{y}{\alpha})) \end{aligned} \tag{2.22}$$

where the second inequality follows from Anderson's original theorem [1], and the third inequality holds simply because $\alpha S \subset S \cap -S$ and $K \subset \frac{K}{\alpha}$. Thus the result is true for simple quasiconcave even functions, and using standard arguments

of dominated and monotone convergence, the result also holds for all nonnegative quasiconcave even Lebesgue-integrable functions. □

The following corollary shows the potential usefulness of Theorem 12 in probability theory. We note that the density function of a uniformly distributed or an n-dimensional Gaussian random vector with mean $\mu = 0$ satisfies the functional conditions of Theorem 12.

Corollary 2 *Let X be a random variable in $\mathbb{R}^n$ whose density function $f(\cdot)$ is an even quasiconcave function. In addition, let Y be an arbitrary random vector independent of X, and let $\beta \in [0,1]$. If $S \subset \mathbb{R}^n$ is a compact convex set which contains the origin in its interior and $\alpha = \mathbf{sym}(0,S)$, then*

$$P(X + \beta Y \in S) \geq \alpha^n P\left(X + \frac{Y}{\alpha} \in S\right) . \tag{2.23}$$

Proof Noting that α does not depend on Y, we have:

$$\begin{aligned} P(X + \beta Y \in S) &= \int \int_{S-\beta y} f(x)dxdP(y) = \int \int_S f(x - \beta y)dxdP(y) \\ &\geq \alpha^n \int \int_S f(x - \frac{y}{\alpha})dxdP(y) = \alpha^n P(X + \tfrac{Y}{\alpha} \in S) . \end{aligned} \tag{2.24}$$

□

2.3 Characterization of symmetry points via the normal cone

Let $S_{opt}(S)$ denote the set of symmetry points of the convex body S. In this section we provide a characterization of $S_{opt}(S)$. From (2.4) and (2.2) we see that $S_{opt}(S)$ is the x-part of the optimal solution of:

$$\begin{aligned} \mathbf{sym}(S) = \max_{x,\alpha} \quad & \alpha \\ s.t. \quad & \alpha(x - S) \subseteq (S - x) \\ & \alpha \geq 0 . \end{aligned} \tag{2.25}$$

For any given $x \in S$ let $\alpha = \mathbf{sym}(x, S)$. Motivated by the set-containment definition of $\mathbf{sym}(x, S)$ in (2.4), let $V(x)$ denote those points $v \in \partial S$ that are also elements of the set $x + \alpha(x - S)$. We call these points the "touching points" of x in S, namely:

$$V(x) := \partial S \cap (x + \alpha(x - S)) \quad \text{where} \quad \alpha = \mathbf{sym}(x, S) \; . \tag{2.26}$$

Let $N_S(y)$ denote the normal cone map for points $y \in S$. We assemble the union of all normal cone vectors of all of the touching points of x and call the resulting set the "support vectors" of x:

$$SV(x) = \{s \in \mathbb{R}^n : \|s\|_2 = 1 \text{ and } s \in N_S(v) \text{ for some } v \in V(x)\} \; . \tag{2.27}$$

The following characterization theorem essentially states that $x^* \in S$ is a symmetry point of S if and only if the origin is in the convex hull of the support vectors of x:

Theorem 13 *Under Assumption A, let $x^* \in S$. The following statements are equivalent:*

(i) $x^* \in S_{opt}(S)$

(ii) $0 \in \mathbf{conv} SV(x^*)$.

The proof of Theorem 13 we will rely on the following technical result:

Lemma 2 *Suppose that S is a convex body in a Euclidean space and $x \in \mathbf{int}\ S$ and $\alpha \geq 0$. Then $\alpha < \mathbf{sym}(x, S)$ if and only if $\alpha(x - S) \subseteq \mathbf{int}\ (S - x)$.*

Proof ($\Rightarrow$) The case $\alpha = 0$ is trivial. For $\alpha > 0$, since $x \in \mathbf{int}\ S$ and S is a convex body, $\alpha < \mathbf{sym}(x, S)$ implies that

$$\alpha(x - S) \subset \mathbf{sym}(x, S)\mathbf{int}\ (x - S) \subseteq \mathbf{int}\ (S - x) \; .$$

($\Leftarrow$) For a fixed value of α, rearrange the subset system to be: $x + \alpha(x - S) \subset \mathbf{int}\ S$. However, S is a compact set, whereby α can be increased to $\alpha + \varepsilon$ for some small

positive value of ϵ and still maintain $x + (\alpha + \varepsilon)(x - S) \subset \mathbf{int}\ S \subset S$, which by (2.4) is equivalent to $\mathbf{sym}(x, S) \geq \alpha + \varepsilon$. □

The proof of Theorem 13 will also use the following construction:

Lemma 3 *Consider the function* $f(\cdot) : \mathbb{R}^n \to \mathbb{R}$ *defined as*

$$f(x) = \sup_{\substack{y \in \partial S \\ s \in N_S(y) \\ \|s\|_2 = 1}} \langle s, x - y \rangle \ . \tag{2.28}$$

Then

(*i*) $f(\cdot)$ *is convex,*

(*ii*) $f(x) = 0$ *for* $x \in \partial S$,

(*iii*) $f(x) > 0$ *for* $x \notin S$,

(*iv*) $f(x) < 0$ *for* $x \in \mathbf{int}\ S$, *and*

(*v*) $\{s : \|s\|_2 = 1,\ s \in N_S(x)\} \subset \mathbf{conv}\partial f(x)$ *for* $x \in \partial S$.

Proof As the supremum of affine functions, $f(\cdot)$ is convex, which shows (*i*). For $x \in \partial S$, $f(x) \geq 0$. For (y, s) feasible for (2.28), $\langle s, x - y \rangle \leq 0$ for all $x \in S$ by definition of the normal cone, whereby $f(x) = 0$, which shows (*ii*). For $x \in \mathbf{int}\ S$, there exists $\delta > 0$ such that $B_2(x, \delta) \subset S$. Let (y, s) be feasible for (2.28), then $\langle s, x - y \rangle = \langle s, (x + \delta s - y) - \delta s \rangle \leq \langle s, -\delta s \rangle = -\delta$, which then implies that $f(x) \leq -\delta$ and shows (*iv*).

For $x \notin S$, there exists a hyperplane strictly separating x from S. That is, there exists $\bar{s}$ satisfying $\|\bar{s}\|_2 = 1$ such that $\langle \bar{s}, x \rangle > \max_y \{\langle \bar{s}, y \rangle : y \in S\}$, and let $\bar{y}$ be an optimal solution of this problem. Then $(\bar{y}, \bar{s})$ is feasible for (2.28) and it follows that $f(x) \geq \langle \bar{s}, x - \bar{y} \rangle > 0$, showing (*iii*). For $x \in \partial S$ and any $s \in N_S(x)$ satisfying $\|s\|_2 = 1$, it follows that for all w that $f(w) \geq \langle s, w - x \rangle = f(x) + \langle s, w - x \rangle$, thereby showing (*v*). □

Proof *of Theorem 13.* Suppose that $x^* \in S_{opt}(S)$. From (2.4) and Lemma 3 it follows that x^* is a solution together with $\alpha^* := \mathbf{sym}(S)$ of the following optimization problem:

$$\begin{aligned} \mathbf{sym}(S) = \max_{x,\alpha} \quad & \alpha \\ s.t. \quad & f(x - \alpha(y - x)) \leq 0 \quad \text{for all } y \in S \ . \end{aligned} \tag{2.29}$$

The necessary optimality conditions for this problem imply that

$$0 \in \sum_{v \in V(x^*)} \lambda_v s_v$$

where $s_v \in \partial f(v)$ for all v, for some λ satisfying $\lambda \geq 0$, $\lambda \neq 0$. Observe for $v \in \partial S$ and $s \in \partial f(v)$ that $0 \geq f(w) \geq f(v) + \langle s, w - v \rangle = \langle s, w - v \rangle$ for all $w \in S$, which implies that $s \in N_S(v)$, and so

$$0 \in \sum_{v \in V(x^*)} \lambda_v s_v$$

where $s_v \in N_S(v)$ for all v, which implies (ii).

Conversely, suppose that $\alpha^* = \mathbf{sym}(x^*, S)$, and note that for any $v \in V(x^*)$, $0 \notin \partial f(v)$ (otherwise f would be nonnegative which contradicts Assumption A and Lemma 3). Therefore $0 \in \mathbf{conv} SV(x^*)$ implies that $\mathbf{cone} SV(x^*)$ contains a line. Let $y \in S$ be given and define $d := y - x^*$. Since $SV(x^*)$ contains a line, there exists $s \in SV(x^*)$ for which $\langle s, d \rangle \geq 0$. Let v be the touching point corresponding to s, i.e., $v \in V(x^*)$ and $s \in N_S(v)$; then $v \in \partial S$ and $v = x^* - \alpha^*(w - x^*)$ for some $w \in S$ (from (2.4)). From (v) of Lemma 3 we have $s \in \partial f(v)$, whereby $s \in \partial f(v)$. Thus, using the subgradient inequality,

$$\begin{aligned} f(y - \alpha^*(w - y)) &= f(v + (y - \alpha^*(w - y) - v)) \\ &\geq f(v) + \langle s, y - \alpha^*(w - y) - v \rangle \\ &= \langle s, d \rangle (1 + \alpha^*) \geq 0 \ , \end{aligned} \tag{2.30}$$

which shows that $y - \alpha^*(w - y) \notin \textbf{int } S$. This implies that

$$-\alpha^*\Big(S - y\Big) \nsubseteq \textbf{int } \Big(S - y\Big) .$$

Then Lemma 2 implies $\textbf{sym}(y, S) \leq \alpha^*$ for all $y \in S$ proving the optimality of x^*. □

We close this subsection with some properties of the set of symmetry points $S_{opt}(S)$. Note that $S_{opt}(S)$ is not necessarily a singleton. To see how multiple symmetry points can arise, consider $S := \{x \in \mathbb{R}^3 : x_1 \geq 0, x_2 \geq 0, x_1 + x_2 \leq 1, 0 \leq x_3 \leq 1\}$, which is the cross product of a 2-dimensional simplex and a unit interval. Therefore $\textbf{sym}(S) = \min\{\frac{1}{2}, 1\} = \frac{1}{2}$ and $S_{opt}(S) = \{x \in \mathbb{R}^3 : x_1 = x_2 = \frac{1}{3}, x_3 \in [\frac{1}{3}, \frac{2}{3}]\}$.

Proposition 6 *Under Assumption A, $S_{opt}(S)$ is a compact convex set with no interior. If S is a strictly convex set, then $S_{opt}(S)$ is a singleton.*

Proof The convexity of $S_{opt}(S)$ follows directly from the quasiconcavity of $\textbf{sym}(\cdot, S)$, see Theorem 2. Let $\alpha := \textbf{sym}(S)$, and suppose that there exists $\hat{x} \in \textbf{int } S_{opt}(S)$. This implies that there exists $\delta > 0$ such that $\textbf{sym}(x, S) = \alpha$ for all $x \in B(\hat{x}, \delta) \subset S_{opt}(S)$. Then for all d satisfying $\|d\| \leq 1$ we have:

$$\alpha(\hat{x} + \delta d - S) \subseteq S - (\hat{x} + \delta d)$$

which implies that

$$\alpha(\hat{x} - S) + B(0, \delta(1 + \alpha)) \subseteq S - \hat{x} .$$

Using Lemma 2, this implies $\alpha < \textbf{sym}(\hat{x}, S)$, which is a contradiction.

For last statement, suppose $x_1, x_2 \in S_{opt}(S)$ and $x_1 \neq x_2$. Since any strict convex combination of elements of S must lie in the interior of S, for any $\gamma \in (0, 1)$ it follows that

$$(\gamma x_1 + (1 - \gamma)x_2) - \alpha(S - (\gamma x_1 + (1 - \gamma)x_2)) \subseteq \textbf{int } S .$$

Again using Lemma 2, it follows that $\textbf{sym}(\gamma x_1 + (1 - \gamma)x_2, S) > \alpha$, which is also a contradiction.□

Remark 6 *In [35], Klee proved the following notable relation between* $\mathbf{sym}(S)$ *and the dimension of* $S_{opt}(S)$*:*

$$\frac{1}{\mathbf{sym}(S)} + dim(S_{opt}(S)) \leq n \ ,$$

which implies that multiple symmetry points can only exist in dimensions $n \geq 3$*.*

2.4 Computing a symmetry point of S when S is Polyhedral

Our interest in this section lies in computing an ε-approximate symmetry point of S, which is a point $x \in S$ that satisfies:

$$\mathbf{sym}(x, S) \geq (1 - \varepsilon)\mathbf{sym}(S) \ .$$

We focus on the polyhedral case; more specifically, we study the problem in which the convex set of interest is given by the convex hull of finitely many points or by the intersection of finitely many half-spaces.

Although the symmetry function is invariant under equivalent representations of the set S, the question of computing the symmetry of a point in a general convex set is not, as the following example indicates.

Example 1 *Let* $C^n = \{x \in \mathbb{R}^n : \|x\|_\infty \leq 1\}$ *be the* $n-$*dimensional hypercube. Let* v *be a vertex of* C^n*, and define* $H = \{x \in \mathbb{R}^n : \langle x, v\rangle \leq n - 1/2\}$*, and define* $S := C^n \cap H$*. Then* $\mathbf{sym}(0, S) = 1 - 1/2n$ *is obtained by considering the vertex* $-v$*. Assume that* S *is given only by a membership oracle and note that* H *cuts off a pyramid from* S *that is completely contained in exactly one of the* 2^n *orthants of* $\mathbb{R}^n$*. Since we can arbitrarily choose the vertex* v*, in the worst case any deterministic algorithm will need to verify every single orthant to show that* $\mathbf{sym}(0, S) < 1$*, leading to an exponential complexity in the dimension* n*.*

We point out that the convex set S defined in Example 1 could be represented by a separation oracle and the same conclusion would also follow. This suggests that more structure is needed for the representation of S in order to compute an ε-approximate symmetry point of S. In the following two subsections we consider the cases when S is given as the convex hull of finitely many points (Section 2.4.1), and as the intersection of finitely many half-spaces (Section 2.4.2).

2.4.1 S Represented by the Convex Hull of Points

In this subsection we assume that S is given as the convex hull of m given points $w^1, \ldots, w^m \in \mathbb{R}^n$, i.e., $S = \mathbf{conv}\{w^1, \ldots, w^m\}$. Given $x \in S$ and a nonnegative scalar α, it follows from (2.4) that $\mathbf{sym}(x, S) \geq \alpha$ if and only if

$$(1+\alpha)x - \alpha w^i \in S = \mathbf{conv}\left\{w^j : j = 1, \ldots, m\right\} \text{ for every } i = 1, \ldots, m,$$

which can be checked by solving a system of linear inequalities. It follows that $\mathbf{sym}(S)$ is the optimal value of the following optimization problem:

$$\begin{array}{lll} \max\limits_{\alpha,x,\lambda,\nu} & \alpha & \\ \text{s.t.} & (1+\alpha)x - \alpha w^i = \sum\limits_{k=1}^{m} \lambda_k^i w^k, & i = 1, \ldots, m \\ & x = \sum\limits_{k=1}^{m} \nu_k w^k & \\ & e^T\lambda^i = 1, \lambda^i \geq 0\,, & i = 1, \ldots, m \\ & e^T\nu = 1, \nu \geq 0\,, & \end{array} \tag{2.31}$$

which is almost a linear program. Note that the constraints "$x = \sum_{k=1}^m \nu_k w^k$, $e^T\nu = 1$, $\nu \geq 0$" of (2.31) simply state that x must lie in the convex hull of the points $w^1, \ldots, w^m$. However, dividing the first set of constraints by $(1+\alpha)$ one obtains for a given i:

$$x = \left(\frac{\alpha}{1+\alpha} w^i + \frac{1}{1+\alpha} \sum_{k=1}^{m} \lambda_k^i w^k\right)\,,$$

which shows that these constraints themselves imply that x is in the convex hull of $w^1, \ldots, w^m$, and so the former set of constraints can be eliminated. Furthermore, setting $y = (1+\alpha)x$, it follows that $\mathbf{sym}(S)$ is the optimal value of the linear program:

$$\begin{array}{lll} \max\limits_{\alpha,y,\lambda} & \alpha & \\ \text{s.t.} & y - \alpha w^i = \sum\limits_{k=1}^{m} \lambda_k^i w^k, & i = 1, \ldots, m \\ & e^T \lambda^i = 1, \lambda^i \geq 0, & i = 1, \ldots, m \ , \end{array} \tag{2.32}$$

and that any optimal solution $(\alpha^*, y^*, \lambda^*)$ of (2.32) yields $\mathbf{sym}(S) = \alpha^*$ and $x^* = y^*/(1+\alpha^*)$ is a symmetry point of S.

Formulation (2.32) has m^2 nonnegative variables and $mn+m$ equality constraints. Moreover, the analytic center for the slice of the feasible region on the level set corresponding to $\alpha = 0$ is readily available for this formulation by setting

$$\alpha = 0, \quad y = \frac{1}{m} \sum_{k=1}^{m} w^k, \quad \lambda^i = \frac{1}{m} e, \ i = 1, \ldots, m \ ,$$

and therefore (2.32) lends itself to solution by interior-point methods so long as m is not too large.

If m is large it might not be attractive to solve (2.32) directly, and in order to develop a more attractive approach to computing $\mathbf{sym}(S)$ we proceed as follows. Based on (2.31) we can compute $\mathbf{sym}(x, S)$ by simply fixing x. Thus, for each $i = 1, \ldots, m$ define

$$\begin{array}{lll} f_i(x) = \max\limits_{\alpha_i, \lambda^i} & \alpha_i & \\ \text{s.t.} & (1+\alpha_i)x - \alpha_i w^i = \sum\limits_{k=1}^{m} \lambda_k^i w^k \\ & e^T \lambda^i = 1, \lambda^i \geq 0 \ , \end{array} \tag{2.33}$$

and it follows that $\mathbf{sym}(x, S) = \min_{i=1,\ldots,m} f_i(x)$. Dividing the first constraint by $(1+\alpha_i)$ and defining $\theta_i = \frac{\alpha_i}{1+\alpha_i}$ and noting that maximizing θ_i is equivalent to maximizing

α_i, it follows that (2.33) is equivalent to:

$$\begin{aligned} h_i(x) = \max_{\theta_i, \lambda^i} \quad & \theta_i \\ \text{s.t.} \quad & \theta_i w^i + \sum_{k=1}^{m} \lambda_k^i w^k = x \\ & e^T \lambda^i = 1 - \theta_i, \lambda^i \geq 0 \ . \end{aligned} \tag{2.34}$$

Now note that $h_i(x)$ is a concave function, whereby

$$h(x) := \min_{i=1,\ldots,m} h_i(x)$$

is also a concave function, and furthermore

$$\frac{\mathbf{sym}(x, S)}{1 + \mathbf{sym}(x, S)} = \min_{i=1,\ldots,m} h_i(x) = h(x) \ .$$

Moreover, given a value of x, the computation of $h(x)$ and the computation of a subgradient of $h(\cdot)$ at x is easily accomplished by solving the m linear programs (2.34) which each have m nonnegative variables and $n + 1$ equality constraints. Therefore the problem of maximizing $h(x)$ is suitable for classical nondifferentiable optimization methods such as bundle methods, see [29] for example.

2.4.2 S Represented by linear inequalities

In this subsection we assume that S is given as the intersection of m inequalities, i.e., $S := \{x \in \mathbb{R}^n : Ax \leq b\}$ where $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$. We present two methods for computing an ε-approximate symmetry point of S. The first method is based on approximately solving a single linear program with $m^2 + m$ inequalities. For such a method, an interior-point algorithm would require $O(m^6)$ operations per Newton step, which is clearly unattractive. Our second method involves solving $m + 1$ linear programs each of which involves m linear inequalities in n unrestricted variables. This method is more complicated to evaluate, but is clearly more attractive should one want to compute an ε-approximate symmetry point in practice.

Let $\bar{x} \in S$ be given, and let $\alpha \leq \mathbf{sym}(\bar{x}, S)$. Then from the definition of $\mathbf{sym}(\cdot, S)$ in (2.1) we have:

$$A(\bar{x} + v) \leq b \quad \Rightarrow \quad A(\bar{x} - \alpha v) \leq b \ ,$$

which we restate as:

$$Av \leq b - A\bar{x} \quad \Rightarrow \quad -\alpha A_{i\cdot} v \leq b_i - A_{i\cdot}\bar{x} \ , i = 1, \ldots, m \ . \tag{2.35}$$

Now apply a theorem of the alternative to each of the $i = 1, \ldots, m$ implications (2.35). Then (2.35) is true if and only if there exists an $m \times m$ matrix Λ of multipliers that satisfies:

$$\begin{aligned} \Lambda A &= -\alpha A \\ \Lambda(b - A\bar{x}) &\leq b - A\bar{x} \\ \Lambda &\geq 0 \ . \end{aligned} \tag{2.36}$$

Here "$\Lambda \geq 0$" is componentwise for all m^2 components of Λ. This means that $\mathbf{sym}(\bar{x}, S) \geq \alpha$ if and only if (2.36) has a feasible solution. This also implies that $\mathbf{sym}(S)$ is the optimal objective value of the following optimization problem:

$$\begin{aligned} \max_{x,\Lambda,\alpha} \quad & \alpha \\ \text{s.t.} \quad & \Lambda A = -\alpha A \\ & \Lambda(b - Ax) \leq b - Ax \\ & \Lambda \geq 0 \ , \end{aligned} \tag{2.37}$$

and any solution $(x^*, \Lambda^*, \alpha^*)$ of (2.37) satisfies $\mathbf{sym}(x^*, S) = \alpha^*$. Using the first

equalities constraint, we can reformulate the inequalities as:

$$\begin{array}{ll} \max\limits_{x,\Lambda,\alpha} & \alpha \\ \text{s.t.} & \Lambda A = -\alpha A \\ & \Lambda b - \alpha A x \leq b - Ax \\ & \Lambda \geq 0 \ , \end{array} \tag{2.38}$$

Notice that (2.38) is not a linear program. In order to solve this problem, we propose a potential reduction interior point method similar to [63, 56]. The overall computational complexity bound obtained for this method was $O(m^8 \ln(m/\varepsilon))$ which is not satisfactory. The details of the analysis can be found in the Appendix C.1.

To convert (2.38) to a linear program, we make the following change of variables:

$$\gamma = \frac{1}{\alpha} \ , \ \Pi = \frac{1}{\alpha}\Lambda \ , \ y = \frac{1+\alpha}{\alpha}x \ ,$$

which can be used to transform (2.38) to the following linear program:

$$\begin{array}{ll} \min\limits_{y,\Pi,\gamma} & \gamma \\ \text{s.t.} & \Pi A = -A \\ & \Pi b + Ay - b\gamma \leq 0 \\ & \Pi \geq 0 \ . \end{array} \tag{2.39}$$

If (y^*, Π^*, γ^*) is a solution of (2.39), then $\alpha^* := 1/\gamma^* = \mathbf{sym}(S)$ and $x^* := \frac{1}{1+\gamma^*} y^* \in S_{opt}(S)$. Notice that (2.39) has $m^2 + m$ inequalities and mn equations. Suppose we know an approximate analytic center x^a of S. Then it is possible to develop an interior-point method approach to solving (2.39) using information from x^a, and one can prove that a suitable interior-point method will compute an ε-approximate symmetry point of S in $O\left(m \ln\left(\frac{m}{\varepsilon}\right)\right)$ iterations of Newton's method. However, due to the $m^2 + m$ inequalities, each Newton step requires $O(m^6)$ operations, see Appendix C.1.1.

In order to improve on the previous approach, we define the following scalar

quantities $\delta_i^*, i = 1, \ldots, m$:

$$\begin{array}{rll} \delta_i^* := & \max\limits_x & -A_{i.}x \\ & \text{s.t.} & Ax \leq b \ , \end{array} \tag{2.40}$$

and notice that $b_i + \delta_i^*$ is the range of $A_{i.}x$ over $x \in S$ unless the i^{th} constraint is never active. We compute $\delta_i^*, i = 1, \ldots, m$ by solving m linear programs whose feasible region is exactly S. It then follows directly from Proposition 1 that

$$\mathbf{sym}(x, S) = \min_{i=1,\ldots,m} \left\{ \frac{b_i - A_{i.}x}{\delta_i^* + A_{i.}x} \right\} \ . \tag{2.41}$$

We now use (2.41) to create another single linear program to compute $\mathbf{sym}(S)$ as follows. Let $\delta^* := (\delta_1^*, \ldots, \delta_m^*)$ and consider the following linear program that uses δ^* in the data:

$$\begin{array}{ll} \max\limits_{x,\check{\theta}} & \check{\theta} \\ \text{s.t.} & Ax + \check{\theta}(\delta^* + b) \leq b \ . \end{array} \tag{2.42}$$

Proposition 7 *Let $(x^*, \check{\theta}^*)$ be an optimal solution of the linear program (2.42). Then x^* is a symmetry point of S and $\mathbf{sym}(S) = \frac{\check{\theta}^*}{1-\check{\theta}^*}$.*

Proof Suppose that $(x, \check{\theta})$ is a feasible solution of (2.42). Then $\frac{1}{\check{\theta}} \geq \frac{\delta_i^* + b_i}{b_i - A_{i.}x}$, whereby

$$\frac{1-\check{\theta}}{\check{\theta}} = \frac{1}{\check{\theta}} - 1 \geq \frac{\delta_i^* + A_{i.}x}{b_i - A_{i.}x} \ ,$$

and so

$$\frac{b_i - A_{i.}x}{\delta_i^* + A_{i.}x} \geq \frac{\check{\theta}}{1-\check{\theta}} \ , i = 1, \ldots, m \ .$$

It then follows from Proposition 1 that $\mathbf{sym}(x, S) \geq \dfrac{\check{\theta}}{1-\check{\theta}}$, which implies that $\mathbf{sym}(S) \geq \frac{\check{\theta}^*}{1-\check{\theta}^*}$. The proof of the reverse inequality follows similarly. □

This yields the following "exact" method for computing $\mathbf{sym}(S)$ and a symmetry point x^*:

Exact Method:

Step 1 For $i = 1, \ldots, m$ solve the linear program (2.40) for δ_i^*.

Step 2 Let $\delta^* := (\delta_1^*, \ldots, \delta_m^*)$. Solve the linear program (2.42) for an optimal solution $(x^*, \check{\theta}^*)$. Then $x^* \in S_{opt}(S)$ and $\mathbf{sym}(S) = \frac{\check{\theta}^*}{1-\check{\theta}^*}$.

This method involves the exact solution of $m + 1$ linear programs. The first m linear programs can actually be solved in parallel, and their optimal objective values are used in the data for the $(m+1)^{\text{st}}$ linear program. The first m linear programs each have m inequalities in n unrestricted unknowns. The last linear program has m inequalities and $n + 1$ unrestricted unknowns, and could be reduced to n unknowns using variable elimination if so desired. The details can be found in the Appendix C.1.2.

Remark 7 *Although* $\mathbf{sym}(S)$ *can be computed via linear programming when* S *is represented either as a convex hull of points or as the intersection of half-spaces, the latter case appears to be genuinely easier; indeed, the Exact Method solves a sequence of* $m + 1$ *linear programs of size* $m \times n$ *when* S *is given by half-spaces, instead of a single linear program with* m^2 *inequalities when* S *is represented as the convex hull of points. It is an open question whether there is a more efficient scheme than solving (2.32) for computing* $\mathbf{sym}(S)$ *when* S *is represented as the convex hull of points.*

Table 2.4.2 summarizes the complexity results of the three approaches mentioned before. Their analysis can be found in the appendix.

	Computational Complexity ($O^*(\cdot)$)			
Method	outer	inner	overall	Comments
potential reduction	m^2	m^6	m^8	quasi-convex
primal-dual	m	m^6	m^7	linear program
dual	$m^{1.5}$	m^3	$m^{4.5}$	decomposable

Chapter 3

Projective pre-conditioners for improving the behavior of a homogeneous conic linear system

In this chapter, our interest lies in behavioral and computational characteristics of the following homogeneous convex feasibility problem in conic linear form:

$$F : \begin{cases} Ax &= 0 \\ x &\in C\backslash\{0\}\,, \end{cases} \tag{3.1}$$

where A is a linear operator and C is a convex cone.

It is well known that the standard form conic feasibility problem

$$\begin{cases} \bar{A}x &= \bar{b} \\ x &\in K \end{cases}$$

is a special case of F under the assignments $C \leftarrow K \times \mathbb{R}_+$, $A \leftarrow [\bar{A}, -\bar{b}]$ and the qualification that we seek solutions in the interiors of the cones involved. Furthermore, this setting is general enough to encompass convex optimization as well.

In the context of interior-point methods (IPMs), the system F has good computational complexity if an IPM for solving F has a good iteration bound. We also say

that F has good geometric behavior if the width of the cone of feasible solutions of F is large, equivalently if F has a solution x whose relative distance from ∂C is large. Choose a point $\bar{s} \in \text{int}C^*$, and note that F is equivalent to the normalized problem $F_{\bar{s}}: \ Ax = 0, \ \bar{s}^T x = 1, \ x \in C$. We show that both the computational complexity and the geometry behavior of F can be bounded as a function of only two quantities: (i) the symmetry of the so-called *image set* $H_{\bar{s}} := \{Ax : \bar{s}^T x = 1, \ x \in C\}$ of $F_{\bar{s}}$ about the origin, denoted by $\mathbf{sym}(0, H_{\bar{s}})$, and the complexity value ϑ of the barrier function for C. These results are shown in Section 3.2 after some initial definitions and analysis are developed in Section 3.1.

In Section 3.3 we present a general theory for transforming the normalized homogeneous conic system $F_{\bar{s}}$ to an equivalent system via projective transformation. Such a projective transformation serves to pre-condition the conic system into a system that has both geometric and computational properties with certain guarantees; we use the term "projective pre-conditioner" to describe such a projective transformation. Under the assumption that F has an interior solution, there must exist projective pre-conditioners that transform $F_{\bar{s}}$ into equivalent systems that are solvable in strongly-polynomial time in m and ϑ. The quality of a projective pre-conditioner depends on the ability to compute a point $\hat{v}$ that is "deep" in the set $H_{\bar{s}}^\circ = \{v : \bar{s} - A^T v \in C^*\}$. Several constructive approaches for computing such points are discussed, including a stochastic method based on a geometric random walk.

The geometric random walk approach is further developed in Section 3.4, with associated complexity analysis. In Section 3.5 we present results from computational experiments designed to assess the practical viability of the projective pre-conditioner method based on geometric random walks. We generated 300 linear programming feasibility problems (100 each in three sets of dimensions) designed to be poorly behaved. We present computational evidence that the method is very effective; for the 100 problems of dimension 1000×5000 the average IPM iterations decreased by 46% and average total running time decreased by 33%, for example.

Section 3.6 contains summary conclusions and next steps.

We point out that a very different pre-conditioner for F was proposed in [13] that

is a linear (not projective) transformation of the range space of A and that aims to improve Renegar's condition measure $\mathcal{C}(A)$ (but does not improve the complexity of the original problem or the geometry of the feasible region).

3.0.3 Notation

Let $e = (1, \ldots, 1)^T \in \mathbb{R}^d$ denote the vector of ones in dimension d. Given a closed convex set $S \subset \mathbb{R}^d$ with $0 \in S$, the *polar* of S is $S^\circ := \{y \in \mathbb{R}^d : y^T x \leq 1 \text{ for all } x \in S\}$ and satisfies $S^{\circ\circ} = S$, see Rockafellar [53]. Given a closed convex cone $K \subset \mathbb{R}^d$, the (positive) *dual cone* of K is $K^* := \{y \in \mathbb{R}^d : y^T x \geq 0 \text{ for all } x \in K\}$ and satisfies $K^{**} = K$, also see [53]. For a general norm $\|\cdot\|$, let $B(c, r)$ and $\text{dist}(x, T)$ denote the ball of radius r centered at c and the distance from a point x to a set T, respectively.

3.1 Normalization and $\bar{s}$-norm, behavioral measures, and barrier calculus

Regarding the conic feasibility problem (3.1), we make the following assumptions:

Assumption 1 *C is a* regular *cone, i.e.,* $\text{int}C \neq \emptyset$ *and C contains no line.*

Assumption 2 *F has a solution, i.e.,*

$$\mathcal{F} := \{x \in \mathbb{R}^n : Ax = 0,\ x \in C \backslash \{0\}\} \neq \emptyset .$$

Assumption 3 $\text{rank}A = m$.

3.1.1 Normalization of F and a class of norms that are linear on C

Let $\bar{s} \in \text{int}C^*$ be chosen, then $x \in C\backslash\{0\}$ if and only if $x \in C$ and $\bar{s}^T x > 0$, whereby we can write F equivalently as the normalized problem:

$$F_{\bar{s}} : \left\{ \begin{array}{rcl} Ax & = & 0 \\ \bar{s}^T x & = & 1 \\ x & \in & C, \end{array} \right.$$

whose feasible region is $\mathcal{F}_{\bar{s}} := \{x \in \mathbb{R}^n : Ax = 0,\ x \in C,\ \bar{s}^T x = 1\}$.

Given the regular cone C and $\bar{s} \in \text{int}C^*$, the linear functional:

$$f(x) := \bar{s}^T x$$

behaves like a norm when restricted to $x \in C$, namely $f(x)$ is (trivially) convex and positively homogeneous on C, and $f(x) > 0$ for $x \in C\backslash\{0\}$. The natural norm that agrees with $f(x) := \bar{s}^T x$ on C is:

$$\begin{array}{rcll} \|x\|_{\bar{s}} & := & \min\limits_{x^1, x^2} & \bar{s}^T(x^1 + x^2) \\ & & \text{s.t.} & x^1 - x^2 = x \\ & & & x^1 \in C \\ & & & x^2 \in C\ , \end{array}$$

and note that $\|x\|_{\bar{s}} = \bar{s}^T x$ for $x \in C$. We refer to this norm as the "$\bar{s}$-norm." Indeed, $\|x\|_{\bar{s}}$ is an exact generalization of the L_1 norm in the case when $C = \mathbb{R}^n_+$ and $\bar{s} = e$:

$$\begin{array}{rcll} \|x\|_1 & := & \min\limits_{x^1, x^2} & e^T(x^1 + x^2) \\ & & \text{s.t.} & x^1 - x^2 = x \\ & & & x^1 \in \mathbb{R}^n_+ \\ & & & x^2 \in \mathbb{R}^n_+\ . \end{array}$$

We will make extensive use of the family of norms $\|\cdot\|_{\bar{s}}$ herein. In the case when C is a self-scaled cone, both $\|\cdot\|_{\bar{s}}$ and its dual norm have convenient explicit formulas, for details see Section 2 of [18].

3.1.2 Measuring the behavior of F: geometry and complexity

A natural way to think of "good" geometric behavior of F is in terms of the existence of a solution x of F that is nicely interior to the cone C. However, due to the homogeneity of F any solution $x \in \text{int}C$ can be re-scaled by a positive constant to yield a solution that is arbitrarily far from the boundary of C. Given a norm $\|\cdot\|$ on $\mathbb{R}^n$, we therefore consider the following measure of distance of x from the boundary of C that is invariant under positive scalings:

$$\text{reldist}(x, \partial C) := \frac{\text{dist}(x, \partial C)}{\|x\|} , \tag{3.2}$$

where $\text{dist}(x, S) := \inf_{y \in S} \|x - y\|$. We define the "width" or "min-width" of the feasible region $\mathcal{F}$ under the norm $\|\cdot\|$ to be the quantity τ_F defined by:

$$\tau_F = \max_{x \in \mathcal{F}} \{\text{reldist}(x, \partial C)\} = \max_{\substack{Ax = 0 \\ x \in C \setminus \{0\}}} \{\text{reldist}(x, \partial C)\} . \tag{3.3}$$

Note that τ_F is larger to the extent that F has a solution of small norm whose distance from the boundary of C is large. τ_F is a variation on the notion of the "inner measure" of Goffin [22] when the norm is Euclidean, and has also been used in similar format in [19, 17].

As is customary, we will measure the computational behavior of F using a worst-case computational complexity upper bound on the number of iterations that a suitably designed interior-point method (IPM) needs to compute a solution of F.

We will show that both the geometry measure τ_F and the computational complexity can be bounded as simple functions of the *symmetry* of the origin in the *image set* of $F_{\bar{s}}$, which we now define.

The *image set* $H = H_{\bar{s}}$ of $F_{\bar{s}}$ is defined as:

$$H = H_{\bar{s}} := \{Ax : x \in C \ , \ \bar{s}^T x = 1\} \ .$$

Note that the assumption that F has a solution implies that $0 \in H$.

We consider the *symmetry* of a point in a convex set as presented in Chapter 1. For the reader convenience, we recall its definition. Let $S \subset \mathbb{R}^d$ be a convex set. Define:

$$\mathbf{sym}(\bar{x}, S) := \max \{t \mid y \in S \Rightarrow \bar{x} - t(y - \bar{x}) \in S\} \ ,$$

which essentially measures how symmetric S is about the point x. Define

$$\mathbf{sym}(S) := \max_{\bar{x} \in S} \mathbf{sym}(\bar{x}, S) \ ,$$

and x^* is called a *symmetry point* of S if $\mathbf{sym}(x^*, S) = \mathbf{sym}(S)$.

3.1.3 Logarithmically-homogeneous barrier calculus

We presume that we have a ϑ-logarithmically homogeneous (self-concordant) barrier function $f(\cdot)$ for C, see [46].

Remark 8 *We will use the following properties of a ϑ-logarithmically homogeneous barrier:*
(i) $\bar{u} \in \text{int}C$ if and only if $-\nabla f(\bar{u}) \in \text{int}C^*$
(ii) $f^*(s) := -\inf_{x \in \text{int}C}\{s^T x + f(x)\}$ is a ϑ-logarithmically homogeneous barrier for C^*
(iii) $\bar{s} \in \text{int}C^*$ if and only if $-\nabla f^*(\bar{s}) \in \text{int}C$
(iv) $-\nabla f(\bar{u})^T \bar{u} = \vartheta$ and $-\nabla f^*(\bar{s})^T \bar{s} = \vartheta$ for $\bar{u} \in \text{int}C$ and $\bar{s} \in \text{int}C^*$
(v) $\nabla f(\bar{u}) = -H(\bar{u})\bar{u}$ for $\bar{u} \in \text{int}C$, where $H(\cdot)$ is the Hessian of the barrier $f(\cdot)$
(vi) $\bar{u} = -\nabla f^*(\bar{s})$ if and only if $\bar{s} = -\nabla f(\bar{u})$
(vii) $-\nabla f(\bar{u})^T y \geq \sqrt{y^T H(\bar{u}) y}$ for $\bar{u} \in \text{int}C$ and $y \in C$.

Properties (i)-(vi) above are restatements of results in [46] or [52], whereas (vii) is borrowed from the proof of Lemma 5 of [48].

3.2 Behavioral bounds on F

Let $\bar{s} \in \text{int}C^*$ be chosen, and let $F_{\bar{s}}$ be as defined in Section 3.1.1 with image set $H = H_{\bar{s}}$ as in Section 3.1.2. The following result shows that the width τ_F of the feasible region $\mathcal{F}$ is linearly related to the symmetry of 0 in the image set H.

Theorem 14 *Let $\bar{s} \in \text{int}C^*$ be chosen. Under the norm $\|\cdot\|_{\bar{s}}$, the width τ_F of $\mathcal{F}$ satisfies:*

$$\left(\frac{1}{\vartheta}\right)\frac{\mathbf{sym}(0,H_{\bar{s}})}{1+\mathbf{sym}(0,H_{\bar{s}})} \leq \tau_F \leq \frac{\mathbf{sym}(0,H_{\bar{s}})}{1+\mathbf{sym}(0,H_{\bar{s}})}.$$

In particular, $\frac{1}{2\vartheta}\mathbf{sym}(0,H_{\bar{s}}) \leq \tau_F \leq \mathbf{sym}(0,H_{\bar{s}})$.

Remark 9 *The left-hand bound in the theorem depends on the complexity parameter ϑ of the barrier function $f(\cdot)$, which seems a bit unnatural since the width τ_F is a geometric object that should not directly depend on the barrier function. If we use the universal barrier of Nesterov and Nemirovskii [46], we can replace ϑ by* CONST $\times\, n$ *for the large absolute constant* CONST *of the universal barrier. Alternatively, we can replace ϑ by the complexity value ϑ^* of an optimal barrier for C.*

Our next result shows that the computational complexity of a standard interior-point method (IPM) for computing a solution of F also depends only on $\mathbf{sym}(0,H)$ and ϑ. In order to establish this result we first develop the model that will be solved by the IPM.

Let $\bar{s} \in C^*$ be chosen, and assign $\bar{x} \leftarrow -\frac{1}{\vartheta}\nabla f^*(\bar{s})$. It follows from Remark 8 that $\bar{x} \in \text{int}C$ and $\bar{s}^T\bar{x} = 1$. Construct the simple optimization problem:

$$\begin{array}{lll} \text{OP}: \quad \theta^* := \max\limits_{x,\theta} & \theta & \\ & Ax + (A\bar{x})\theta = 0 & \\ & \bar{s}^T x = 1 & \\ & x \in C, & \end{array} \tag{3.4}$$

and note that $(x, \theta) = (\bar{x}, -1)$ is a feasible solution of OP. Furthermore, $(\bar{x}, -1)$ is the *analytic center* associated with OP for the barrier function $f(\cdot)$, i.e., $(\bar{x}, -1)$ is the optimal solution of the problem of minimizing the barrier function $f(x)$ over the feasible region of OP. We will therefore use $(\bar{x}, -1)$ as a starting point with which to initiate a standard primal feasible interior-point method for approximately following the central path $(x(\eta), \theta(\eta))$ of the parameterized barrier problem:

$$\begin{array}{lllll} \mathrm{OP}_\eta: & \max\limits_{x,\theta} & -f(x) & +\eta \cdot \theta & \\ & & Ax & +(A\bar{x})\theta & = 0 \\ & & \bar{s}^T x & & = 1 \\ & & x & & \in \mathrm{int}C \end{array} \tag{3.5}$$

for an increasing sequence of values of $\eta > 0$, until we have computed a point (x, θ) that satisfies $\theta \geq 0$, whereby $(x + \theta\bar{x})$ is a feasible solution of F. The details of the algorithm scheme are presented in Algorithm A in the Appendix, where we also prove the following complexity bound for the method:

Theorem 15 *Let $\bar{s} \in \mathrm{int}C^*$ be chosen. The standard primal-feasible interior-point Algorithm A applied to (3.5) will compute $\tilde{x}$ satisfying $A\tilde{x} = 0, \tilde{x} \in \mathrm{int}C$ in at most*

$$\left\lceil 9\sqrt{\vartheta} \ln\left(11\vartheta\left(1 + \frac{1}{\mathbf{sym}(0, H_{\bar{s}})}\right)\right)\right\rceil$$

iterations of Newton's method. Furthermore, under the norm $\|\cdot\|_{\bar{s}}$, $\tilde{x}$ will also satisfy

$$\mathrm{reldist}_{\bar{s}}(\tilde{x}, \partial C) \geq \frac{1}{1.2\vartheta + 0.2} \cdot \tau_F \ .$$

(Note that the complexity bound is trivially valid even when $\mathbf{sym}(0, H_{\bar{s}}) = 0$, using the standard convention that $1/0 = \infty$.) Taken together, Theorems 14 and 15 present bounds on both behavioral measures that are simple functions of the complexity value ϑ of the barrier function and the symmetry of 0 in the image set $H = H_{\bar{s}}$. Furthermore, Algorithm A will compute a feasible point whose relative distance from ∂C is within a factor of $O(\vartheta)$ of the maximum relative distance from ∂C over all

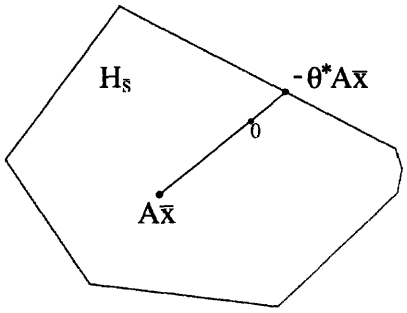

Figure 3-1: The image set $H_{\bar{s}}$ and the points $A\bar{x}, 0$, and $-\theta^* A\bar{x}$.

feasible points.

Figure 3-1 can be used to gain some intuition on the complexity result of Theorem 15. The figure portrays the image set $H_{\bar{s}}$, which by assumption contains 0. Furthermore, $A\bar{x} \in H_{\bar{s}}$ by design of $\bar{x}$. The optimal value θ^* of (3.4) is the largest value of θ for which $-\theta A\bar{x} \in H_{\bar{s}}$. Also notice in Figure 3-1 that $\theta^* \geq \mathbf{sym}(0, H_{\bar{s}})$, and so θ^* will be large if $\mathbf{sym}(0, H_{\bar{s}})$ is large. Since the interior-point algorithm starts at the analytic center $(\bar{x}, -1)$ where $\theta = -1$ and will stop when the current iterate (x, θ) satisfies $\theta \geq 0$, it follows from the linear convergence theory of interior-point methods that the iteration bound will be proportional to the logarithm of the ratio of the initial optimality gap divided by the optimality gap at the stopping point. This ratio is simply $(1 + \theta^*)/\theta^*$, which is bounded above by $1 + 1/(\mathbf{sym}(0, H_{\bar{s}}))$.

Note that one can view $\mathbf{sym}(0, H_{\bar{s}})$ as a condition number of sorts associated with $F_{\bar{s}}$, see [13]. In the next section, we will show how projective transformations can be used to modify $\mathbf{sym}(0, H)$ and hence improve the behavior (both geometry and computational complexity) of $F_{\bar{s}}$.

3.2.1 Proof of Theorem 14

The proof of Theorem 14 is derived from the following two lemmas which we will prove in turn. For $\bar{u} \in \text{int}C$ define the ellipsoidal norm induced by $H(\bar{u})$ by $\|v\|_{\bar{u}} := \sqrt{v^T H(\bar{u}) v}$. Let $B_{\bar{s}}(c, r)$ and $B_{\bar{u}}(c, r)$ denote the balls centered at c of radius r in the norms $\|\cdot\|_{\bar{s}}$ and $\|\cdot\|_{\bar{u}}$, respectively. Note that $B_{\bar{u}}(c, r)$ is an ellipsoid whereas

$B_{\bar{s}}(c,r)$ is not ellipsoidal, i.e., these two norms are not part of the same family. The following lemma shows that $\|\cdot\|_{\bar{s}}$ and $\|\cdot\|_{\bar{u}}$ are within a factor of ϑ of one another if $\bar{u} = -\nabla f^*(\bar{s})$.

Lemma 4 *Let $\bar{s} \in \text{int}C^*$ be chosen, and define $\bar{u} := -\nabla f^*(\bar{s})$. Then*

$$B_{\bar{u}}(0, 1/\vartheta) \subset B_{\bar{s}}(0,1) \subset B_{\bar{u}}(0,1) \quad \text{and} \quad \text{reldist}_{\bar{s}}(\bar{u}, \partial C) \geq \frac{1}{\vartheta} .$$

Lemma 5 *Let $\bar{s} \in \text{int}C^*$ be chosen and define $\bar{x} := -\nabla f^*(\bar{s})/\vartheta$. Then under the norm $\|\cdot\|_{\bar{s}}$,*

$$\text{reldist}_{\bar{s}}(\bar{x}, \partial C) \left(\frac{\mathbf{sym}(0, H_{\bar{s}})}{1 + \mathbf{sym}(0, H_{\bar{s}})} \right) \leq \text{reldist}_{\bar{s}}(\bar{x}, \partial C) \left(\frac{\theta^*}{1+\theta^*} \right) \leq \tau_F \quad \textit{and}$$

$$\tau_F \leq \frac{\mathbf{sym}(0, H_{\bar{s}})}{1 + \mathbf{sym}(0, H_{\bar{s}})} \leq \frac{\theta^*}{1+\theta^*} .$$

Proof of Theorem 14: Define $\bar{x} := -\nabla f^*(\bar{s})/\vartheta$. Then $\bar{x}$ is a positive scaling of $\bar{u}$ defined in Lemma 4, and so $\text{reldist}_{\bar{s}}(\bar{x}, \partial C) = \text{reldist}_{\bar{s}}(\bar{u}, \partial C) \geq \frac{1}{\vartheta}$. Substituting this inequality into the first inequality of Lemma 5 yields the first inequality of the theorem, and the second inequality of the theorem is simply the third inequality of Lemma 5.□

Proof of Lemma 4: Suppose that x satisfies $\|x\|_{\bar{u}} \leq 1$. Then $x^1 := \frac{1}{2}(\bar{u} + x)$ and $x^2 := \frac{1}{2}(\bar{u} - x)$ satisfy $x^1, x^2 \in C$ from the theory of self-concordance, and $x^1 - x^2 = x$, whereby from the definition of the $\bar{s}$-norm we have $\|x\|_{\bar{s}} \leq \bar{s}^T(x^1 + x^2) = \bar{s}^T\bar{u} = \vartheta$. Therefore $B_{\bar{u}}(0,1) \subset B_{\bar{s}}(0,\vartheta)$, which is equivalent to the first set inclusion of the lemma.

Let $L := \{x \in C : \bar{s}^T x = 1\}$. For any $x \in L$ we have $1 = \bar{s}^T x = -\nabla f(\bar{u})^T x \geq \sqrt{x^T H(\bar{u}) x} = \|x\|_{\bar{u}}$, where the second equality and the inequality follow from (vi) and (vii) of Remark 8, respectively. Similarly, if $x \in -L$ then $\|x\|_{\bar{u}} \leq 1$ as well. Noticing that $B_{\bar{s}}(0,1)$ is the convex hull of L and $-L$, it follows that $x \in B_{\bar{s}}(0,1)$ implies $\|x\|_{\bar{u}} \leq 1$, or equivalently, $B_{\bar{s}}(0,1) \subset B_{\bar{u}}(0,1)$.

Last of all, it follows from the theory of self-concordance that $\bar{u} + B_{\bar{u}}(0,1) \subset C$, whereby $\bar{u} + B_{\bar{s}}(0,1) \subset \bar{u} + B_{\bar{u}}(0,1) \subset C$. Therefore

$$\text{reldist}_{\bar{s}}(\bar{u}, \partial C) \geq \frac{1}{\|\bar{u}\|_{\bar{s}}} = \frac{1}{\bar{s}^T \bar{u}} = \frac{1}{\vartheta} ,$$

where the last equality follows from (iv) of Remark 8. □

We are grateful to Nesterov [45] for contributing to a strengthening of a previous version of Lemma 4 and its proof.

Proof of Lemma 5: Recall from Remark 8 that $\bar{x} \leftarrow -\frac{1}{\vartheta}\nabla f^*(\bar{s})$ satisfies $\bar{x} \in \text{int}C$ and $\bar{s}^T\bar{x} = 1$. Therefore $A\bar{x} \in H_{\bar{s}}$ and hence $-\mathbf{sym}(0, H_{\bar{s}})A\bar{x} \in H_{\bar{s}}$, whereby there exists $x \in C$ satisfying $\bar{s}^T x = 1$ and $Ax = -\mathbf{sym}(0, H_{\bar{s}})A\bar{x}$, and therefore θ^* of (3.4) must satisfy:

$$\theta^* \geq \mathbf{sym}(0, H_{\bar{s}}) . \tag{3.6}$$

Therefore the second and third inequalities of the lemma imply the first and fourth inequalities of the lemma. To prove the second inequality, let (x^*, θ^*) be an optimal solution of (3.4), and let $x = \frac{x^* + \theta^* \bar{x}}{1+\theta^*}$. Then $x \in C$, $\|x\|_{\bar{s}} = \bar{s}^T x = 1$, and $Ax = 0$, and by construction

$$B_{\bar{s}}\left(x, \frac{\theta^*}{1+\theta^*}\text{dist}(\bar{x}, \partial C)\right) \subset C ,$$

whereby we have

$$\tau_F \geq \text{reldist}(x, \partial C) \geq \frac{\theta^*}{1+\theta^*}\text{dist}(\bar{x}, \partial C) = \frac{\theta^*}{1+\theta^*}\text{reldist}(\bar{x}, \partial C) ,$$

which demonstrates the second inequality. To prove the third inequality of the lemma, let $\hat{x} \in C$ satisfy the maximization for which τ_F is defined in (3.3), whereby without loss of generality $\|\hat{x}\|_{\bar{s}} = \bar{s}^T\hat{x} = 1$, $A\hat{x} = 0$, and $B_{\bar{s}}(\hat{x}, \tau_F) \subset C$. Let $y \in H_{\bar{s}}$ be given, whereby $y = Ax$ for some $x \in C$ satisfying $\|x\|_{\bar{s}} = \bar{s}^T x = 1$, and define:

$$\tilde{x} := \frac{\hat{x} - \tau_F x}{1 - \tau_F} ,$$

and it follows that $\tilde{x} \in C$, $\bar{s}^T \tilde{x} = 1$, and hence

$$-\frac{\tau_F}{1-\tau_F} y = A\tilde{x} \in H_{\bar{s}} \ .$$

Therefore $\mathbf{sym}(0, H_{\bar{s}}) \geq \frac{\tau_F}{1-\tau_F}$, and rearranging this last inequality yields the third inequality of the lemma. □

3.3 Pre-conditioning F by projective transformation of $F_{\bar{s}}$

Herein we present a systematic approach to transforming the problem $F_{\bar{s}}$ to an equivalent problem $F_{\hat{s}}$ for a suitably chosen vector $\hat{s} \in \text{int}C^*$, with the goal of improving the symmetry of 0 in the associated image set $H_{\hat{s}}$. Recall the following relevant facts about the symmetry function $\mathbf{sym}(\cdot)$ presented in Chapter 1:

Remark 10 *Let $S \subset \mathbb{R}^m$ be a nonempty closed bounded convex set. The following properties of* $\mathbf{sym}(\cdot)$ *hold:*
(i) Let $\mathcal{A}(x) := Mx + g$, M nonsingular. If $\bar{x} \in S$, then $\mathbf{sym}(\mathcal{A}(\bar{x}), \mathcal{A}(S)) = \mathbf{sym}(\bar{x}, S)$.
(ii) If $0 \in S$, then $\mathbf{sym}(0, S) = \mathbf{sym}(0, S^\circ)$.
(iii) $\mathbf{sym}(S) \geq \frac{1}{m}$.

Under Assumption 2, $0 \in H_{\bar{s}}$, whereby $H_{\bar{s}}$ is a closed convex set containing the origin. Therefore $H_{\bar{s}}^\circ$, the polar of $H_{\bar{s}}$, is also a closed convex set containing the origin, and $H_{\bar{s}}^{\circ\circ} = H_{\bar{s}}$. In fact, there is a simple form for $H_{\bar{s}}^\circ$ given by the following:

Proposition 8 *Let $\bar{s} \in \text{int}C^*$ be chosen. Then $H_{\bar{s}}^\circ = \{v \in \mathbb{R}^m : \bar{s} - A^T v \in C^*\}$.*

Proof We have:

$$\begin{aligned} H_{\bar{s}}^{\circ} &= \{v : v^T y \leq 1 \text{ for all } y \in H_{\bar{s}}\} \\ &= \{v : v^T Ax \leq 1 \text{ for all } x \text{ that satisfy } \bar{s}^T x = 1, x \in C\} \\ &= \{v : v^T Ax \leq \bar{s}^T x \text{ for all } x \text{ that satisfy } \bar{s}^T x = 1, x \in C\} \\ &= \{v : (\bar{s} - A^T v)^T x \geq 0 \text{ for all } x \in C\} \\ &= \{v : \bar{s} - A^T v \in C^*\} . \end{aligned}$$

□

It is curious to note from Proposition 8 that while checking membership in $H_{\bar{s}}$ is presumably not easy (validating that $0 \in H_{\bar{s}}$ is an equivalent task to that of solving F), the set $H_{\bar{s}}^{\circ}$ is in fact easy to work with in at least two ways. First, $0 \in \text{int}H_{\bar{s}}^{\circ}$, so we have a known point in the interior of $\text{int}H_{\bar{s}}^{\circ}$. Second, checking membership in $H_{\bar{s}}^{\circ}$ is a relatively simple task if we have available a membership oracle for C^*.

Motivated by Theorems 14 and 15 which bound the geometric and computational behavior of F in terms of the symmetry of the origin in $H_{\bar{s}}$, we now consider replacing $\bar{s}$ by some other vector $\hat{s} \in \text{int}C^*$ with the goal of improving $\mathbf{sym}(0, H_{\hat{s}})$. We proceed as follows. Taking $\bar{s} \in \text{int}C^*$ as given, suppose we choose some $\hat{v} \in \text{int}H_{\bar{s}}^{\circ} = \{v \in \mathbb{R}^m : \bar{s} - A^T v \in C^*\}$, and define $\hat{s} := \bar{s} - A^T \hat{v}$, therefore $\hat{s} \in \text{int}C^*$. We replace $\bar{s}$ by $\hat{s}$, obtaining the modified normalized feasibility problem:

$$F_{\hat{s}} : \left\{ \begin{array}{rcl} Ax & = & 0 \\ \hat{s}^T x & = & 1 \\ x & \in & C , \end{array} \right.$$

with modified image set:

$$H_{\hat{s}} = \{Ax : x \in C \ , \ \hat{s}^T x = 1\}$$

and polar set:

$$H_{\hat{s}}^{\circ} = \{v \in \mathbb{R}^m : \hat{s} - A^T v \in C^*\} .$$

The following shows that $\mathbf{sym}(0, H_{\hat{s}})$ inherits the symmetry of $\hat{v}$ in the original polar image set $H_{\bar{s}}^{\circ}$.

Theorem 16 *Let $\bar{s} \in \text{int}C^*$ be given. Let $\hat{v} \in \text{int}H_{\bar{s}}^{\circ}$ be chosen and define $\hat{s} := \bar{s} - A^T\hat{v}$. Then*

$$\mathbf{sym}(0, H_{\hat{s}}) = \mathbf{sym}(\hat{v}, H_{\bar{s}}^{\circ}) \ .$$

Proof We have:

$$\begin{aligned} H_{\bar{s}}^{\circ} - \{\hat{v}\} &= \{u = v - \hat{v} : \bar{s} - A^T v \in C^*\} = \{u : \bar{s} - A^T\hat{v} - A^T u \in C^*\} \\ &= \{u : \hat{s} - A^T u \in C^*\} = H_{\hat{s}}^{\circ} \ . \end{aligned}$$

It then follows from (ii) of Remark 10 that

$$\mathbf{sym}(0, H_{\hat{s}}) = \mathbf{sym}(0, H_{\hat{s}}^{\circ}) = \mathbf{sym}(0, H_{\bar{s}}^{\circ} - \{\hat{v}\}) = \mathbf{sym}(\hat{v}, H_{\bar{s}}^{\circ}) \ ,$$

where the last equality above readily follows from (i) of Remark 10. □

Note that the following projective transformations map $F_{\bar{s}}$ and $F_{\hat{s}}$ onto one another:

$$x' \leftarrow \frac{x}{\hat{s}^T x} \text{ and } x \leftarrow \frac{x'}{\bar{s}^T x'} \ . \tag{3.7}$$

Furthermore, Theorem 16 has an interesting interpretation in the context of projective transformations and polarity theory which we will discuss in Subsection 3.3.2.

Our present goal, however, is to use Theorem 16 constructively to develop a method for transforming $F_{\bar{s}}$. Suppose we can compute a point $\hat{v} \in H_{\bar{s}}^{\circ}$ with good symmetry in $H_{\bar{s}}^{\circ}$; letting $\alpha := \mathbf{sym}(\hat{v}, H_{\bar{s}}^{\circ})$ we seek $\hat{v}$ for which $\alpha > \mathbf{sym}(0, H_{\bar{s}})$ and is relatively large, for example, $\alpha = \Omega(1/m)$. Then replace $\bar{s}$ by $\hat{s} := \bar{s} - A^T\hat{v}$ and work instead with with $F_{\hat{s}}$. Theorem 16 states that the transformed system will have $\mathbf{sym}(0, H_{\hat{s}}) = \alpha$, i.e., the transformed system will take on the symmetry of $\hat{v}$ in $H_{\bar{s}}^{\circ}$. This is most important, since it then follows from Theorems 14 and 15 that the transformed system $F_{\hat{s}}$ will have geometry and complexity behavior that will depend on α as well. We formalize this method and the above conclusion as follows:

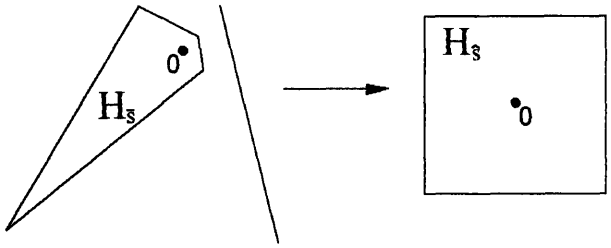

Figure 3-2: Projective transformation of the image set $H_{\bar{s}}$ to improve the symmetry of 0 in the transformed image set.

Projective Pre-conditioning Method (PPM)

Step 1. Construct $H_{\bar{s}}^{\circ} := \{v \in \mathbb{R}^m : \bar{s} - A^T v \in C^*\}$

Step 2. Find a suitable point $\hat{v} \in H_{\bar{s}}^{\circ}$ (with hopefully good symmetry in $H_{\bar{s}}^{\circ}$)

Step 3. Compute $\hat{s} := \bar{s} - A^T \hat{v}$

Step 4. Construct the transformed problem:

$$F_{\hat{s}} : \begin{cases} Ax &= 0 \\ \hat{s}^T x &= 1 \\ x &\in C \end{cases} \tag{3.8}$$

Step 5. The transformed image set is $H_{\hat{s}} := \{Ax \in \mathbb{R}^m : x \in C \ , \ \hat{s}^T x = 1\}$, and $\mathbf{sym}(0, H_{\hat{s}}) = \mathbf{sym}(\hat{v}, H_{\bar{s}}^{\circ})$.

Figure 3-2 illustrates the strategy of the Projective Pre-Conditioning Method. On the left part of the figure is the image set $H_{\bar{s}}$, and notice that $H_{\bar{s}}$ is not very symmetric about the origin, i.e., $\mathbf{sym}(0, H_{\bar{s}}) << 1$. However, under the projective transformation given by the projective plane in the slanted vertical line, $H_{\bar{s}}$ is transformed to a box that is perfectly symmetric about the origin, i.e., $\mathbf{sym}(0, H_{\hat{s}}) = 1$. (In general, of course, we can at best attain $\mathbf{sym}(0, H_{\hat{s}}) = 1/m$.)

The following corollary follows from Theorem 16 and the above discussion, using Theorems 14 and 15:

Corollary 3 *Let $\bar{s} \in \text{int}C^*$ be chosen, and suppose that the Projective Pre-conditioning*

Method has been run, and let $\alpha := \mathbf{sym}(\hat{v}, H_{\bar{s}}^{\circ})$. *Under the norm* $\|\cdot\|_{\hat{s}}$, *the width* τ_F *of* $\mathcal{F}$ *satisfies:*

$$\left(\frac{1}{\vartheta}\right)\frac{\alpha}{1+\alpha} \;\leq\; \tau_F \;\leq\; \frac{\alpha}{1+\alpha}\,,$$

and the standard primal-feasible interior-point Algorithm A applied to (3.8) will compute $\tilde{x}$ *satisfying* $A\tilde{x} = 0, \tilde{x} \in \mathrm{int}C$ *in at most*

$$\left\lceil 9\sqrt{\vartheta}\ln\left(11\vartheta\left(1+\frac{1}{\alpha}\right)\right)\right\rceil$$

iterations of Newton's method. Furthermore, under the norm $\|\cdot\|_{\hat{s}}$, $\tilde{x}$ *will also satisfy*

$$\mathrm{reldist}_{\hat{s}}(\tilde{x}, \partial C) \geq \frac{1}{1.2\vartheta + 0.2}\cdot\tau_F\,.$$

Let us now presume that F has an interior solution, whereby $0 \in \mathrm{int}H_{\bar{s}}$ and $H_{\bar{s}}^{\circ}$ will be bounded and $\mathbf{sym}(0, H_{\bar{s}}) = \mathbf{sym}(0, H_{\bar{s}}^{\circ}) > 0$. Furthermore, we know from (iii) of Remark 10 that there exists a point v whose symmetry value in $H_{\bar{s}}^{\circ}$ is at least $1/m$. Notice that if we can generate a point $\hat{v} \in H_{\bar{s}}^{\circ}$ with very good symmetry in $H_{\bar{s}}^{\circ}$ in the sense that $\alpha := \mathbf{sym}(\hat{v}, H_{\bar{s}}^{\circ}) = \Omega(1/m)$, we can then compute $\tilde{x}$ of Corollary 3 using at most $O\left(\sqrt{\vartheta}\ln\left(\vartheta\cdot m\right)\right)$ Newton steps, which is strongly polynomial-time. And even if we merely satisfy $\alpha := \mathbf{sym}(\hat{v}, H_{\bar{s}}^{\circ}) > \mathbf{sym}(0, H_{\bar{s}})$, we still may improve the computational effort needed to solve F by working with $F_{\hat{s}}$ rather than $F_{\bar{s}}$.

Of course, the effectiveness of this method depends entirely on the ability to efficiently compute a point $\hat{v} \in H_{\bar{s}}^{\circ}$ with good symmetry. The set $H_{\bar{s}}^{\circ}$ has the convenient representation $H_{\bar{s}}^{\circ} = \{v \in \mathbb{R}^m : \bar{s} - A^T v \in C^*\}$ from Proposition 8; furthermore, we have a convenient point $0 \in \mathrm{int}H_{\bar{s}}^{\circ}$ with which to start a method for finding a point with good symmetry; also, testing membership in $H_{\bar{s}}^{\circ}$ depends only on the capability of testing membership in C^*. Thus, the relative ease with which we can work with $H_{\bar{s}}^{\circ}$ suggests that excessive computation might not be necessary in order to compute a point $\hat{v}$ with good symmetry in $H_{\bar{s}}^{\circ}$. We explore several different approaches for computing such a point $\hat{v}$ in the following subsection.

3.3.1 Strategies for computing points with good symmetry in $H_{\bar{s}}^{\circ}$

In this subsection we presume that F has an interior solution, whereby $0 \in \text{int}H_{\bar{s}}$ and $H_{\bar{s}}^{\circ}$ will be bounded and $\mathbf{sym}(0, H_{\bar{s}}^{\circ}) > 0$. Recall that a symmetry point of a convex set $S \subset \mathbb{R}^m$ is a point x^* whose symmetry is optimal on S. From (iii) of Remark 10 we know that $\mathbf{sym}(x^*, S) \geq 1/m$. When $C = \mathbb{R}_+^n$, a symmetry point of $H_{\bar{s}}^{\circ}$ can be computed by approximately solving $n+1$ linear programs using the method developed in Section 2.4.2 of Chapter 1. Thus even for the case of $C = \mathbb{R}_+^n$ the computational burden of finding a point with guaranteed good symmetry appears to be excessive. In fact, the seemingly simpler task of just evaluating $\mathbf{sym}(x, S)$ at a particular point $x = \bar{x}$ might be hard for a general convex body S, see Section 2.4 of Chapter 1. Therefore, one is led to investigate heuristic and/or probabilistic methods, or perhaps methods that compute other types of points that lie "deep" in a convex body. Table 3.1 presents the symmetry guarantee for several types of deep points in a convex body; the computational effort for $C = \mathbb{R}_+^n$ ($H_{\bar{s}}^{\circ}$ is the intersection of n half-spaces) is shown in the third column of the table. We now briefly discuss each of these three possible choices for such points.

Analytic center approach

Starting at $v = 0$, we could use the damped Newton method outlined in [46] to compute an approximate analytic center of $H_{\bar{s}}^{\circ}$ using the barrier function $b(v) := f^*(\bar{s} - A^T v)$. We know from the theory of self-concordant barriers that the analytic center v^a of $H_{\bar{s}}^{\circ}$ has symmetry value at least $1/\sqrt{\vartheta(\vartheta - 1)}$ (see Lemma 5 of the Appendix of Nesterov, Todd, and Ye [48]). Each iteration of the algorithm would be of comparable computational burden as an interior-point iteration for the problem OP_μ and so it would probably be wisest to perform only a few iterations and hope that the final iterate v^f would have good symmetry value in $H_{\bar{s}}^{\circ}$ nevertheless.

Löwner-John center approach

The Löwner-John theorem guarantees the existence of an m-rounding of $H_{\bar{s}}^{\circ}$, i.e., an ellipsoid E centered at the origin and a point v^j with the property that $\{v^j\} + E \subset H_{\bar{s}}^{\circ} \subset \{v^j\} + mE$, see [23]. Such a point v^j is called a Löwner-John center and it follows that $\mathbf{sym}(v^j, H_{\bar{s}}^{\circ}) \geq 1/m$. In the case when C is either the nonnegative orthant $\mathbb{R}_+^n$ or is the cartesian product of second-order cones, we can compute such an approximate Löwner-John center by computing the center of the maximum volume inscribed ellipsoid in $H_{\bar{s}}^{\circ}$ via semidefinite programming (see Zhang and Gao [64], for example, for the case when $C = \mathbb{R}_+^n$). The problem with this approach is that the computational effort is likely to be substantially larger than that of solving the original problem $F_{\bar{s}}$, and the approach is limited to the case when C is the cartesian product of half-lines and/or second-order cones.

Center of Mass Approach.

The center of mass (or centroid) of a convex body $S \subset \mathbb{R}^m$ will have guaranteed symmetry of at least $1/m$, see [27]. Even when $C = \mathbb{R}_+^n$, computing the center of mass of $H_{\bar{s}}^{\circ}$ is #P-hard, see [12]. However, if we instead consider nondeterministic algorithms, then the recent work of Lovász and Vempala [41, 39, 42] points the way to computing points near the center of mass with high probability with good theoretical efficiency. This approach will be developed in more detail in Section 3.4.

Central Point	Symmetry Guarantee	Computational Effort when $C = \mathbb{R}_+^n$
Symmetry Point	$1/m$	$\approx$ LP $\times$ $(n+1)$
Analytic Center	$\frac{1}{\sqrt{\vartheta(\vartheta-1)}}$	$\approx$ LP
Löwner-John Center	$1/m$	$\approx$ SDP
Center of Mass	$1/m$	#P-Hard (deterministic) $\approx$ polynomial-time (stochastic)

Table 3.1: Summary Properties of Strategies for Computing Deep Points in $H_{\bar{s}}^{\circ}$

3.3.2 Polarity and projective transformations in Theorem 16

While it is obvious that $F_{\bar{s}}$ and $F_{\hat{s}}$ are related through the pair of projective transformations (3.7), it is perhaps not so obvious that the image sets $H_{\bar{s}}$ and $H_{\hat{s}}$ are related via projective transformations: $H_{\bar{s}}$ maps onto $H_{\hat{s}}$ with the following projective transformations between points $y \in H_{\bar{s}}$ and $y' \in H_{\hat{s}}$:

$$y' = T(y) := \frac{y}{1 - \hat{v}^T y} \quad \text{and} \quad y = T^{-1}(y') = \frac{y'}{1 + \hat{v}^T y'} \,. \tag{3.9}$$

This pair of projective transformations results from a more general theory concerning the polarity construction, translations, and projective transformations as follows, see Grünbaum [26] for example. Let S be a closed convex set containing the origin. ($S = H_{\bar{s}}$ in our context.) Then S° is a closed convex set containing the origin and $S^{\circ\circ} = S$. Let $\hat{v} \in \text{int}S^\circ$ be given. Then $(S^\circ - \{\hat{v}\})$ is the translation of S by $\hat{v}$, and also is a closed convex set containing the origin, and its polar is $(S^\circ - \{\hat{v}\})^\circ$. It is elementary arithmetic to show that S and $(S^\circ - \{\hat{v}\})^\circ$ are related through the projective transformations (3.9), namely $(S^\circ - \{\hat{v}\})^\circ = T(S)$ and $S = T^{-1}((S^\circ - \{\hat{v}\})^\circ)$. In other words, translation of the polar set corresponds to projective transformation of the original set, see Figure 3-3. This correspondence was previously used in [15, 16].

3.4 Approximate center of mass of $H_{\bar{s}}^\circ$ and its symmetry

In this section we present some general results about sampling from the uniform distribution on a given convex body $S \subset \mathbb{R}^d$, which are relevant for our particular case where $S = H_{\bar{s}}^\circ$ and $d = m$. We proceed as follows. A function $f : \mathbb{R}^d \to \mathbb{R}_+$ is said to be logconcave if $\log f$ is a concave function. A random variable $Z \in \mathbb{R}^d$ is called a logconcave random variable if the probability density function of Z is a logconcave function. Note that logconcave random variables are a broad class that includes Gaussian, exponential, and uniformly distributed random variables on convex

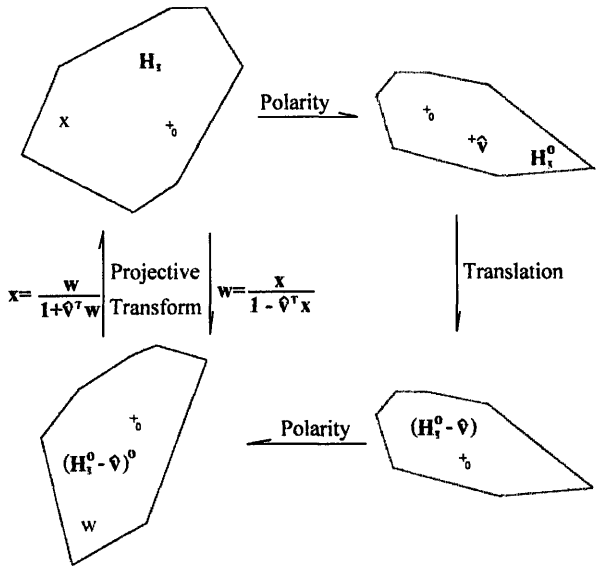

Figure 3-3: Translation of the polar set corresponds to projective transformation of the original set.

sets.

The center of mass (or centroid) and covariance matrix associated with Z are given respectively by

$$\mu_Z := E[Z] \qquad \text{and} \qquad \Sigma_Z := E[(Z - \mu_Z)(Z - \mu_Z)^T] \ .$$

The matrix Σ_Z is symmetric positive semi-definite. If Σ_Z is positive definite it can be used to define the ellipsoidal norm:

$$\|v\|_{\Sigma_Z} := \sqrt{v^T \Sigma_Z^{-1} v} \ .$$

The following are very useful properties of logconcave random variables.

Lemma 6 [36, 50, 51] *The sum of independent logconcave random variables is a logconcave random variable.*

Lemma 7 [41] *Let Z be a logconcave random variable in $\mathbb{R}^d$. Then for any $R \geq 0$:*

$$\mathsf{P}\left(\|Z - \mu_Z\|_{\Sigma_Z} \geq R\sqrt{d}\right) \leq e^{-R}.$$

Now let X be a random variable in $\mathbb{R}^d$ uniformly distributed on a convex body S, i.e., the probability density function of X is given by

$$f(x) = \frac{1}{\mathbf{Vol}(S)} 1_S(x) \ , \tag{3.10}$$

where $1_S(\cdot)$ is the indicator function of the set S. For simplicity, we denote its center of mass and covariance matrix respectively by μ and Σ, and note that Σ is positive definite since S has a non-empty interior. Let $B_\Sigma(x, r)$ denote the ball centered at x with radius r in the norm $\|\cdot\|_\Sigma$.

Lemma 8 [41] *Let X be a random variable uniformly distributed on a convex body $S \subset \mathbb{R}^d$. Then*

$$B_\Sigma\left(\mu \ , \sqrt{(d+2)/d}\right) \subset S \subset B_\Sigma\left(\mu \ , \sqrt{d(d+2)}\right) \ .$$

Assume that we are given M independent uniformly distributed random points $v^1, v^2, \ldots, v^M$ on the convex body S. We define the sample mean the usual way:

$$\hat{v} := \frac{1}{M} \sum_{i=1}^{M} v^i \ .$$

Lemma 9 *Let $\hat{v}$ be the sample mean of M independent uniformly distributed points on the convex body $S \subset \mathbb{R}^d$. Then*

$$\mathbf{sym}(\hat{v}, S) \geq \frac{\sqrt{(d+2)/d} - \|\hat{v} - \mu\|_\Sigma}{\sqrt{d(d+2)} + \|\hat{v} - \mu\|_\Sigma} \ .$$

Proof Consider any chord of S that passes through $\hat{v}$. It is divided by $\hat{v}$ into two segments of length s_1 and s_2. From Lemma 8 it follows that $B_\Sigma\left(\mu \ , \sqrt{(d+2)/d}\right) \subset$

$S \subset B_\Sigma\left(\mu\ , \sqrt{d(d+2)}\right)$. Thus, we can bound the ratio of s_1 to s_2 by

$$\frac{s_1}{s_2} \geq \frac{\sqrt{(d+2)/d} - \|\hat{v}-\mu\|_\Sigma}{\sqrt{d(d+2)} + \|\hat{v}-\mu\|_\Sigma} \ .$$

□

Theorem 17 *Let $\hat{v}$ be the sample mean of M independent uniformly distributed points on the convex body $S \subset \mathbb{R}^d$. Then for any $t \geq 0$ it holds that $\mathsf{P}(\|\hat{v}-\mu\|_\Sigma \geq t) \leq e^{-t\sqrt{\frac{M}{d}}}$.*

Proof Let $Y = \sqrt{M}\hat{v}$. Since $v^1, v^2, \ldots, v^M$ are independent uniformly distributed random variables on S, $E[Y] = \sqrt{M}\mu$ and Σ is the covariance matrix of Y. Moreover, using Lemma 6, Y is a logconcave random variable since it is a sum of independent logconcave random variables. Applying Lemma 7 using $R = t\sqrt{\frac{M}{d}}$ we obtain:

$$\begin{aligned} \mathsf{P}\left(\|\hat{v}-\mu\|_\Sigma \geq t\right) &= \mathsf{P}\left(\|\sqrt{M}\hat{v}-\sqrt{M}\mu\|_\Sigma \geq t\sqrt{M}\right) \\ &= \mathsf{P}\left(\|Y-E[Y]\|_\Sigma \geq R\sqrt{d}\right) \leq e^{-R} = e^{-t\sqrt{\frac{M}{d}}} \ . \end{aligned}$$

□

Corollary 4 *For any $\delta \in (0,1)$ and setting $M = 4d\left(\ln(1/\delta)\right)^2$, we have*

$$\mathbf{sym}(\hat{v}, S) \geq \frac{1}{2d+3}$$

with probability at least $1-\delta$.

Proof Using Theorem 17 with $M = 4d\left(\ln(1/\delta)\right)^2$ and $t = 1/2$ we obtain $\mathsf{P}(\|\hat{v}-\mu\|_\Sigma \geq 1/2) \leq \delta$, whereby $\mathsf{P}(\|\hat{v}-\mu\|_\Sigma \leq 1/2) \geq 1-\delta$. Finally, using Lemma 9 we obtain:

$$\mathbf{sym}(\hat{v}, S) \geq \frac{1-1/2}{d+1+1/2} = \frac{1}{2d+3}$$

with probability at least $1-\delta$. □

Remark 11 *The proof of Corollary 4 can be extended to show that setting* $M = \left(\frac{1+\frac{1-\varepsilon}{d}}{\varepsilon}\right)^2 d\left(\ln(1/\delta)\right)^2)$ *we obtain* $\mathbf{sym}(\hat{v}, S) \geq \frac{1-\varepsilon}{d}$ *with probability at least* $1-\delta$.

Keeping in mind the fact that $\mathbf{sym}(S)$ can only be guaranteed to be at most $1/d$ (and this bound is attained, for example, for a d-dimensional simplex), Corollary 4 gives an upper bound on the number of points that must be sampled to obtain a point $\hat{v}$ whose symmetry is bounded below by $\Omega(1/d)$ with high probability. Specializing to the case of $S = H_{\bar{s}}^{\circ}$ and $d = m$ and presuming that F has an interior solution (and hence $H_{\bar{s}}^{\circ}$ is a convex body), Corollary 4 provides a mechanism for achieving $\mathbf{sym}(\hat{v}, H_{\bar{s}}^{\circ}) = \Omega(1/m)$, and hence achieving $\mathbf{sym}(0, H_{\hat{s}}) = \Omega(1/m)$ with high probability (from Theorem 16). It follows from Corollary 3 and Corollary 4 that in the context of the Projective Pre-conditioning Method presented in Section 3.3, with high probability (i.e., probability at least $1-\delta$) we attain a complexity bound for solving $F_{\hat{s}}$ of

$$\left\lceil 9\sqrt{\vartheta}\ln\left(11\vartheta\left(2m+4\right)\right)\right\rceil$$

iterations of Newton's method. This iteration bound is strongly polynomial-time (with high probability). In order to make Corollary 4 constructive, we need a method for sampling on a convex body that obeys an approximate uniform distribution, which is discussed in the following subsection.

3.4.1 Sampling from the uniform distribution on a convex body

Herein we outline some relevant theory about uniform sampling on a convex body $S \subset \mathbb{R}^d$, see [41], [39], and [42] for recent results on this problem. We describe a standard sampling algorithm specialized to the structure of our application. To generate a random point distributed approximately uniformly on S, we will use a geometric random walk algorithm on S. The implementation of the algorithm requires only the use of a membership oracle for S and an initial point $X^0 \in S$. In the context of the Projective Pre-conditioning Method of Section 3.3, where $S = H_{\bar{s}}^{\circ}$ and $d = m$,

the initial point is $0 \in H^{\circ}_{\bar{s}}$. The requirement of a membership oracle for $S = H^{\circ}_{\bar{s}}$ is met if we have a membership oracle for the dual cone C^* as discussed earlier.

The geometric random walk algorithm known as "Hit-and-Run" generates iterates $X^1, X^2, \ldots$, as follows:

Geometric Random Walk Algorithm

Step 1. Initialize with $X^0 \in S$, $k = 0$

Step 2. Choose s uniformly on the unit sphere S^{d-1} in $\mathbb{R}^d$

Step 3. Let X^{k+1} be chosen uniformly on the line segment $S \cap \{X^k + ts : t \in \mathbb{R}\}$

Step 4. Set $k \leftarrow k + 1$, goto **Step 2**

It is a simple and well known result that this random walk induces a Markov chain whose stationary distribution is the uniform distribution on S. The rate of convergence to the stationary distribution depends on the spectral gap, i.e., the difference between the two largest eigenvalues of the transition kernel. Suppose that we run the algorithm for N iterations. In [39] it is proved that to achieve an ε-approximation to the uniform distribution density function (3.10) in the L_1 norm, it is sufficient that

$$N = O\left(d^3 \left(\frac{R}{r}\right)^2 \ln\left(\frac{R}{\varepsilon \cdot \mathrm{dist}_2(0, \partial S)}\right)\right) ,$$

where r, R satisfy $B_2(w, r) \subset S \subset B_2(v, R)$ for some $w, v \in S$, and $B_2(c, \delta)$, $\mathrm{dist}_2(v, T)$ are the Euclidean ball centered at c with radius δ and the Euclidean distance from v to T, respectively.

Note that Step 3 of the algorithm requires that one computes the end points of the line segment in S that passes through X^k and has direction s. This can be done by binary search using a membership oracle for S. In our case $S = H^{\circ}_{\bar{s}} = \{v : \bar{s} - A^T v \in C^*\}$ and the required membership oracle for S is met if we have a membership oracle for C^*. For self-scaled cones the endpoints computation in Step 3 is a standard computation: when $C = \mathbb{R}^n_+$ the endpoints computation is a min-ratio test, when C is the cross-product of second-order cones the endpoints computation uses the quadratic formula, and when C is the positive semidefinite cone the endpoints

computation is a min-ratio test of the eigenvalues of a matrix obtained after a single Cholesky factorization.

3.5 Computational results on randomly generated poorly-behaved problems

We performed computational experiments to assess the practical viability of the projective pre-conditioning method (PPM). We tested the PPM on 300 artificially generated homogeneous linear programming feasibility problems (i.e., $C = \mathbb{R}^n_+$). These 300 problems were comprised of 100 problems each of dimensions $(m, n) = (100, 500)$, $(500, 2500)$, and $(1000, 5000)$, and were generated so as to guarantee that the resulting problems would be poorly behaved. Each problem is specified by a matrix A and the chosen value of $\bar{s}$. We first describe how A was generated. Given a prespecified density value DENS for A, each element of A was chosen to be 0 with probability $1 - \text{DENS}$, otherwise the element was generated as a standard Normal random variable. We used DENS = 1.0, 0.01, and 0.01 for the problem dimensions $(m, n) = (100, 500)$, $(500, 2500)$, and $(1000, 5000)$, respectively. The vector $\bar{s}$ was chosen in a manner that would guarantee that the problem would be poorly behaved as follows. Starting with $s^0 = e$, we created the polar image set $H^\circ_{s^0} = \{v : A^T v \leq e\}$. We randomly generated a non-zero vector $d \in \mathbb{R}^m$ and performed a min-ratio test to compute $\bar{t} > 0$ for which $\bar{t}A^T d \in \partial H^\circ_{s^0}$. Then $\bar{s}$ is determined by the formula:

$$\bar{s} = s^0 - (1 - 4 \times 10^{-5})\bar{t}A^T d \, .$$

This method is essentially the reverse process of the PPM, and yields $\mathbf{sym}(0, H_{\bar{s}}) \leq 4 \times 10^{-5}$, with resulting poor geometry from Theorem 14.

We implemented the projective pre-conditioning method (PPM) using the following simplified version of the stochastic process described in Section 3.4: starting from $v^0 = 0 \in \text{int}H^\circ_{\bar{s}}$ we take K steps of the geometric random walk algorithm, yielding points $v^1, \ldots, v^K$, and computed $\hat{v} := \frac{1}{K}\sum_{i=1}^K v^i$, and then set $\hat{s} = \bar{s} - A^T\hat{v}$. We set

$K = 30$. It is also well known that this simple method yields $\hat{v} \to \mu$ as $K \to \infty$ and that the convergence results are similar to those described in Section 3.4. Nonetheless, the theoretical analysis is more technical and requires additional notation and so is omitted here. See [14] and [21] for discussion and further references.

We solved the 300 original problem instances of OP (stopping as soon as $\theta \geq 0$), as well as the resulting instances after pre-conditioning, using the interior-point software SDPT3 [58]. Tables 3.2 and 3.3 summarize our computational results. Because these problems are feasibility problems the number of IPM iterations is relatively small, even for the original problems. Notice that average IPM iterations shows a marked decrease in all three dimension classes, and in particular shows a 46% decrease in average IPM iterations for the 100 problem instances of dimension 1000×5000. The total running time (which includes the time for pre-conditioning using the geometric random walk) also shows a marked decrease when using the projective pre-conditioning method, and in particular shows a 33% decrease for the 100 problem instances of dimension 1000×5000. The last two columns of Table 3.2 shows the average value of θ^*. Given that $(x, \theta) = (\bar{x}, -1)$ is a feasible starting point for OP (and for SDPT3), θ^* is a good measure of the computational difficulty of a problem instance – a problem is poorly behaved to the extent that θ^* is close to zero. Here we see, regardless of any IPM, that θ^* increases by a factor of roughly 400, 800, and 600 for the problem dimensions $(m, n) = (100, 500)$, $(500, 2500)$, and $(1000, 5000)$, respectively. These results all demonstrate that by taking only a small number of steps of the geometric random walk algorithm, one can greatly improve the behavior of a poorly-behaved problem instance, and hence improve the practical performance of an IPM for solving the problem instance.

We also explored the sensitivity of the computational performance of the PPM to the number of steps of the random walk. Figure 3-4 shows the median number of IPM iterations as well as the 90% band (i.e., the band excluding the lower 5% and the upper 5%) of IPM iterations for the 100 problems of dimension 100×500 before and after pre-conditioning. Notice that only 10 steps of the random walk are needed to reduce the median and variance of the IPM iterations to a very low level.

Dimensions		Average IPM Iterations		Average Total Running Time (secs.)	
		Original	After	Original	After
m	n	Problem	Pre-conditioning	Problem	Pre-conditioning
100	500	8.52	4.24	0.5786	0.2983
500	2500	9.30	5.17	2.4391	2.0058
1000	5000	9.69	5.20	22.9430	15.3579

Table 3.2: Average Performance of SDPT3 on the 300 Problem Test-bed of Linear Programming Feasibility Problems. Computation was performed on a laptop computer running Windows XP.

Dimensions		Average Value of θ^*	
		Original	After
m	n	Problem	Pre-conditioning
100	500	0.0020	0.8730
500	2500	0.0012	1.0218
1000	5000	0.0019	1.1440

Table 3.3: Average Performance of SDPT3 on the 300 Problem Test-bed of Linear Programming Feasibility Problems. Computation was performed on a laptop computer running Windows XP.

As the number of random walk steps increase, the number of IPM iterations quickly concentrates and converges to a value below the 0.05 quantile for the original problem instances.

Figure 3-5 shows the median value of θ^* as well as the 90% band of θ^* values for the 100 problems of dimension 100×500 before and after pre-conditioning. As discussed earlier, θ^* is a good measure of problem instance behavior: larger values of θ^* indicate that the problem instance is better behaved, especially for computation via an IPM. The figure indicates that there is almost no improvement in of median value of θ^* after 50 steps of the random walk.

Figure 3-6 shows the median total running time as well as the 90% band of total running times for the 100 problems of dimension 100×500 before and after pre-conditioning. Notice that the median running time of the system with pre-conditioning rapidly decreases with a flat bottom in the range 10-100 steps of the random walk, after which the cost of the random walk steps exceeds the average ben-

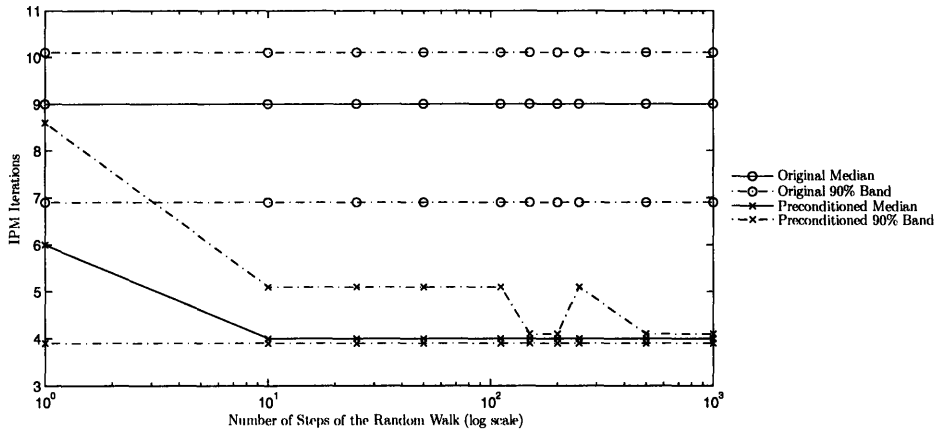

Figure 3-4: IPM iterations versus number of steps of the geometric random walk for the 100 problem instances of dimension 100×500.

efit from computing a presumably better pre-conditioner. Also notice, however, that the variation in running time decreases with the number of random steps, which may offer some advantage in lowering the likelihood of outlier computation times by using more random walk steps.

3.6 Summary/Conclusions/Other matters

In this chapter we have presented a general theory for transforming a normalized homogeneous conic system $F_{\bar{s}}$ to an equivalent system $F_{\hat{s}}$ via projective transformation induced by the choice of a point $\hat{v} \in H^{\circ}_{\bar{s}}$. Such a projective transformation serves to pre-condition the conic system into a system that has both geometric and computational behavior with certain guarantees. We have given a characterization of both the geometric behavior and the computational behavior of the transformed system as a function of the symmetry of $\hat{v}$ in the image set $H^{\circ}_{\bar{s}} = \{v : \bar{s} - A^T v \in C^*\}$. Because $H^{\circ}_{\bar{s}}$ must contain a point v whose symmetry is at least $1/m$, if we can find a point whose symmetry is $\Omega(1/m)$ then we can projectively transform the conic system to one whose geometric and computational complexity behavior will be strongly-polynomial in m and the complexity value ϑ of the barrier function $f(\cdot)$ of the cone C. We have

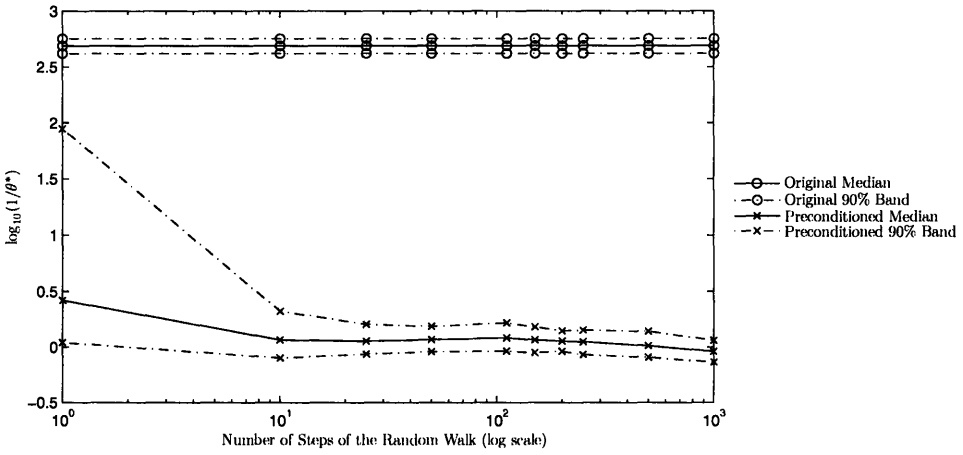

Figure 3-5: $\log(\theta^*)$ versus number of steps of the geometric random walk for the 100 problem instances of dimension 100×500.

presented a method for generating such a point $\hat{v}$ using sampling on geometric random walks on $H_{\bar{s}}^{\circ}$ with associated complexity analysis. Finally, we have implemented this methodology on randomly generated homogeneous linear programming feasibility problems, constructed to be poorly behaved. Our computational results indicate that the projective pre-conditioning methodology holds the promise to markedly reduce the overall computation time for conic feasibility problems; for instance we observe a 46% improvement in average IPM iterations for the 100 problem instances of dimension 1000×5000. The next step in this line of research will be to develop a suitable adaptation of the methods developed herein to solve conic optimization problems, and to test such an adaptation on conic optimization problems that arise in practice.

3.6.1 Infeasible case

The theory presented herein is based on the assumption that F has a solution. When F does not have a solution, then one can consider the alternative/dual system:

$$F_a : \begin{cases} A^T v + s &= 0 \\ s &\in C^* \backslash \{0\} \, . \end{cases}$$

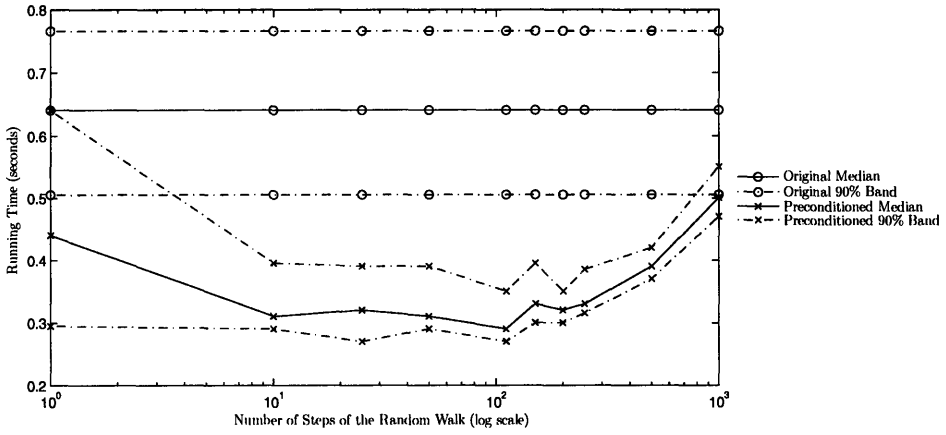

Figure 3-6: Total running time versus number of steps of the geometric random walk for the 100 problem instances of dimension 100×500.

This system can then be re-formatted as:

$$F_a^{'}: \left\{ \begin{array}{rcl} Bs & = & 0 \\ s & \in & C^* \backslash \{0\} \end{array} \right. ,$$

for a suitably computed matrix B whose null-space is the orthogonal complement of the null-space of A. Note that $F_a^{'}$ is of the same format as F and the results for F can be easily adapted to $F_a^{'}$. (Actually, the computation of B is not necessary. Given $\bar{x} \in \text{int}C$, consider the analogous image sets for F_a and $F_a^{'}$ defined as $H_{\bar{x}} := \{A^T v + s : s \in C^*, \bar{x}^T s = 1, v \in \mathbb{R}^m\}$ and $H_{\bar{x}}^{'} := \{Bs : s \in C^*, \bar{x}^T s = 1\}$. Then $\mathbf{sym}(0, H_{\bar{x}}) = \mathbf{sym}(0, H_{\bar{x}}^{'})$ even though $H_{\bar{x}}$ is unbounded, and one can work with $H_{\bar{x}}$ and problem F_a directly.) Nevertheless, it may be more fruitful and revealing to develop a different projective pre-conditioner theory designed directly for the dual form F_a, and this is a direction for future research.

3.6.2 Related complexity matters

Nesterov [45] has suggested the following "dual approach" to solving (3.1): starting at $v^0 = 0$ compute an approximate analytic center v^a of $H_{\bar{s}}^{\circ}$, which is the essentially

unconstrained problem $\min_v\{f^*(\bar{s} - A^T v) : \bar{s} - A^T v \in \text{int}C^*\}$. It is elementary to show from the barrier calculus that as soon as a point v is computed whose Newton step $(\Delta v, \Delta s)$ satisfies $\sqrt{(\Delta s)^T H^*(s)\Delta s} < 1$ (where $s = \bar{s} - A^T v$ and $H^*(s)$ is the Hessian of $f^*(\cdot)$ at s), then the Newton step multipliers yield an interior solution of F. Regarding the complexity of this scheme, it follows from an analysis that is almost identical to that yielding inequality (2.19) of [52] that the number of Newton steps of a short-step IPM to compute an approximate analytic center is:

$$O\left(\sqrt{\vartheta}\ln\left(\frac{\vartheta}{\mathbf{sym}(0, H_{\bar{s}}^{\circ})}\right)\right) = O\left(\sqrt{\vartheta}\ln\left(\frac{\vartheta}{\mathbf{sym}(0, H_{\bar{s}})}\right)\right)$$

(from (ii) of Remark 10), which is of the same order as the complexity of Algorithm A from Theorem 15. These complexity bounds depend on $\mathbf{sym}(0, H_{\bar{s}})$ to bound the complexity of traversing the central path via a short-step method. As is shown in Nesterov and Nemirovski [47], a generically more accurate complexity bound can be found by analyzing the central path via its Riemannian geometry. However, as is demonstrated in the current work, $\mathbf{sym}(0, H_{\bar{s}})$ lends itself to analysis, characterization, and ultimately manipulation and reduction via the constructive projective pre-conditioning method shown herein. An interesting research challenge is to develop analogous tools/methods to work with and reduce the Riemannian distance of the central path as developed in [47].

Chapter 4

Efficiency of a re-scaled perceptron algorithm for conic systems

In this chapter, the object of interest is to compute a solution of the following conic system

$$\begin{cases} Ax & \in & \textbf{int } K \\ x & \in & X \end{cases} \tag{4.1}$$

where X and Y are Euclidean subspaces, $K \subset Y$ is a regular closed convex cone, and $A : X \to Y$ is a linear operator. The goal is to compute an interior element in the cone of feasible solutions denoted by $\mathcal{F} = \{x \in X : Ax \in K\}$. Important special cases of this framework includes feasibility problems for linear programming (LP), second order cone programming (SOCP) and semidefinite positive programming (SDP).

The ellipsoid method ([34]), random walk method ([6]), and interior-point methods (IPMs) ([32], [46]) are examples of methods which solve (4.1) in polynomial time. Nonetheless, these methods differ a lot in their practical performance and representation requirements. For example, a membership oracle suffices for the ellipsoid method and random walk method, while a barrier function for K is required to implement an IPM. The latter is by far the most successful algorithm for conic programming in practice. In particular, applications of SDP range over several fields such as optimal control, eigenvalue optimization, combinatorial optimization and many others.

In the case of $X = \mathbb{R}^n$ and $K = \mathbb{R}^m_+$, we recover the well studied case of a system of linear inequalities. Within this context, the simple perceptron algorithm provides a convergent procedure whose complexity depends on the square of the inverse of the *width* of the cone of feasible solutions τ of (4.1). In [11], the perceptron algorithm is combined with a sequence of re-scalings which gradually increase τ (on average). These re-scalings are constructed based on near feasible points. This improved algorithm reduces the dependence on $1/\tau$ from polynomial to logarithmic.

We will extend [11] to the conic setting. Although the probabilistic analysis is not affected, this is not the case for the remainder of the analysis. The improvement obtained in [11] can be seen to arise from a clever use of a *deep separation oracle*, which is stronger than the usual separation oracle used in the classical perceptron algorithm. The deep separation oracle is the most notable difference between the linear case and the general case. Under the linear programming framework, there is no difference between the implementation of both oracles. Thus, several important issues were not required to be resolved in [11].

Here we investigate in detail the deep separation oracle and reveal its properties and structure. Based on these properties, we propose an iterative scheme for general convex cones which exploits the particular structure of the deep separation oracle. We show that it solves the deep separation oracle in polynomial time, while requiring only a deep separation oracle for the dual cone of K (which is readily available for many cones of interest such as the second order cone and the cone of positive semi-definite matrices).

We start with properties of convex cones that will be used throughout the chapter. Section 4.2 generalizes the classical perceptron algorithm for the conic setting while Section 4.3 extends the re-scaling phase of [11]. In Section 4.4, we go over the probabilistic analysis for sake of completeness. Sections 4.5 and 4.6 are devoted to properties and construction of the deep separation oracle in general.

4.1 Preliminaries

4.1.1 Notation

We will confine our analysis to finite dimensional Euclidean spaces. Let X and Y will be Euclidean spaces with finite dimension n and m. Denote by $\|\cdot\|$ their Euclidean norms, and $\langle\cdot,\cdot\rangle$ their Euclidean inner product. For $\bar{x} \in X$, $B(\bar{x}, r)$ will denote the ball centered at $\bar{x}$ with radius r; analogously for Y. Let $A : X \to Y$ denote a linear operator, and $A^* : Y \to X$ denote the adjoint operator associated with A.

4.1.2 Convex Cones

Definition 1 *A set C is said to be a cone if for every $t \geq 0$ we have $tC \subset C$. Moreover, it is a convex cone if C is also convex.*

Now we introduce several standard concepts associated with convex cones, see [53]. The dual cone of a convex cone C is defined as

$$C^* = \{d : \langle x, d\rangle \geq 0, \text{ for all } x \in C\} \tag{4.2}$$

while **ext** C denote the set of extreme rays of the cone C. A cone is pointed if it contains no lines. We say that C is a *regular* cone if C is a pointed closed convex cone with non-empty interior. It is elementary to show that C is regular if and only if C^* is regular.

Given a regular convex cone C, we use the following conditioning measure.

Definition 2 *If C is a regular cone in X, the* width *of C is given by*

$$\tau_C \triangleq \max_{x,r} \left\{ \frac{r}{\|x\|} : B(x, r) \subset C \right\}$$

Among convex cones, we will be particular interested in the non-negative orthant $\mathbb{R}^m_+$, the second order cone denoted by $Q^n = \{x \in \mathbb{R}^n : \|(x_1, x_2, \ldots, x_{n-1})\| \leq x_n\}$, and the cone of positive semi-definite matrices $S^{k\times k}_+ = \{X \in \mathbb{R}^{k\times k} : X^T =$

X, $\langle v, Xv\rangle \geq 0$, for all $v \in \mathbb{R}^k\}$. Recall that these three cones are self-dual and their width are respectively $1/\sqrt{m}$, $1/\sqrt{2}$, and $1/\sqrt{k}$.

Next we establish a few properties of convex cones. These properties are well-known but they are presented in the format needed for our analysis.

Lemma 10 *Suppose C is a regular convex cone. Then*

$$\textbf{int}\ (C \cap C^*) \neq \emptyset.$$

Proof We proceed by contradiction assuming that **int** $(C \cap C^*) = \emptyset$. Since C is regular, C^* is also regular. Thus, both cones have non-empty interior. In this case, **int** $(C \cap C^*) = \emptyset$ implies that **int** $C \cap$ **int** $C^* = \emptyset$. Therefore, there exists h, $\|h\| \neq 0$,

$$\langle h, x\rangle \geq 0 \text{ for all } x \in C \quad \text{and} \quad \langle h, y\rangle \leq 0 \text{ for all } y \in C^*$$

By definition of C^*, $h \in C^*$. On the other hand, $\langle h, h\rangle \leq 0$ by the second property above. Therefore, $h = 0$ (a contradiction with $\|h\| = 1$). □

The following characterization will be used in our analysis.

Lemma 11 *Let $T = \{A^*\lambda : \lambda \in C^*\}$. Then* **cl** $(T) = \{x : Ax \in C\}^*$.

Proof Denote $\mathcal{M} = \{x : Ax \in C\}^*$. We will prove **cl** $(T) = \mathcal{M}$.

($\subseteq$) Let $\lambda \in C^*$. Then for every x satisfying $Ax \in C$, $\langle x, A^*\lambda\rangle = \langle Ax, \lambda\rangle \geq 0$, since $Ax \in C$ and $\lambda \in C^*$. Thus, **cl** $(T) \subseteq \mathcal{M}$ since $\mathcal{M}$ is closed.

($\supseteq$) Assume that there exists $y \in \mathcal{M} \backslash$ **cl** (T). Thus, there exists h, $\|h\| \neq 0$, satisfying

$$\langle h, y\rangle < 0 \quad \text{and} \quad \langle h, w\rangle \geq 0 \text{ for all } w \in \textbf{cl}\ (T).$$

Notice that $\langle h, A^*\lambda\rangle \geq 0$ for all $\lambda \in C^*$. Thus, $Ah \in C$. On the other hand, since $y \in \mathcal{M}$, $\langle h, y\rangle \geq 0$ contradicting $\langle h, y\rangle < 0$. □

The question of sets of the form T being closed has been recently studied by Pataki [49]. Necessary and sufficient conditions for T to be a closed set are given in

[49] when C^* belongs to af class called "nice cones", a class which includes polyhedra and self-scaled cones. Nonetheless, the set T may fail to be closed even in simple examples.

Example 2 *Let* $C^* = Q^3 = \{(\lambda_1, \lambda_2, \lambda_3) \mid \|(\lambda_1, \lambda_2)\| \leq \lambda_3\}$ *and* $A = \begin{bmatrix} -1 & 0 \\ 0 & 1 \\ 1 & 0 \end{bmatrix}$. *In this case,* $T = \{A^T\lambda \mid \lambda \in C^*\} = \{(-\lambda_1 + \lambda_3, \lambda_2) \mid \|(\lambda_1, \lambda_2)\| \leq \lambda_3\}$. *It is easy to verify that* $(0,1) \notin T$ *but* $(\varepsilon, 1) \in T$ *for every* $\varepsilon > 0$ *(set* $\lambda_1 = \frac{1}{2\varepsilon} - \frac{\varepsilon}{2}$, $\lambda_2 = 1$, *and* $\lambda_3 = \frac{1}{2\varepsilon} + \frac{\varepsilon}{2}$*) which shows that* T *is not closed.*

Let $\bar{z}$ be a unit vector which achieves the width of a cone C as defined in Definition 2. The next two lemmas characterize points that are *deep* in C.

Lemma 12 *If* $y \notin C - \sigma\bar{z}$, *there exists* $d \in C^*$, $\|d\| = 1$, *and* $\langle d, y\rangle < -\sigma\tau_C$.

Proof Since $y + \sigma\bar{z} \notin C$, there exists d, $\|d\| = 1$, such that $\langle d, y + \sigma\bar{z}\rangle < 0$ and $\langle d, x\rangle \geq 0$ for all $x \in C$. This automatically implies that $d \in C^*$. Since $\bar{z} - \tau_C d \in C$, $\langle d, \bar{z} - \tau_C d\rangle \geq 0$ which implies that $\langle d, \bar{z}\rangle \geq \tau_C$. Thus, $\langle d, y\rangle < -\sigma \langle d, \bar{z}\rangle \leq -\sigma\tau_C$. □

Lemma 13 $B(z, r) \subseteq C$ *if and only if* $\langle d, z\rangle \geq r\|d\|$ *for all* $d \in C^*$.

Proof Suppose $B(z,r) \subset C$. Let $d \in C^*$. Then, $z - r\frac{d}{\|d\|} \in C$ and since $d \in C^*$, $\left\langle d, z - r\frac{d}{\|d\|}\right\rangle \geq 0$. Thus, $\langle d, z\rangle \geq r\frac{\langle d,d\rangle}{\|d\|} = r\|d\|$. Conversely, suppose $\langle d, z\rangle \geq r\|d\|$ for every $d \in C^*$. Let v such that $\|v\| \leq r$. Assume $z + v \notin C$, then there exists $d \in C^*$, $\langle d, z + v\rangle < 0$. Therefore

$$\langle d, z\rangle < -\langle d, v\rangle \leq r\|d\|,$$

which contradicts $\langle d, z\rangle \geq r\|d\|$. □

4.1.3 Oracles

In our complexity analysis, we will distinguish three different types of oracles.

Definition 3 *A* membership oracle *for a convex set $S \subset \mathbb{R}^n$ is a subroutine that given a point $x \in \mathbb{R}^n$, identifies if $x \in S$ or if $x \notin S$.*

Definition 4 *A* separation oracle *for a convex set $S \subset \mathbb{R}^n$ is a subroutine that given a point $x \in \mathbb{R}^n$, identifies if $x \in S$ or returns a vector $d \in \mathbb{R}^n$, $\|d\| = 1$, such that*

$$\langle d, x\rangle \leq \langle d, y\rangle \text{ for all } y \in S.$$

Definition 5 *For a fixed positive scalar t, a* deep separation oracle *for a cone $C \subset \mathbb{R}^n$ is a subroutine that given a non-zero point $x \in \mathbb{R}^n$, either*

*(I) correctly identifies that $\dfrac{\langle d, x\rangle}{\|d\|\|x\|} \geq -t$ for all $d \in$ **ext** C^**

or

(II) returns a vector $d \in C^*$, $\|d\| = 1$, satisfying $\dfrac{\langle d, x\rangle}{\|d\|\|x\|} \leq -t$.

Definitions 3 and 4 are standard in the literature while Definition 5 is not. The notion of a *deep separation oracle* was motivated by the need to proper extend the analysis in [11] for conic systems. For instance, we point out that conditions (I) and (II) are not exclusive.

4.2 Perceptron algorithm for conic systems

We use the following notation to represent respectively the feasible cone of solutions for (4.1) and its dual cone

$$\mathcal{F} = \{x \in X : Ax \in K\} \quad \text{and} \quad \mathcal{F}^* = \{s \in X^* : \langle s, x\rangle \geq 0, \text{ for all } x \in \mathcal{F}\}.$$

Also, let $(\tau_{\mathcal{F}}, \bar{z}) \in \arg\max_{r,x} \{r : \|x\| = 1, B(x, r) \subseteq \mathcal{F}\}$; $\tau_{\mathcal{F}}$ is the width if the feasibility cone and $\bar{z}$ is the center of $\mathcal{F}$.

The perception algorithm was proposed to solve a homogeneous system of linear inequalities ((4.1) with $K = \mathbb{R}^m_+$). It is well-known that it has finite termination in at most $1/\tau^2_{\mathcal{F}}$ iterations. Nonetheless, this bound can be exponentially large in the bit model.

Our starting point in this study is to show that the perceptron algorithm can be easily extended to the case of conic systems like (4.1).

Perceptron Algorithm for Conic Systems

Step 1. Let x be the origin in X.

Step 2. If $Ax \in \mathbf{int}\ K$ Stop.

Step 3. Call separation oracle for $\mathcal{F}$ at x.

The oracle returns $d \in \mathcal{F}^*$, $\|d\| = 1$, such that $\langle d, x \rangle \leq 0$, set $x \leftarrow x + d$.

Step 4. Goto **Step 2.**

Lemma 14 *The perceptron algorithm for conic systems returns a feasible point in at most* $\lfloor 1/\tau_{\mathcal{F}}^2 \rfloor$ *iterations.*

Proof Consider the potential function $\langle x, \bar{z} \rangle / \|x\| \leq 1$. If the algorithm does not stop at **Step 2**, we update x to $x + d$, whereby

$$\langle x + d, \bar{z} \rangle = \langle x, \bar{z} \rangle + \langle d, \bar{z} \rangle \geq \langle x, \bar{z} \rangle + \tau_{\mathcal{F}},$$

and

$$\|x + d\|^2 = \langle x, x \rangle + 2 \langle x, d \rangle + \langle d, d \rangle \leq \langle x, x \rangle + 1,$$

since $\langle x, d \rangle \leq 0$, $\langle d, d \rangle = 1$, and $\langle d, \bar{z} \rangle \geq \tau_{\mathcal{F}}$ from Lemma 13.

After k iterations, the potential function is at least $k\tau_{\mathcal{F}}/\sqrt{k}$. After more than $\lfloor 1/\tau_{\mathcal{F}}^2 \rfloor$ iterations, the potential function would be greater than one, a contradiction. Thus, the algorithm must terminate after at most $\lfloor 1/\tau_{\mathcal{F}}^2 \rfloor$ iterations returning a feasible $x \in \mathbf{int}\ \mathcal{F}$. □

In the linear case ($K = \mathbb{R}_+^m$), the separation oracle it suffices to check each constraint individually. In the general case, we cannot consider one component of Ax at a time and the complexity of the separation oracle will depend on the structure of the cone K. Using a separation oracle for K itself, if $Ax \notin K$ one can find a non-zero vector $h \in Y^*$ separating Ax from K, i.e. $\langle h, Ax \rangle < 0$ and $h \in K^*$, therefore $d = A^*h \in X^*$ is a non-zero vector which separates x from the cone of feasible solution

$\mathcal{F}$. Thus, the perceptron algorithm can be properly implemented if a separation oracle for K is available.

Example 3 *Consider the case of $K = S_+^{k\times k}$ being the cone of positive semi-definite matrices and $A : \mathbb{R}^n \to S^{k\times k}$ a linear operator. In order to compute a direction $d \in \mathcal{F}^*$, we start by computing any eigenvector v of the symmetric matrix Ax associated with a non-positive eigenvalue. Then the vector $d = A^*(vv^T)$ is such that*

$$\langle d, x\rangle = \langle A^*(vv^T), x\rangle = \langle vv^T, Ax\rangle = \mathbf{tr}(vv^T Ax) = \mathbf{tr}(v^T(Ax)v) \leq 0,$$

and for all $y \in \mathcal{F}$ we have:

$$\langle d, y\rangle = \langle vv^T, Ay\rangle = \mathbf{tr}(v^T(Ay)v) \geq 0,$$

i.e., $d \in \mathcal{F}^*$, and $\langle d, x\rangle \leq 0$. If (4.1) has a solution it easily follows that $d \neq 0$ whereby $d/\|d\|$ can be used in **Step 3** *of the Perceptron Algorithm.*

4.3 Re-scaling the conic system

In order to obtain an algorithm with a better complexity result with respect to the width of the cone of solutions $\mathcal{F}$, we will systematically re-scale the system (4.1) through the use of a suitably constructed random vector that approximates $\bar{z}$.

For motivation purposes suppose that we are given a point $\bar{v} \in \mathbf{int}\ (\mathcal{F} \cap \mathcal{F}^*)$. The existence of such point is guaranteed under our regularity conditions, see Lemma 10. Now, for $\mu > 0$, let $\hat{A} = A \circ (I + \mu\bar{v}\bar{v}^T)$ define a new conic system (4.1) where solutions are denoted by $\hat{\mathcal{F}}$. Then for any $x \in \mathcal{F}$, we have that

$$A(I + \mu\bar{v}\bar{v}^T)x = Ax + \mu A\bar{v}(\bar{v}^T x) \in K.$$

That is, the set of solutions $\hat{\mathcal{F}}$ contains the set of solutions $\mathcal{F}$ of the original conic system. In fact, one can easily show that if $\tau_{\mathcal{F}} > 0$ we have $\tau_{\hat{\mathcal{F}}} > \tau_{\mathcal{F}}$. Actually, letting $\mu \to \infty$, $\tau_{\mathcal{F}}$ can be made arbitrary close to one as pointed out in [11] for the linear

case. Nonetheless, that approach is not practical since one must know a solution for the original system to implement it. Surprisingly, it turns out that nearly feasible solutions will suffice to construct transformations which will increase the width of $\mathcal{F}$.

The linear transformation that will use was first proposed in [11] for the case of linear inequalities. Here, we will extend these ideas to our conic setting. In Table 4.3 we describe a generalization of the algorithm proposed in [11].

The perceptron improvement phase requires a deep separation oracle for $\mathcal{F}^*$ instead of the separation oracle as required by the perceptron algorithm. We point out that now one needs to balance the values of $\|d\|$ and $\langle d, \bar{x} \rangle$ at the same time. This suggests that additional work is required beyond a deep separation oracle for K itself which is usually easier to implement. Again, in the linear case investigated in [11], there is no difference between the separation oracle and the deep separation oracle due to the particular structure of the cone $\mathbb{R}^m_+$. A more detailed study of the deep separation oracle is deferred to the Sections 4.5 and 4.6.

In this section we assume that a deep separation oracle is available and proceed to analyze the impact of the re-scaling phase on the width of the cone of solutions. The next lemma quantifies such impact.

Lemma 15 *Let $\tau_{\mathcal{F}}$, $\tau_{\hat{\mathcal{F}}}$ be the width of the feasibility cones $\mathcal{F}$, $\hat{\mathcal{F}}$ of two consecutive iterations of the re-scaled perception algorithm and A, $\hat{A}$ are the respective linear operators. Then*

$$\tau_{\hat{\mathcal{F}}} \geq \frac{(1-\sigma)}{\sqrt{1+3\sigma^2\|\hat{z}\|}}\tau_{\mathcal{F}}$$

where $\hat{z} = \bar{z} + \frac{1}{2}\left(\tau_{\mathcal{F}} - \left\langle \frac{x}{\|x\|}, \bar{z} \right\rangle\right)\frac{x}{\|x\|}$, and x is the output of the perceptron improvement phase.

Proof At the end of the perception improvement phase, we have a vector x satisfying

$$\frac{\langle d, x \rangle}{\|d\|\|x\|} \geq -\sigma \quad \text{for all } d \in \textbf{ext } \mathcal{F}^*.$$

Step 1 Initialization
Step 1.1 Set $B = I$ and $\sigma = 1/(32n)$.

Step 2 Perceptron Algorithm for Conic Systems
Step 2.1 Let x the origin in X. Repeat at most $(1/\sigma^2)$ times.
Step 2.2 If $Ax \in$ **int** K Stop.
Step 2.3 Call separation oracle for $\mathcal{F}$ at x.
The oracle returns $d \in \mathcal{F}^*$, $\|d\| = 1$, such that $\langle d, x \rangle \leq 0$, set $x \leftarrow x + d$.

Step 3 Stop Criteria
Step 3.1 If $Ax \in$ **int** K then output Bx and **stop**.

Step 4 Perceptron Improvement Phase
Step 4.1 Let x be a random unit vector in X
Step 4.2 Repeat at most $(1/\sigma^2)\ln(n)$ times:
Call deep separation oracle for $\mathcal{F}$ at x with $t = \sigma$.
If oracle returns condition (I), end **Step 4**.
If oracle returns a vector $d \in \mathcal{F}^*$, $\|d\| = 1$, such that $\langle d, x \rangle \leq -\sigma\|x\|$, set $x \leftarrow x - \langle d, x/\|x\| \rangle\, d$.
If $x = 0$, go back to **Step 4.1**.
Step 4.3 Call deep separation oracle for $\mathcal{F}$ at x with $t = \sigma$.
If oracle returns condition (II), restart at **Step 4.1**.

Step 5 Stop Criteria
Step 5.1 If $Ax \in$ **int** K then output Bx and **stop**.

Step 6 Re-scaling the System
Step 6.1 Update $A \leftarrow A \circ \left(I + \frac{xx^T}{\langle x, x \rangle}\right)$ and $B \leftarrow B \circ \left(I + \frac{xx^T}{\langle x, x \rangle}\right)$.

Step 7 Loop
Step 7.1 Goto **Step 2**.

Table 4.1: One iteration of the re-scaled perceptron algorithm is one pass of **Step 2-6**.

Let $\bar{x} = x/\|x\|$. Then $\langle d, \bar{x}\rangle \geq -\sigma\|d\|$ for all $d \in \mathbf{ext}\ \mathcal{F}^*$. From Lemma 13, it holds that

$$\frac{\langle d, \bar{z}\rangle}{\|d\|\|\bar{z}\|} = \frac{\langle d, \bar{z}\rangle}{\|d\|} \geq \tau_{\mathcal{F}} \quad \text{for all } d \in \mathcal{F}^*,$$

i.e. $\langle d, \bar{z}\rangle \geq \tau_{\mathcal{F}}\|d\|$ for all $d \in \mathcal{F}^*$.

From Lemma 11 it holds that

$$\langle \lambda, A\bar{z}\rangle = \langle A^*\lambda, \bar{z}\rangle \geq \tau_{\mathcal{F}}\|A^*\lambda\| \quad \text{for all} \quad \lambda \in K^*.$$

Note that $\hat{z} = \bar{z} + \frac{1}{2}(\tau_{\mathcal{F}} - \langle \bar{x}, \bar{z}\rangle)\bar{x}$ and let $\hat{\tau} := \frac{(1-\sigma)}{\sqrt{1+3\sigma^2}}\tau_{\mathcal{F}}$.

We want to show that

$$\langle v, \hat{z}\rangle \geq \hat{\tau}\|v\| \quad \text{for all} \quad v \in \mathbf{ext}\ \mathcal{F}^*. \tag{4.3}$$

If condition (4.3) is true then, by convexity of the function $f(v) = \hat{\tau}\|v\| - \langle v, \hat{z}\rangle$, it will also be true that $\langle v, \hat{z}\rangle \geq \hat{\tau}\|v\|$ for any $v \in \mathcal{F}^*$. Then from Lemma 13 it follows that $B(\hat{z}, \hat{\tau}) \subset \mathcal{F}$, whereby

$$\tau_{\hat{\mathcal{F}}} \geq \frac{\hat{\tau}}{\|\hat{z}\|} \quad \text{as desired.}$$

Let v be an extreme ray of $\mathcal{F}^*$. Using Lemma 11, there exist a sequence $\{\lambda^i\}_{i\geq 1}$, $\lambda^i \in K^*$, $A^*\lambda^i \to v$ as $i \to \infty$. Since (4.3) is trivially true for $v = 0$, we can assume that $v \neq 0$ and hence $A^*\lambda^i \neq 0$ for i large enough.

Next note that

$$\|\hat{A}^*\lambda^i\|^2 = \|A^*\lambda^i\|^2 + 2\left\langle A^*\lambda^i, \bar{x}\right\rangle^2 + \langle \bar{x}, \bar{x}\rangle \left\langle A^*\lambda^i, \bar{x}\right\rangle^2 = \|A^*\lambda^i\|^2 \left(1 + 3\left(\frac{\langle A^*\lambda^i, \bar{x}\rangle}{\|A^*\lambda^i\|}\right)^2\right)$$

and

$$\begin{aligned}
\left\langle \hat{A}^*\lambda^i, \hat{z}\right\rangle &= \left\langle A^*\lambda^i, \hat{z}\right\rangle + \langle \bar{x}, \hat{z}\rangle \left\langle A^*\lambda^i, \bar{x}\right\rangle \\
&= \left\langle A^*\lambda^i, \bar{z}\right\rangle + (\tau_{\mathcal{F}} - \langle \bar{x}, \bar{z}\rangle)\left\langle A^*\lambda^i, \bar{x}\right\rangle + \langle \bar{x}, \bar{z}\rangle \left\langle A^*\lambda^i, \bar{x}\right\rangle \\
&\geq \tau_{\mathcal{F}}\|A^*\lambda^i\| + \tau_{\mathcal{F}}\left\langle A^*\lambda^i, \bar{x}\right\rangle \\
&= \tau_{\mathcal{F}}\left(1 + \frac{\langle A^*\lambda^i, \bar{x}\rangle}{\|A^*\lambda^i\|}\right)\|A^*\lambda^i\|.
\end{aligned} \tag{4.4}$$

Therefore $\frac{\left\langle \hat{A}^*\lambda^i, \hat{z} \right\rangle}{\|\hat{A}^*\lambda^i\|} \geq \tau_{\mathcal{F}} \frac{1+t_i}{\sqrt{1+3t_i^2}}$ where $t_i = \frac{\langle A^*\lambda^i, \bar{x} \rangle}{\|A^*\lambda^i\|}$. Note that $t_i \leq 1$ and

$$\langle v, \bar{x} \rangle \geq -\sigma \|v\| \text{ since } v \in \textbf{ext } \mathcal{F}^*$$

and so $\frac{\langle v, \bar{x} \rangle}{\|v\|} \geq -\sigma$ and by continuity, for any $\varepsilon > 0$, there exists i such that $t_i \geq -\sigma - \varepsilon$. Thus, $t_i \in [-\sigma - \varepsilon, 1]$ for i large enough.

For $t \in [0, 1]$, we have that $\frac{1+t}{\sqrt{1+3t^2}} \geq \frac{1+t}{\sqrt{1+2t+t^2}} = 1$.

For $t \in [-\sigma - \varepsilon, 0]$, the function $g(t) = \frac{1+t}{\sqrt{1+3t^2}} \geq \frac{1-\sigma-\varepsilon}{\sqrt{1+3(\sigma+\varepsilon)^2}}$ since

$$\frac{dg(t)}{dt} = \frac{1-3t}{(1+3t^2)^{3/2}} \geq 0$$

for $t \in [-\sigma - \varepsilon, 0]$, that is, $g(t)$ is increasing on $[-\sigma - \varepsilon, 0]$.

Therefore, for i large enough we have

$$\frac{\left\langle \hat{A}\lambda^i, \hat{z} \right\rangle}{\|\hat{A}^*\lambda^i\|} \geq \tau_{\mathcal{F}} \frac{(1-\sigma-\varepsilon)}{\sqrt{1+3(\sigma+\varepsilon)}}.$$

Passing to the limit on $i \to \infty$, so $\lambda^i \to v$, we have

$$\frac{\langle v, \hat{z} \rangle}{\|v\|} \geq \tau_{\mathcal{F}} \frac{(1-\sigma-\varepsilon)}{\sqrt{1+3(\sigma+\varepsilon)}}$$

and so

$$\frac{\langle v, \hat{z} \rangle}{\|v\|} \geq \tau_{\mathcal{F}} \frac{(1-\sigma)}{\sqrt{1+3\sigma}} = \hat{\tau}.$$

□

4.4 Probabilistic analysis

As mentioned before, the probabilistic analysis of our conic framework is similar to the analysis with linear inequalities in [11]. Although a few changes are required, all the main ideas are still valid. For sake of completeness, we go over some results of [11].

Our exposition intentionally separates the probabilistic analysis from the remanning sections.

The first lemma of this section was established in [7] to the linear case and here is generalized for the conic framework. Roughly speaking, it shows how the perceptron improvement phase can generate near feasible solutions given that it is started with a good initial point, which happens with at least a fixed probability.

Lemma 16 *Let z be a feasible solution for (4.1) with norm one. With probability at least $\frac{1}{8}$, the perception improvement phase returns a vector x, $\|x\| \leq 1$, such that*

(i) $\langle d, x\rangle \geq -\sigma\|x\|$ *for every* $d \in \mathbf{ext}\ \mathcal{F}^*$, $\|d\| = 1$;

(ii) $\langle z, x/\|x\|\rangle \geq \frac{1}{\sqrt{n}}$.

Proof If x^0 is a random unit vector in $\mathbb{R}^n$, that is the initial value of x at Step 4.1, then with probability at least $1/8$ we have $\langle z, x^0\rangle \geq 1/\sqrt{n}$ [11]. Notice that in the perceptron improvement phase

$$\langle x - \langle d, x\rangle\, d, z\rangle = \langle x, z\rangle - \langle d, x\rangle\, \langle d, z\rangle \geq \langle x, z\rangle$$

since $\langle d, x\rangle \leq 0$ and $\langle d, z\rangle \geq 0$ (since $d \in \mathcal{F}^*$ and $z \in \mathcal{F}$). Thus, the inner product in $\langle z, x\rangle$ does not decrease. Also, in each inner iteration of the perceptron improvement phase (Step 4), the norm of x does decrease by at least a constant factor:

$$\begin{aligned}
\langle x - \langle x, d\rangle\, d, x - \langle x, d\rangle\, d\rangle &= \langle x, x\rangle - 2\,\langle d, x\rangle^2 + \langle d, x\rangle^2\,\langle d, d\rangle \\
&= \langle x, x\rangle - \langle d, x\rangle^2 = \langle x, x\rangle - \langle d, x/\|x\|\rangle^2\,\langle x, x\rangle \\
&\leq \langle x, x\rangle\,(1 - \sigma^2).
\end{aligned}$$

where we used that $0 \geq \langle d, x/\|x\|\rangle \geq -\sigma$ and $\|d\| = 1$.

Thus, after $\lceil (1/\sigma^2)\ln(n)\rceil$ iterations, we would have that $\frac{\langle x, z\rangle}{\|x\|} > 1$ which is a contradiction since z is a unit vector. Therefore, we terminate with a vector x satisfying (i) and (ii) with probability at least $1/8$. $\square$

Lemma 16 establishes that points obtained after the perceptron improvement phase are near feasible for the current conic system. The next lemma clarifies the implications of using these near feasible points to re-scale the conic system.

Lemma 17 *Suppose that $\tau_{\mathcal{F}}, \sigma \leq 1/32n$ and A is the linear operator of the current iteration. Let $\hat{A}$ be the linear operator obtained after one iteration of the perceptron improvement phase. If we denote by $\tau_{\hat{\mathcal{F}}}$ the width of the cone of feasible solutions for the new system associated with $\hat{A}$, then*

(i) $\tau_{\hat{\mathcal{F}}} \geq \left(1 - \frac{1}{32n} - \frac{1}{512n^2}\right)\tau_{\mathcal{F}}$;

(ii) With probability at least $\frac{1}{8}$, $\tau_{\hat{\mathcal{F}}} \geq \left(1 + \frac{1}{3n}\right)\tau_{\mathcal{F}}$.

Proof Let x be the output of the perceptron improvement phase. For simplicity, let $\tau := \tau_{\mathcal{F}}$, $\hat{\tau} := \tau_{\hat{\mathcal{F}}}$, and $\bar{x} = x/\|x\|$. Using Lemma 15, we have that

$$\hat{\tau} \geq \frac{(1-\sigma)}{\sqrt{1+3\sigma}\|\hat{z}\|}\tau$$

where $\hat{z} = \bar{z} + \frac{1}{2}(\tau - \langle \bar{x}, \bar{z} \rangle)\bar{x}$. Next note that

$$\|\hat{z}\|^2 = 1 + (\tau - \langle \bar{x}, \bar{z} \rangle) + \frac{1}{4}(\tau - \langle \bar{x}, \bar{z} \rangle)^2 = 1 + \frac{\tau^2}{4} + \langle \bar{z}, \bar{x} \rangle \left(\frac{\tau}{2} - \frac{3}{4}\langle \bar{z}, \bar{x} \rangle\right).$$

Following [11], consider two cases. First assume that $|\langle \bar{z}, \bar{x} \rangle| < 1/\sqrt{n}$ which happens with probability at most $7/8$. Then viewing the above as a quadratic function in $\langle \bar{z}, \bar{x} \rangle$ which is maximal when $\langle \bar{z}, \bar{x} \rangle = \tau/3$

$$\|\hat{z}\|^2 \leq 1 + \frac{\tau^2}{4} + \frac{\tau^2}{12} = 1 + \frac{\tau^2}{3}.$$

Thus, we have that

$$\hat{\tau} \geq \tau(1-\sigma)\left(1 - \frac{3\sigma^2}{2}\right)\left(1 - \frac{\tau^2}{6}\right) \geq \tau\left(1 - \frac{1}{32n} - \frac{1}{512n^2}\right),$$

since τ and σ are less or equal to $\frac{1}{32n}$, and $\frac{1}{\sqrt{1+t}} \geq 1 - \frac{t}{2}$.

The second case assumes that $|\langle \bar{z}, \bar{x} \rangle| \geq 1/\sqrt{n}$, which happens with probability at least $1/8$. In this case, the quadratic function in $\langle \bar{z}, \bar{x} \rangle$ will be maximal at $\langle \bar{z}, \bar{x} \rangle = \frac{1}{\sqrt{n}}$ which yields

$$\|\hat{z}\|^2 \leq 1 - \frac{3}{4n} + \frac{\tau}{2\sqrt{n}} + \frac{\tau^2}{4}.$$

Now we have (again using $\frac{1}{\sqrt{1+t}} \geq 1 - \frac{t}{2}$) that

$$\hat{\tau} \geq \tau(1-\sigma)\left(1 - \frac{3\sigma^2}{2}\right)\left(1 + \frac{3}{8n} - \frac{\tau}{4\sqrt{n}} - \frac{\tau^2}{8}\right) \geq \tau\left(1 + \frac{1}{3n}\right).$$

□

Before proceeding to the main complexity result, we establish that each perceptron improvement phase will terminate rapidly with high probability.

Lemma 18 *The perceptron improvement phase terminates in at most* $\lceil 8 \ln \frac{1}{\delta} \rceil$ *iterations with probability at least* $1 - \delta$.

Proof Note that if we start with a vector x such that $\langle z, x \rangle \geq \frac{1}{\sqrt{n}}$, we terminate the procedure, and this happens with probability at least $1/8$. Thus, the probability of not terminating after k iterations is at most $(7/8)^k$. A simple computation yields $k = 8 \ln 1/\delta$ to bound this probability by δ. □

The following theorem bounds the number of overall iterations and the number of oracle calls made by the algorithm.

Theorem 18 *Assuming that (4.1) is feasible, the algorithm properly terminates in at most*

$$T = \max\left\{4096 \ln\left(\frac{1}{\delta}\right), 1000n \ln\left(\frac{1}{32n\tau_{\mathcal{F}}}\right)\right\} = O\left(n \ln \frac{1}{\tau_{\mathcal{F}}} + \ln \frac{1}{\delta}\right)$$

iterations of the re-scaled algorithm with probability at least $1 - \delta$. *Moreover, the algorithm makes at most* $O(T\ n^2\ \ln(n)\ \ln(1/\delta))$ *calls of a deep separation oracle for* $\mathcal{F}$ *and at most* $O(T\ n^2)$ *calls of a separation oracle for* $\mathcal{F}$ *with probability* $1 - \delta$.

Proof This proof will follow closely the proof of Theorem 3.4 in [11]. Let τ_k denote the width of the feasible cone after k iterations and let z^k denote a unit vector that achieves this width. Moreover, let T denote the total number of iterations, and set $\sigma = 1/(32n)$.

Note that if $\tau_k \geq \frac{1}{\sigma}$, the (classical) perceptron algorithm will produce a strictly feasible solution in at most $(1/\sigma^2) = 1024n^2$ iterations, where each iteration makes only one call to the separation oracle. For each perceptron improvement phase, with probability $1 - \delta$, we need at most $\lceil 8\ln(1/\delta) \rceil$ trials to succeed by Lemma 18. In each trial the number of deep separation oracle calls is bounded by $(1/\sigma^2)\ln(n) = 1024n^2\ln(n)$. Thus it suffices to bound the number of iterations T until which $\tau_T \geq 1/\sigma^2$.

Let $U_i = 1_{\langle x, z^k \rangle \geq \frac{1}{\sqrt{n}}}$, where x is a random unit vector on X, be a binary random variable and $U = \sum_{i=1}^{T} U_i$. Note that Lemma 17 (ii) implies that $U_i = 1$ yields $\tau_{i+1} \geq \tau_i(1 + 1/(3n))$. Furthermore, Lemma 17 (ii) implies that $P(U_i = 1) \geq 1/8$ whereby $E[U] \geq T/8$. Now, the Chernoff bound yields

$$P(U < (1-\epsilon)E[U]) \leq e^{-\epsilon^2 E[U]/2} \leq e^{-\epsilon^2 T/16}.$$

In order to bound this probability by δ and setting $\epsilon = 1/16$, we need $T \geq 4096\ln(1/\delta)$.

Thus, with probability at least $1 - \delta$, using Lemma 17 (i) and (ii), we have

$$\begin{aligned}
\tau_T &\geq \tau_0 \left(1 + \tfrac{1}{3n}\right)^U \left(1 - \tfrac{1}{32n} - \tfrac{1}{512n^2}\right)^{T-U} \\
&\geq \tau_0 \left(1 + \tfrac{1}{3n}\right)^{T(1-\epsilon)/8} \left(1 - \tfrac{1}{32n} - \tfrac{1}{512n^2}\right)^{T - T(1-\epsilon)/8} \\
&\geq \tau_0 \left(1 + \tfrac{1}{3n}\right)^{\frac{15T}{128}} \left(1 - \tfrac{1}{32n} - \tfrac{1}{512n^2}\right)^{\frac{113T}{128}} \\
&\geq \tau_0 e^{T/1000n}.
\end{aligned}$$

Setting $T \geq 1000n\ln(1/(32n\tau_0))$ we obtain $\tau_T \geq 1/32n$. □

4.5 Properties of the deep separation oracle for $\mathcal{F}$

Up to now we assumed that a deep separation oracle for $\mathcal{F}$ was readily available. Here, we will study in more detail this crucial step of the perceptron improvement phase.

Although it is tempting to try to solve this problem by constructing the following optimization problem

$$\begin{cases} \min\limits_{d} & \langle d, \bar{x} \rangle \\ & \|d\| \leq 1 \\ & d \in \mathcal{F}^*, \end{cases} \tag{4.5}$$

this is far from being necessary. In fact, any algorithm that approximate (4.5) within a constant factor (or a factor bounded by a polynomial in the dimension) would lead to an efficient algorithm. Moreover, even in the case of linear inequalities considered in [11], this approach would lead to a nontrivial problem. This has important algorithmic implications since it is sufficient to verify each constraint individually, in contrast to consider all possible normalized convex combinations of the constrains.

Another strong evidence for that is the dual problem associated with (4.5). This dual can be formulated as the following projection problem

$$\begin{aligned} \max_{\beta} \quad & -\|\bar{x} - \beta\| \\ & A\beta \in K \\ & \beta \in X. \end{aligned} \tag{4.6}$$

This suggest that both problems (4.5) and (4.6) are no easier than our original problem (4.1).

From the definition of the deep separation oracle, it suffices to search only over the extreme rays of $\mathcal{F}^*$. In fact, any set $\mathcal{U}$ that contains $\mathbf{ext}\ \mathcal{F}^*$ also suffices. A fundamental question is how to choose a set $\mathcal{U}$ which leads to an efficient implementation of the oracle. An elementary result that has important algorithmic consequences is the following.

Proposition 9 *If the cone of interest can be written as $K = K_1 \times K_2$, then*

$$\textbf{ext } \mathcal{F}^* \subseteq \{A^*(\textbf{ext } K_1 \times \{0\}), A^*(\{0\} \times \textbf{ext } K_2)\}.$$

This property is particularly interesting since in many cases we can rewrite the cone K as the cartesian products of simpler cones (as is the case for linear inequalities). Thus, one can decompose the problem for each one of these simpler cones.

Example 4 *Consider the following system*

$$\begin{cases} A_0 x & \in & \mathbb{R}^m_+ \\ A_i x & \in & Q^{n_i} \quad i = 1, \ldots, L \\ x & \in & S^{k \times k}_+. \end{cases}$$

The implementation of the deep separation oracle decomposes into smaller cones. For $A_1 x \in \mathbb{R}^m_+$, we recover the setup of [11] and one can check each constraint individually. For each $A_i x \in Q^{n_i}$, we show in Theorem 19 how to solve it efficiently. Finally, $x \in S^{k \times k}_+$ reduced to the computation of the minimum eigenvalue of the matrix x.

Example 5 *We consider the implication of the previous theorem for a deep separation oracle in the case of SDP. Recall the maximum separation oracle (4.5). Using the representation of Lemma 11, we have that (4.5) is*

$$\begin{cases} \min\limits_{\lambda} & \langle \lambda, A\bar{x} \rangle \\ & \|A^*\lambda\| \leq 1 \\ & \lambda \succcurlyeq 0 \end{cases} \tag{4.7}$$

where λ is a positive semi-definite matrix. As suggested before, this is not easier than our original problem. We can restrict the minimization to the extreme rays of $\mathcal{F}^*$. In particular, it would be sufficient to consider only rank one matrices in (4.7):

$$\begin{cases} \min\limits_{v} & v^T(A\bar{x})v \\ & \sum_{i=1}^{n}(v^T A_i v)^2 \leq 1 \\ & v \in \mathbb{R}^m. \end{cases} \tag{4.8}$$

Although this seems to be a simpler problem, we were not able to solve it (or approximate its solution) efficiently via a simple method.

4.6 A deep separation oracle for general cones

The results of the previous section suggest that there is no simple deep separation oracles for general cones. Moreover, the deep separation oracle can be cast as a conic system of the form of (4.1). Motivated by these reasons, we propose to implement the deep separation oracle based on simple iterative methods. In particular, we have two candidates: the perceptron algorithm, and the re-scaled perceptron algorithm itself, the first being more convenient if we can ensure that the width of the system (HS) is large and the latter when no such guarantee is available.

Of course, we need to provide oracles needed to implement each one of these methods for the new system (HS). It turns out that we will be able to implement this new deep separation oracle efficiently requiring only a deep separation oracle for the cone K^*. We point out that for the second order cone and the cone of positive semi-definite matrices these oracles are simple to construct.

The problem of interest is to compute an interior solution of the following feasibility system,

$$(HS)(t) \begin{cases} -\langle \lambda, Ax \rangle - t\mu & \geq & 0 & (a) \\ -\|A^*\lambda\| + \mu & \geq & 0 & (b) \\ \lambda & \in & K^* & (c) \\ (\lambda, \mu) & \in & Y \times \mathbb{R} & \end{cases} \tag{4.9}$$

where t is associated with the violation of the first inequality. The perceptron improvement phase has been using $t = \sigma$, but it will be convenient to work with a

generic t instead.

We start by converting (4.9) to the standard form of (4.1). Define the following linear map $\tilde{A} : Y \times \mathbb{R} \to \mathbb{R} \times Q^{n+1} \times Y$ as follows.

$$\tilde{A} \begin{pmatrix} \lambda \\ \mu \end{pmatrix} = \begin{pmatrix} -\langle Ax, \lambda \rangle - t\mu \\ A^*\lambda \\ \mu \\ \lambda \end{pmatrix}.$$

Thus, the system (4.9) is equivalent to

$$\tilde{A} \begin{pmatrix} \lambda \\ \mu \end{pmatrix} \in \widetilde{K} = \mathbb{R}_+ \times Q^{n+1} \times K^* \tag{4.10}$$

Note that given an interior solution $(\hat{\lambda}, \hat{\mu})$, one can easily construct a solution for the original deep separation oracle using $\frac{\hat{\lambda}}{\hat{\mu}}$. This is well defined because we must have $\hat{\mu} \geq \|A^*\hat{\lambda}\| > 0$ in any feasible solution due to the constraint (a) in (4.9).

4.6.1 Separation oracle for (4.10)

The construction of a separation oracle for $(HS)(t)$ is simple. Since the cone of interest $\widetilde{K}$ is the cartesian product of three simpler cones, we can consider each one separately. It is clear that if a point cannot be separated for each cone individually, this point must be feasible.

- **Cone (a):** this cone is associated with a linear inequality and the normal vector associated with it is $(-Ax, -t)^T$.

- **Cone (b):** this is a second order cone constraint. We can derive the following normal vector

$$\begin{cases} (0,1)^T, & \text{if } \|A^*\lambda\| = 0, \\ \left(-\dfrac{AA^*\lambda}{\|A^*\lambda\|}, 1\right)^T, & \text{otherwise.} \end{cases}$$

• **Cone (c):** the last cone is K^* itself and, by assumption, we have a deep separation separation oracle.

4.6.2 Deep separation oracle for (4.10)

In order to construct a deep separation oracle for (HS) we will take advantage of the particular structure of the problem. The new cone of interest is $\widetilde{K} = \mathbb{R}_+ \times Q^{n+1} \times K^*$. Thus, based on the Corollary 9, the deep separation oracle for (4.10) can be implemented for each cone independently.

• **Cone (a):** this cone is associated with a linear inequality and its normal vector is $(-Ax, -t)^T$.

• **Cone (b):** this is a particular second order cone constraint. By assumption, we have $\|A^*\bar{\lambda}\| > \bar{\mu}$ otherwise the constraint is not violated. In order to solve it we will formulate the following optimization problem

$$\begin{aligned} \min_{v,w} \quad & \langle v, A^*\bar{\lambda} \rangle + w\bar{\mu} \\ & \|Av\|^2 + w^2 \leq 1 \\ & \|v\| \leq w. \end{aligned}$$

Since we can restrict to the extreme rays of the second order cone, we can assume that the last inequality is binding. This problem simplifies to the following optimization problem:

$$\begin{aligned} \min_{v} \quad & \langle v, A^*\bar{\lambda} \rangle + \bar{\mu}\|v\| \\ & v^T(A^*A + I)v \leq 1 \end{aligned} \tag{4.11}$$

We obtain the following result based on Newton's method and binary search. We extend the analysis of Ye [62], which in turn was based on a work of Smale [55], to the problem (4.11).

Theorem 19 *We have three cases for the problem (4.11) with a precision of $\varepsilon > 0$:*
(i) if $\bar{\mu} = 0$, (4.11) can be solved in $O(n^3)$;
(ii) if $\bar{\mu} > 0$, (4.11) can be solved in $O(n^3 + n^2 \ln\ln(1/\varepsilon))$;

(iii) if $\bar{\mu} < 0$, (4.11) can be solved in $O(n^4 + n^2 \ln\ln(1/\varepsilon))$.

The proof of Theorem 19 can be found in the Appendix E.2.

Remark 12 *Note that one does not need to consider the case of $\bar{\mu} < 0$. Let $\bar{y} = (\bar{\lambda}, \bar{\mu})^T$ be the current point to be separated. In the case of $\bar{\mu} < 0$, we can (weakly) separate $\bar{y}$ by the hyperplane defined by $d = (0, 1)$. After making the update, we have the next iterate $\bar{y}^+ = \bar{y} - \langle d, \bar{y}\rangle\, d = (\bar{\lambda}, 0)^T$. However, if z achieves the width associated with the cone of solutions of $(HS)(\sigma)$, we have*

$$\langle z, \bar{y}^+\rangle = \langle z, \bar{y}\rangle - \langle d, \bar{y}\rangle\, \langle d, z\rangle \geq \langle z, \bar{y}\rangle \quad \textit{and} \quad \|\bar{y}^+\|^2 \leq \|\bar{y}\|^2.$$

Thus, the algorithm can continue to operate as is as long as we do not count this simple update to the total of updates made in the perceptron improvement phase. Since the number of these updates is bounded by $(1/\sigma^2)\ln(n)$, we can make at most $(1/\sigma^2)\ln(n)$ simple updates during each perceptron improvement phase.

- **Cone (c):** the last cone is K^* itself. A deep separation oracle is available for K^* by assumption.

4.6.3 Bounding Complexity and Infeasibility

One needs to address the complexity of solving each one of these problems. In fact, we need to deal not only with the inconvenient case of the system (HS) being infeasible, but also the cases where the width of the cone of solutions might be arbitrary close to zero. It turns out that is possible to efficiently solve both cases simply by exploiting the structure of our problem.

Consider the following "basic" feasibility conic system

$$(B)\begin{cases} -\|A^*\lambda\| + \mu \;\geq\; 0 & (a) \\ \lambda \;\in\; K^* & (c) \\ (\lambda, \mu) \;\in\; Y \times \mathbb{R}, \end{cases} \tag{4.12}$$

which will be the same for all x that we encounter in (HS).

Lemma 19 *Let $\tau_{(B)}$ be the width of the cone formed by the solutions of (4.12). Then,*

$$\tau_{(B)} \geq \frac{\tau_{K^*}}{4\max\{\|A^*\|, 1\}}.$$

Proof Let $(z_K^*, \tau_{K^*}) \in \arg\max_{r,x}\{r : \|x\| = 1, B(x,r) \subset K^*\}$.

Letting $a = 4\max\{\|A^*\|, 1\}$, define $z_{(B)} = (z_K^*/a, \sqrt{1 - 1/a^2})$, $\|z_{(B)}\| = 1$. Then, for any $(h, v) \in Y \times \mathbb{R}$, $\|(h, v)\| \leq 1$,

$$\left\| A^* \left(\frac{z_K^*}{a} + \frac{\tau_{K^*}}{a} h \right) \right\| \leq \frac{\|z_K^*\|}{4} + \frac{\tau_{K^*}}{4}\|h\| \leq \frac{1}{2} \leq \sqrt{1 - 1/a^2} - \tau_{K^*}\frac{|v|}{4}$$

Thus, every point in $B\left(z_{(B)}, \frac{\tau_{K^*}}{4\max\{\|A^*\|, 1\}}\right)$ is feasible for (4.12).□

The previous lemma shows that the conic system (B) cannot have an arbitrary bad conditioning. In turn, we show that if the conic system $(HS)(\sigma)$ is feasible, the conic system $(HS)(\sigma/2)$ also cannot have an arbitrary bad conditioning. This will lead to a efficient approximation algorithm for the deep separation oracle.

Lemma 20 *Let τ be the width of $\mathcal{F}$. Assume that $(HS)(\sigma)$ is feasible and let $t \in (0, \sigma)$. The width of the cone associated with $(HS)(t)$ is at least*

$$\tau_{(HS)(t)} \geq \tau \frac{\tau_{K^*}(\sigma - t)}{32\max\{\|A^*\|, 1\}\max\{\|A\|, 1\}}$$

Proof Denote as $\mathcal{F}_{(B)}$ (respectively $\mathcal{F}_{(HS)(t)}$) the cone of solutions for the (B) conic system (respectively $(HS)(t)$ conic system). Let $\hat{z} = (\hat{\lambda}, \hat{\mu}) \in \mathcal{F}_{(HS)(\sigma)}$ be a point with norm one and let $z_{(B)}$, $\|z_{(B)}\| = 1$, be such that $B\left(z_{(B)}, \tau_{(B)}\right) \subset \mathcal{F}_{(B)}$.

We have that $\mathcal{F}_{(HS)(t)} \subset \mathcal{F}_{(HS)(\sigma)} \subset \mathcal{F}_{(B)}$. Thus $\hat{z}$ is also feasible for $(HS)(t)$ and (B). Therefore, for any $\alpha \in [0, 1]$

$$B\left((1 - \alpha)\hat{z} + \alpha z_{(B)}, \alpha\tau_{(B)}\right) \subset \mathcal{F}_{(B)}.$$

Next consider the linear constraint in $(HS)(t)$, denoted here by $p(z) \geq 0$. By definition, $p(\hat{z}) \geq (\sigma - t)\hat{\mu}$. Therefore, any element $z \in B\left(\hat{z}, \frac{(\sigma-t)\hat{\mu}}{\max\{\|A\|,1\}}\right)$ must satisfy $p(z) \geq 0$.

By choosing $\alpha = \frac{(\sigma - t)\hat{\mu}}{4\max\{\|A\|,1\}}$, we have

$$B\left((1-\alpha)\hat{z} + \alpha z_{(B)}, \alpha\tau_{(B)}\right) \subset B\left(\hat{z}, \alpha(\|z_{(B)}\| + \|\hat{z}\|) + \alpha\tau_{(B)}\right) \subset B\left(\hat{z}, \frac{(\sigma - t)\hat{\mu}}{\max\{\|A\|, 1\}}\right),$$

since $\|z_{(B)}\| + \|\hat{z}\| \leq 2$ and $\tau_{(B)} \leq 1$.

Thus, any element $z \in B\left((1-\alpha)\hat{z} + \alpha z_{(B)}, \alpha\tau_{(B)}\right)$ is feasible for $(HS)(t)$. Note that $\hat{\mu} \geq \|A^*\hat{\lambda}\|$ where $A^*\lambda \in \mathcal{F}^*$ and $\hat{\mu}^2 + \|\hat{\lambda}\|^2 = 1$. Let z^* be a unit vector that achieves the width τ of the cone $\mathcal{F}$. Using Lemma 13 with $z = z^*$, $C = \mathcal{F}$ and $d = A^*\hat{\lambda} \in \mathcal{F}^*$, we have that

$$\hat{\mu} \geq \|A^*\hat{\lambda}\| \geq \left\langle z^*, A^*\hat{\lambda} \right\rangle \geq \tau\|\hat{\lambda}\|.$$

In fact, it holds that $\hat{\mu} \geq \tau/2$ since $\hat{\mu}$ and $\|\hat{\lambda}\|$ cannot be both less than $1/2$. Therefore,

$$\tau_{(HS)(t)} \geq \alpha\tau_{(B)} = \tau\frac{(\sigma - t)}{8\max\{\|A\|, 1\}}\tau_{(B)}.$$

The result follows from Lemma 19. □

Several important remarks should be made with respect to this result and its implications to the re-scaled perceptron algorithm. The conditioning of the system $(HS)(t)$ is comparable to the conditioning of the original system. The key difference is that the deep separation oracle is implementable for $(HS)(t)$ as long as a deep separation oracle is available for K^*. As a side note, we point out that one can always re-scale A so that $\|A\| = 1$ without affecting the analysis.

Remark 13 *Suppose we choose to use $(HS)(t)$ for some $t < \sigma$ instead of $(HS)(\sigma)$. This will increase the number of inner iterations in the perceptron improvement phase to $(1/t^2)\ln(n)$.*

An important issue is to detect infeasibility of the system $(HS)(\sigma)$. The algorithm to solve $(HS)(t)$ can be run for a pre-specified amount of iterations. It will perform correctly, with high probability, if we pre-specify correctly a bound for the width of

the original system. On the other hand, if one does not find a solution for the global system, we have an upper bound on the width of the original system.

We point out that as we re-scale the original system (4.1), we are also improving the conditioning of the system $(HS)(t)$ when $(HS)(\sigma)$ is feasible. Finally, the following corollaries summarize the complexity results to solve the deep separation oracle.

Corollary 5 *Let $t < \sigma$ be fixed. Suppose that the conic system $(HS)(\sigma)$ is feasible. Then the perceptron algorithm will find a feasible solution for $(HS)(t)$ in at most*

$$\left\lceil \left(\frac{32 \max\{\|A^*\|, 1\} \max\{\|A\|, 1\}}{\tau \ \tau_{K^*}(\sigma - t)} \right)^2 \right\rceil$$

iterations.

Corollary 6 *Let $t < \sigma$ be fixed. Suppose that the conic system $(HS)(\sigma)$ is feasible. Then the re-scaled perceptron algorithm will find a feasible solution for $(HS)(t)$ in at most*

$$O\left(n \ln \frac{\max\{\|A^*\|, 1\} \max\{\|A\|, 1\}}{\tau \ \tau_{K^*}(\sigma - t)} + \ln \frac{1}{\delta} \right)$$

iterations with probability at least $1 - \delta$.

Corollary 7 *Let $t < \sigma$ be fixed. Suppose that the re-scaled perceptron algorithm cannot find a solution for the system $(HS)(t)$ after k iterations and $(HS)(\sigma)$ is feasible. Then the width of the original system is at most*

$$O\left(\frac{\max\{\|A^*\|, 1\} \max\{\|A\|, 1\}}{\delta^{1/n} \tau_{K^*}(\sigma - t)} e^{-k/n} \right)$$

with probability at least $1 - \delta$.

Chapter 5

Norm-induced densities and testing the boundedness of a convex set

We say that a probability density function f is norm-induced if

$$f_t(x) = \begin{cases} \dfrac{e^{-t\|x\|}dx}{\int_K e^{-t\|y\|}dy}, & x \in K \\ 0, & x \notin K \end{cases} \tag{5.1}$$

for some norm $\|\cdot\|$, where K is a convex set in $\mathbb{R}^n$, and $t > 0$ is a parameter. In this chapter we will explore the connection between these densities and geometric properties of the convex set K itself (usually as a function of the parameter t). Our results make use of the geometry of unbounded convex sets. For instance, given an arbitrary unbounded convex set we show that most its points are contained in any *enlargement* of its recession cone. That simple geometric phenomenon motivates many of our results.

Moreover, a density function of the form (5.1) is logconcave and proper by construction. Thus, a random variable whose density distribution function is f_t can be efficiently simulated (at least approximately) by geometric random walk algorithms [40]. In turn theoretical results on f_t can be used to construct (implementable) algorithms to test properties of K itself.

Herein we develop an algorithm to test if a given convex set $K \subset \mathbb{R}^n$ is bounded or

unbounded. In either case the algorithm will provide us with an associated certificate of boundedness or unboundedness based on the properties established in Section 5.2. We emphasize that algorithms for this problem will be closely related to the representation used to describe the set K. Our interest lies in cases in which the convex set is given only by a membership oracle. This (minimal assumption) framework is the standard framework in several applications in the Computer Science literature. Furthermore, it covers many other problems of interest like convex feasibility.

The decision problem of testing for boundedness has a variety of interesting consequences. In recent years, several probabilistic methods have been proposed to compute quantities of interest like centroid, volume [43], convex optimization [31], and many others [6], in the context of convex bodies. In all these cases, boundedness of a convex set is a fundamental assumption for whose testing our algorithm provides a constructive approach. Khachiyan establishes the equivalence between a strongly polynomial algorithm for linear programming and a strongly polynomial algorithm for testing unboundedness of a convex set associated with a system of inequalities [33]. Moreover, linear homogeneous conic feasibility problems of the form

$$\begin{cases} Ax & \in & C \\ x & \in & \mathbb{R}^n \backslash \{0\} \end{cases} \tag{5.2}$$

can be converted into our framework by defining

$$K = \{x \in \mathbb{R}^n : Ax + h \in C\} \tag{5.3}$$

for any $h \in \mathbf{int}\ C$. In this case, $0 \in \mathbf{int}\ K$, and the recession cone of K coincides with the set of feasible solutions of the original system (5.2). In this case, only a membership oracle for the cone C would be required to conduct the algorithm.

The implementability of our algorithm relies on the ability to sample random variables distributed according to a probability density f_t. Over the last decade many important developments on sampling from log-concave densities, most notably via geometric random walks, have been observed. In particular, the *hit-and-run*

random walk has been extensively analyzed and polynomial rates of convergence have been established for this particular random walk under the log-concavity assumption [37, 41, 42, 40]. Besides, the homotopy analysis proposed here is similar to the analysis done by Lovász and Vempala in [43] of the algorithm they called Reverse Annealing, which was applied to the problem of computing the volume of a (bounded) convex body. However, our approach differs from [43] with two respects: by using a different density family, and by dealing explicitly with the possible unboundedness of K.

In the presence of additional structure, other algorithms are available in the literature. For example, assuming that a *self-concordant barrier* function is available for K, minimizing such function leads to appropriate certificates of boundedness or unboundedness (note that the minimum is finite only if K is bounded). That idea was used first by de Ghellink and Vial in [9] for linear programming and more recently by Nesterov, Todd and Ye on [48] for nonlinear programming problems. Moreover, if K is given explicitly by a set of linear inequalities, one can identify an element of the recession cone by solving a linear programming problem.

We emphasize that none of these approaches extends to the membership oracle framework. In fact, negative results do exist for approximating the diameter of a convex set, which is a closely related problem. Lovász and Simonovits [38] show that no polynomial time algorithm (deterministic or probabilistic) can approximate the diameter of a convex set within a factor of $\sqrt{n}$ in polynomial time under the membership oracle framework. Thus, it is notable that, as we show, testing if a convex set is unbounded is solvable in polynomial time.

An outline of the Chapter is as follows. Section 5.1 illustrates the geometric intuition underlying many results. Then we establish many properties relating the density functions (5.1) and the convex set K in Section 5.2. The algorithm to test boundedness is presented in Section 5.3 and its analysis is presented in the following sections. Finally, Appendix F contains the details on how to implement the hit-and-run geometric random walk efficiently for the density functions used in the algorithm.

5.0.4 Preliminaries, definitions, and notation

Recall that a real-valued function $\|\cdot\| : \mathbb{R}^n \to \mathbb{R}_+$ is said to be a norm if:

(i) $\|x\| = 0$ only if $x = 0$;

(ii) $\|tx\| = |t|\|x\|$ for any $t \in \mathbb{R}$;

(iii) $\|x + y\| \leq \|x\| + \|y\|$ for any $x, y \in \mathbb{R}^n$.

For a given norm, we can define a unit ball

$$B_{\|\cdot\|}(x, r) = B(x, r) = \{y \in \mathbb{R}^n : \|y - x\| \leq r\}, \tag{5.4}$$

and a unit sphere $S^{n-1} = \{y \in \mathbb{R}^n : \|y\| = 1\}$.

The Euclidean inner product is denoted by $\langle \cdot, \cdot \rangle$ while $\|\cdot\|_2 = \sqrt{\langle \cdot, \cdot \rangle}$ denotes the Euclidean norm induced by it. For $x \in \mathbb{R}^n$ and $r > 0$, let $B_2(x, r)$ denote the Euclidean ball centered at x with radius r, i.e., $B_2(x, r) = \{y \in \mathbb{R}^n : \|x - y\|_2 \leq r\}$. The unit Euclidean sphere of $\mathbb{R}^n$ is denoted by S_2^{n-1}, i.e., $S_2^{n-1} = \{y \in \mathbb{R}^n : \|y\| = 1\}$.

The dual norm of $\|\cdot\|$ induced by $\langle \cdot, \cdot \rangle$ is defined by

$$\|s\|_* = \max\{\langle s, x \rangle : x \in B(0, 1)\}, \tag{5.5}$$

for which we can also define an unit ball $B_*(s, r) = \{w \in \mathbb{R}^n : \|w - s\|_* \leq r\}$ and an unit sphere S_*^{n-1}. Note that we have that $|\langle s, x \rangle| \leq \|s\|_* \|x\|$. The dual norm completely defines the original norm, since we have

$$\|x\| = \max\{\langle s, x \rangle : s \in B_*(0, 1)\}. \tag{5.6}$$

That is, the dual norm of the dual norm is the original norm. Recall that the dual norm of the Euclidean norm is also the Euclidean norm, which is said to be self-dual.

A set K is convex if $x, y \in S$ implies $\alpha x + (1 - \alpha)y \in S$ for every $\alpha \in [0, 1]$. C is a cone if $x \in C$ implies $\alpha x \in C$ for every $\alpha \geq 0$. If C is a convex cone, the *width* of

C is given by

$$\tau_C = \max_{x,\tau}\{\tau : B(x,\tau) \subset C,\ \|x\| = 1\}, \tag{5.7}$$

the largest ball contained in C centered at a unit vector (in the appropriate norm). L is a subspace of $\mathbb{R}^n$ if $x, y \in L$ implies that $\alpha x + \beta y \in L$ for every $\alpha,\ \beta \in \mathbb{R}$.

For a set S, the operations **conv**(S), **cone**(S), **ext** (S), **int** (S), **cl** (S), **diam**(S), and **Vol**(S) denote, respectively, the convex hull, conic hull, extreme points, interior, closure, diameter, and volume of S (see [53] for complete definitions). Also, for $x \in \mathbb{R}^n$, let **dist**$(x, S) = \inf\{\|x - y\| : y \in S\}$ denote the distance of x to S. For a scalar u, set $(u)_+ = \max\{0, u\}$, and for a matrix M denote by $\lambda_{\max}(M)$ (respectively $\lambda_{\min}(M)$) its maximum (respectively minimum) eigenvalue.

A *membership oracle* for a set S is any algorithm, that given any point $x \in \mathbb{R}^n$, correctly identifies if $x \in S$ or not. Let 1_S denote the indicator function of the set S, that is, $1_S(x) = 1$ if $x \in S$ and $1_S(x) = 0$ otherwise.

With respect to the complexity notation, $g(n)$ is said to be $O(f(n))$ if there exists a constant M such that $g(n) \leq Mf(n)$, while $g(n, m)$ is $O^*(f(n, m))$ if there exists constants M and k such that $g(n) \leq Mf(n, m)\ln^k n$.

5.0.5 Logconcave densities: concepts and notation

We define π_f as the probability measure associated with a probability density function f (i.e, $\pi_f(S) = \int_S f(x)dx$), $E_f[\cdot]$ as the expectation with respect to f, and z_f as the mean of a random variable whose probability density function is f. The following class of functions plays a central role in the sampling literature.

Definition 6 *A function $f : \mathbb{R}^n \to \mathbb{R}_+$ is logconcave if for any two points $x,\ y \in \mathbb{R}^n$ and any $\lambda \in (0, 1)$,*

$$f(\lambda x + (1 - \lambda)y) \geq f(x)^\lambda f(y)^{1-\lambda}.$$

Definition 6 implies that $\ln f$ is a concave function and, in particular, the support of f is a convex set. We say that a random variable is logconcave if its probability density function is a logconcave function. Gaussians, exponential and uniforms densities are a few classical examples of logconcave densities.

There are a variety of metrics available for probability densities. Here, we will make use of two of them: the *total variation* norm, defined as

$$\|f - g\|_{TV} = \frac{1}{2} \int_{\mathbf{R}^n} |f(x) - g(x)| dx,$$

and the L^2 *norm of* f *with respect to* g, defined as

$$\|f/g\| = E_f \left[\frac{f(x)}{g(x)} \right] = \int_{\mathbf{R}^n} \frac{f(x)}{g(x)} f(x) dx = \int_{\mathbf{R}^n} \left(\frac{f(x)}{g(x)} \right)^2 g(x) dx.$$

The following useful concept is associated with the covariance matrix induced by f.

Definition 7 *A density function* f *with expectation* z_f *is said to be* C*-isotropic if for every vector* v*,*

$$\frac{\|v\|_2^2}{C} \leq \int_{\mathbf{R}^n} \langle v, x - z_f \rangle^2 f(x) dx \leq \|v\|_2^2 C,$$

equivalently, any eigenvalue λ *of the covariance matrix of* f *satisfies:* $\frac{1}{C} \leq \lambda \leq C$.

A function f is said to be in isotropic position if it is 1-isotropic, that is, its covariance matrix is the identity. Thus, any density can be made isotropic by a linear transformation of the space.

5.1 On the geometry of unbounded convex sets

In this section we revisit a classical representation theorem for closed convex sets and we provide a new set inclusion characterization for such sets which will be key in our analysis.

Let $K \subset \mathbb{R}^n$ be a closed convex set. As a matter of convenience, assume K is full dimensional, as one can always work with the affine hull of K, but at considerable notational and expositional expense. For the sake of exposition, in this section we will assume that K contains no lines (in the upcoming sections we will drop such assumption).

As is standard in convex theory, the set of all directions of half-lines contained in K defines the *recession cone* of K denoted by C_K, i.e.,

$$C_K = \{d \in \mathbb{R}^n : K + d \subseteq K\}, \tag{5.8}$$

which is a closed convex cone ([53] Corollary 8.3.2). Moreover, it is well-known that K is unbounded only if $C_K \neq \{0\}$ ([53] Theorem 8.4).

Under this framework, K contains at least one extreme point ([53] Corollary 18.5.3). Thus, the "Minkowski-Hirsch-Hoffman-Goldman-Tucker theorem" for representing closed, line-free, convex sets applies to K.

Theorem 20 *Any closed line-free convex set $K \subset \mathbb{R}^n$ can be decomposed into the sum of two sets, the recession cone of K and the convex hull of extreme points. That is,*

$$K = C_K + \mathbf{conv}\Big(\mathbf{ext}\ (K)\Big). \tag{5.9}$$

In order to develop intuition on the relation between high-dimensional convex sets and volume, we need to understand how to draw pictures of what high-dimensional convex sets look like. The intuition for convex bodies (bounded convex sets with nonempty interior) was first suggested by Milman. The fact that the volume of parallel intersections of half-spaces with K decays exponentially fast after passing the median level must be taken into account. As suggested by Milman, small dimensional pictures of a high-dimensional convex body should have a hyperbolic form, see Figure 5-1.

However, our concern here is to extend such intuition to unbounded convex sets. In this context a similar (concentration) phenomenon will also be observed. Assuming that the recession cone has positive width, "most of the points" of the set are in the recession cone. (Note that one needs to be careful when quantifying "most of the points", since the volume is infinite.) Again small dimensional pictures of high-dimensional unbounded convex sets must have a hyperbolic form, see Figure 5-2.

In fact, even if the recession cone has zero width, "most of the points" of K will

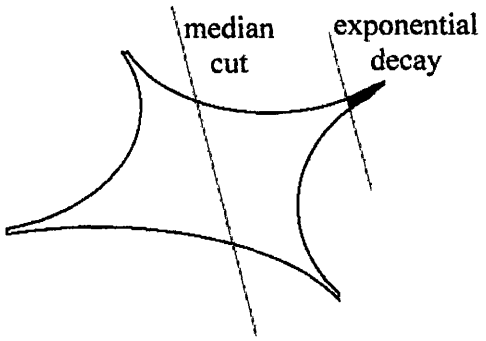

Figure 5-1: High-dimensional bounded convex sets and volume. Exponential decay as we move away the median cut.

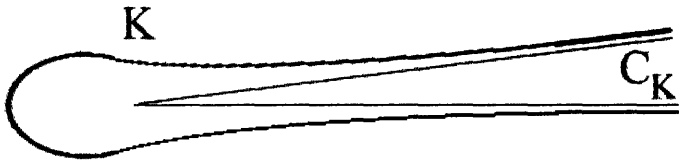

Figure 5-2: High-dimensional unbounded convex sets and volume. Most of the mass concentrates around the recession cone.

be contained in an ϵ-*enlargement* of the recession cone, where the latter is formally defined as:

$$C_K^\epsilon = \mathbf{cone}\left\{S^{n-1} \cap (C_K + B(0,\epsilon))\right\}. \tag{5.10}$$

The following properties of the ϵ-enlargement of the recession cone follow directly from the definition.

Lemma 21 *For $\epsilon > 0$ sufficiently small, we have that:*
(i) C_K^ϵ is a closed convex cone;
(ii) if $C_K = \{0\}$, then $C_K^\epsilon = \{0\}$;
(iii) if $C_K \neq \{0\}$, then $\tau_{C_K^\epsilon} \geq \tau_{C_K} + \epsilon$.

Now we are in position to obtain a set inclusion characterization of the aforementioned geometric phenomenon. That will motivate most of the analysis in the sections to come.

Theorem 21 *Let K be a convex set. Then, for any $\epsilon > 0$, there exists a finite number R_ϵ such that*

$$K \subset C_K^\epsilon + B(0, R_\epsilon). \tag{5.11}$$

Proof Without loss of generality, assume $0 \in K$. Suppose that there exists a sequence $\{x^j\}$ of elements of K such that $\mathbf{dist}(x^j, C_K^\epsilon) > j$. The normalized sequence has a convergent subsequence, $\{d^{k_j} = x^{k_j}/\|x^{k_j}\|\}$, to a point d. Since K is convex, closed and $0 \in K$, d^{k_j} and d are in K. In fact, $d \in C_K$. For any $\epsilon > 0$, there exists a number j_0 such that

$$\|d^{k_j} - d\| < \epsilon \quad \text{for all } j > j_0.$$

Therefore, we have $d^{k_j} \in C_K^\epsilon$ for $j > j_0$. In addition, $x^{k_j} \in C_K^\epsilon$, since C_K^ϵ is a cone, contradicting $\mathbf{dist}(x^j, C_K^\epsilon) > j$ for every j. $\square$

The ϵ-enlargement of the recession cone captures all points of K except for a bounded subset. Thus, assuming that K is unbounded, such geometric phenomenon

implies that "most points" of K will be contained in $K \cap C_K^\epsilon$. Moreover, if $\tau_{C_K} > 0$, most points of K will actually be contained in C_K itself.

Theorem 22 *Let C_K be a convex cone with strictly positive width τ_{C_K}. Then,*

$$\mathbf{Vol}_{n-1}\left(C_K^\epsilon \cap S^{n-1}\right) \leq \left(1 + \frac{2\epsilon}{\tau_{C_K}}\right)^n \mathbf{Vol}_{n-1}\left(C_K \cap S^{n-1}\right)$$

Proof By definition of τ_{C_K}, $B(z, \tau_{C_K}) \subset C_K$ for some $z \in S^{n-1}$. Since C_K and C_K^ϵ are cones,

$$\frac{\mathbf{Vol}_{n-1}\left(C_K^\epsilon \cap S^{n-1}\right)}{\mathbf{Vol}_{n-1}\left(C_K \cap S^{n-1}\right)} = \frac{\mathbf{Vol}\left(C_K^\epsilon \cap B(0,1)\right)}{\mathbf{Vol}\left(C_K \cap B(0,1)\right)} \leq \frac{\mathbf{Vol}\left(\left(C_K \cap B(0,1)\right) + B(0,\epsilon)\right)}{\mathbf{Vol}\left(C_K \cap B(0,1)\right)},$$

since $\left(C_K^\epsilon \cap B(0,1)\right) \subset C_K \cap B(0,1) + B(0,\epsilon)$. Next, we have $B(z/2, \tau_{C_K}/2) \subset C_K \cap B(0,1)$, so that

$$\left(C_K \cap B(0,1)\right) + B(0,\epsilon) \subset \left(C_K \cap B(0,1)\right) + \frac{2\epsilon}{\tau_{C_K}}\left(C_K \cap B(0,1) - \frac{z}{2}\right),$$

and the result follows. □

Corollary 8 *Let X be a random variable uniformly distributed on a set A, where the latter satisfies*

$$C_K \cap S^{n-1} \subseteq A \subseteq C_K^\epsilon \cap S^{n-1}.$$

Then,

$$P\left(X \in C_K\right) \geq \left(1 - \frac{2\epsilon}{\tau_{C_K}}\right)^n.$$

Taken together, Theorem 21 and Corollary 8 imply that for $\epsilon \leq \tau_{C_K}/8n$ a constant fraction of the points in C_K^ϵ are also in C_K. Thus, finding points in C_K is easily achieved by finding random points in $C_K^\epsilon \cap K$.

5.2 Probabilistic Properties

In this section we will establish several properties relating a random variable X_t, whose density function f_t is given by (5.1), and a variety of geometric properties of

the set K itself. For the sake of exposition, we will assume the following:

Assumption 4 *K is a closed convex set that contains the origin and has nonempty interior.*

Assumption 5 *There exists a positive number R such that $K \subseteq C_K + B(0, R)$.*

Assumption 4 is needed to ensure that K has positive n-dimensional volume (possibly infinite). Nonetheless, one can always work with the affine hull of K if one uses the appropriate k-dimensional volume. As expected, the origin could be replaced by any other point in K. All results could be restated for that particular point (if we translate the density f_t appropriately).

The focus on convex sets that satisfies Assumption 5 is not restrictive. In light of Section 5.1 we know that such inclusion can always be obtained if the recession cone is properly enlarged. In the case of unbounded convex sets whose recession cone has positive width, Corollary 8 shows how to relate the enlarged cone with the original recession cone. On the other hand, for unbounded convex sets whose recession cone has zero width, we know that the recession cone has zero measure for any f_t. Finally, we point out that Assumption 5 is satisfied by any bounded convex set with $R = \mathbf{diam}(K)$ and $C_K = \{0\}$, or any polyhedral convex set.

As anticipated, we will show that the probability of the event $\{X_t \in C_K\}$ will be large for "small" values of t. To do so, we exploit the spherical symmetry of f_t (i.e., $f_t(x) = f_t(y)$ if $\|x\| = \|y\|$) and the geometric phenomenon induced by the representation of Assumption 5. That symmetry will allow us to connect the volume of relevant sets with the probability of the event $\{X_t \in C_K\}$. Lemma 22 below properly quantifies this notion.

Lemma 22 *Suppose that K is an unbounded convex set, let $X_t \sim f_t$ as in (5.1), and let $\rho > 6R/\tau_{C_K}$. Then*

$$P(X_t \in C_K \mid \|X_t\| = \rho\ ) \geq \left(1 - \frac{6R}{\tau_{C_K} \cdot \rho}\right)^n.$$

Proof Note that restricted on $\|X_t\| = \rho$, the density f_t is constant. Thus, we have

$$P(X_t \in C_K \mid \|X_t\| = \rho\ ) = \frac{\mathbf{Vol}_{n-1}(C_K \cap (\rho S^{n-1}))}{\mathbf{Vol}_{n-1}(K \cap (\rho S^{n-1}))}.$$

For any a, b, define the sets $K_{[a,b]} := \{y \in K : a \leq \|y\| \leq b\}$ and $K_a := K_{[a,a]} = (aS^{n-1}) \cap K$. Set $\bar{\rho} = \rho/3 \geq 2R/\tau_{C_K}$, and consider the following set inclusions

$$I_{[\bar{\rho},\rho]} := \{y \in C_K : \bar{\rho} \leq \|y\| \leq \rho\} \subset K_{[\bar{\rho},\rho]} \subset O_{[\bar{\rho},\rho]} := \{y \in B(0,R) + C_K : \bar{\rho} \leq \|y\| \leq \rho\}.$$

We first show that $O_{[\bar{\rho},\rho]} \subset I_{[\bar{\rho},\rho]} + B(0, 2R)$. For $y \in O_{[\bar{\rho},\rho]}$, we have $y = v + w$, where $\|v\| \leq R$, $w \in C_K$, $\|v + w\| \in [\bar{\rho}, \rho]$, and hence $\|w\| \in [\bar{\rho} - R, \rho + R]$. Therefore $w \in I_{[\bar{\rho}-R,\rho+R]} \subset I_{[\bar{\rho},\rho]} + B(0, R)$, and the result follows since $v \in B(0, R)$.

Now, take $z \in C_K$, $\|z\| = 1$, such that $B(z, \tau_{C_K}) \subset C_K$. Thus,

$$B\left(\frac{(\bar{\rho}+\rho)}{2} z, \frac{(\bar{\rho}+\rho)}{2} \tau_{C_K}\right) \subset C_K.$$

Observe that $\rho \geq 3\bar{\rho}$ implies that $\bar{\rho} \leq \frac{(\bar{\rho}+\rho)}{2} - \frac{(\bar{\rho}+\rho)}{4}$ and $\frac{(\bar{\rho}+\rho)}{2} + \frac{(\bar{\rho}+\rho)}{4} \leq \rho$, so we have

$$B\left(\frac{(\bar{\rho}+\rho)}{2} z, \frac{(\bar{\rho}+\rho)}{4} \tau_{C_K}\right) = B\left(\frac{(\bar{\rho}+\rho)}{2} z, 2\kappa R\right) \subset I_{[\bar{\rho},\rho]}, \tag{5.12}$$

where $\kappa = \frac{(\bar{\rho}+\rho)}{8} \frac{\tau_{C_K}}{R}$. Thus, $B(w, 2R) \subset \frac{1}{\kappa} I_{[\bar{\rho},\rho]}$ for $w = \frac{\bar{\rho}+\rho}{2\kappa} z$.

$$\begin{aligned} \frac{\mathbf{Vol}(I_{[\bar{\rho},\rho]})}{\mathbf{Vol}(K_{[\bar{\rho},\rho]})} &\geq \frac{\mathbf{Vol}(I_{[\bar{\rho},\rho]})}{\mathbf{Vol}(O_{[\bar{\rho},\rho]})} \geq \frac{\mathbf{Vol}(I_{[\bar{\rho},\rho]})}{\mathbf{Vol}(I_{[\bar{\rho},\rho]} + B(0, 2R))} \\ &\geq \frac{\mathbf{Vol}(I_{[\bar{\rho},\rho]})}{\mathbf{Vol}(I_{[\bar{\rho},\rho]} + \frac{1}{\kappa} I_{[\bar{\rho},\rho]})} = \frac{1}{\left(1 + \frac{1}{\kappa}\right)^n} \geq \left(1 - \frac{1}{\kappa}\right)^n. \end{aligned}$$

We will complete the proof in three steps. For $s \in [\bar{\rho}, \rho]$, consider the sets I_s and O_s.

First note that

$$\frac{\mathbf{Vol}_{n-1}(I_s)}{\mathbf{Vol}_{n-1}(O_s)} = \frac{\mathbf{Vol}_{n-1}(C_K \cap S^{n-1})}{\mathbf{Vol}_{n-1}(C_K + B(0, R/s) \cap S^{n-1})} = \frac{\mathbf{Vol}_{n-1}(C_K \cap S^{n-1})}{\mathbf{Vol}_{n-1}(C_K + B(0, R/s) \cap S^{n-1})}$$

is a nondecreasing function of s. Second, observe that

$$\frac{\mathbf{Vol}(I_{[\bar{\rho},\rho]})}{\mathbf{Vol}(O_{[\bar{\rho},\rho]})} = \frac{\int_{\bar{\rho}}^{\rho} \mathbf{Vol}_{n-1}(I_s)ds}{\int_{\bar{\rho}}^{\rho} \mathbf{Vol}_{n-1}(O_s)ds}.$$

Next, we will make use of the following remark.

Remark 14 *For any $a < b$ and any two positive functions g and h such that*

$$\frac{g(s)}{h(s)} \text{ is nondecreasing for } s \in [a,b], \text{ we have } \frac{g(a)}{h(a)} \leq \frac{\int_a^b g(s)}{\int_a^b h(s)}ds \leq \frac{g(b)}{h(b)}.$$

Third, applying Remark 14 with $g(s) = \mathbf{Vol}_{n-1}(I_s)$, $h(s) = \mathbf{Vol}_{n-1}(O_s)$, $a = \bar{\rho}$ and $b = \rho$ to obtain

$$P(X_t \in C_K \mid \|X_t\| = \rho\ ) = \frac{\mathbf{Vol}_{n-1}(I_\rho)}{\mathbf{Vol}_{n-1}(K_\rho)} \geq \frac{\mathbf{Vol}_{n-1}(I_\rho)}{\mathbf{Vol}_{n-1}(O_\rho)} \geq \frac{\mathbf{Vol}(I_{[\bar{\rho},\rho]})}{\mathbf{Vol}(O_{[\bar{\rho},\rho]})} \geq \left(1 - \frac{1}{\kappa}\right)^n.$$

Finally, since $\bar{\rho} = \rho/3$ we have $\kappa = \frac{\tau_{C_K}\rho}{6R}$. □

Proof (OF REMARK 14) Simply note that $\frac{g(s)}{h(s)} \leq \frac{g(b)}{h(b)}$ for $s \in [a,b]$. Then

$$\int_a^b g(s)ds = \int_a^b h(s)\frac{g(s)}{h(s)}ds \leq \left(\frac{g(b)}{h(b)}\right)\int_a^b h(s)ds,$$

and a similar argument yields the lower bound. □

Since the bound obtained in Lemma 22 is monotone in ρ, we have the following useful corollary.

Corollary 9 *For unbounded K, we have*

$$P(X_t \in C_K \mid \|X_t\| \geq \rho) \geq \left(1 - \frac{6R}{\tau_{C_K} \cdot \rho}\right)_+^n.$$

The previous results used conditioning on the event that the random points have large norm. It is natural to bound the probability of these conditional events as well.

Lemma 23 *Suppose that K is unbounded, and that f_t is given by (5.1). For a random variable X_t distributed according to f_t we have that*

$$P\left(\|X_t\| \geq \rho\right) \geq 1 - 4et\rho.$$

Proof We can assume that $t\rho < 1/4e$, since the bound is trivial otherwise. We use the notation introduced in the proof of Lemma 22, $K_{[a,b]} = \{y \in K : a \leq \|y\| \leq b\}$.

Assuming that K is unbounded, there exists $z \in C_K$, $\|z\| = 1$. Then,

$$K_{[0,\rho]} + (\kappa + 2\rho)z \subset K_{[\kappa,\kappa+3\rho]} \quad \text{for any} \quad \kappa \in \mathbb{R}_+, \tag{5.13}$$

since for each $x \in K_{[0,\rho]}$, we have

$$\begin{aligned} \|x + (\kappa + 2\rho)z\| &\leq \|x\| + (\kappa + 2\rho)\|z\| \leq \kappa + 3\rho, \\ \|x + (\kappa + 2\rho)z\| &\geq (\kappa + 2\rho)\|z\| - \|x\| \geq \kappa + \rho, \end{aligned}$$

and $x + (\kappa + 2\rho)z \in K$ (the latter follows from $z \in C_K$).

So, for any odd integer $m \geq 3$, $K_{[\rho,m\rho]} = \bigcup_{i=1}^{(m-1)/2} K_{[(2i-1)\rho,(2i+1)\rho]}$, where the union is disjoint (except on a set of measure zero), and by (5.13) we have

$$\mathbf{Vol}(K_{[\rho,m\rho]}) = \sum_{i=1}^{(m-1)/2} \mathbf{Vol}(K_{[(2i-1)\rho,(2i+1)\rho]}) \geq \frac{(m-1)}{2}\mathbf{Vol}(K_{[0,\rho]}).$$

Now, define m such that $t = \frac{1}{m\rho}$. By assumption we have $t\rho < 1/4e$, which implies

$m \geq 4e$ (again, we assume that m in an odd integer for convenience). Thus,

$$\begin{aligned}
P(\|X_t\| \geq \rho) &= \frac{\int_{K_{[\rho,\infty]}} e^{-t\|x\|}dx}{\int_{K_{[0,\rho]}} e^{-t\|x\|}dx + \int_{K_{[\rho,\infty]}} e^{-t\|x\|}dx} \\
&\geq \frac{\int_{K_{[\rho,m\rho]}} e^{-t\|x\|}dx}{\int_{K_{[0,\rho]}} e^{-t\|x\|}dx + \int_{K_{[\rho,m\rho]}} e^{-t\|x\|}dx} \\
&\geq \frac{\mathbf{Vol}(K_{[\rho,m\rho]})e^{-mt\rho}}{\mathbf{Vol}(K_{[0,\rho]}) + \mathbf{Vol}(K_{[\rho,m\rho]})e^{-mt\rho}} \\
&\geq \frac{(m-1)\mathbf{Vol}(K_{[\rho,m\rho]})e^{-mt\rho}}{2\mathbf{Vol}(K_{[\rho,m\rho]}) + (m-1)\mathbf{Vol}(K_{[\rho,m\rho]})e^{-mt\rho}} \\
&\geq \frac{(m-1)e^{-mt\rho}}{2 + (m-1)e^{-mt\rho}}.
\end{aligned}$$

Using the definition of m, we have

$$\begin{aligned}
P(\|X_t\| \geq \rho) &\geq \frac{(\frac{1}{t\rho} - 1)e^{-1}}{2 + (\frac{1}{t\rho} - 1)e^{-1}} = \frac{(\frac{1}{t\rho} - 1)}{2e + (\frac{1}{t\rho} - 1)} \\
&\geq 1 - \frac{2e}{\frac{1}{t\rho} - 1} = 1 - \frac{2et\rho}{1 - t\rho} \geq 1 - 4et\rho.
\end{aligned}$$

□

Lemma 22 quantifies the geometric notion mentioned earlier (motivated on Figure 5-2), that most points in K outside a compact set are in C_K if its width is positive. On the other hand, Lemma 23 shows that the norm of X_t is likely to be greater than $0.01/t$. Taken together, they lead to a lower bound on the probability of the event $\{X_t \in C_K\}$.

Theorem 23 *Suppose that K is unbounded, and f_t is defined as in (5.1). Let*

$$t \leq \frac{\delta^2 \tau_{C_K}}{96enR},$$

and let X_t be a random variable distributed according to f_t. Then

$$P(X_t \in C_K) \geq 1 - \delta.$$

Proof All that is needed is to combine Lemmas 22 and 23. Clearly,

$$P(X_t \in C_K) \geq P(\|X_t\| \geq \rho)\, P(X_t \in C_K \mid \|X_t\| \geq \rho).$$

By Lemmas 22 and 23, we have

$$P(\|X_t\| \geq \rho) \geq 1 - 4et\rho \quad \text{and} \quad P(X_t \in C_K \mid \|X_t\| \geq \rho) \geq \left(1 - \frac{6R}{\tau_{C_K} \cdot \rho}\right)^n.$$

It is sufficient to ensure that

$$1 - 4et\rho > 1 - \frac{\delta}{2} \qquad \text{and} \qquad \left(1 - \frac{6R}{\tau_{C_K} \cdot \rho}\right)^n > 1 - \frac{\delta}{2}$$

for some $\rho \geq 6R/\tau_{C_K}$. Noting that $(1-x)^{1/n} \leq 1 - \frac{x}{n}$ for all $x \in [0,1]$, the second relation also holds since if

$$1 - \frac{6R}{\tau_{C_K} \cdot \rho} > 1 - \frac{\delta}{2n}.$$

It suffices to choose $\rho = \frac{12nR}{\delta \tau_{C_K}}$ for the second relation to hold, and the first relation holds since

$$4et\rho \leq \frac{4\delta^2 \tau_{C_K} e 12nR}{96enR\delta\tau_{C_K}} = \frac{\delta}{2}.$$

□

Theorem 23 characterizes the behavior of random variables $X_t \sim f_t$ for values of t that are "relatively small" with respect to its support K. It is natural to ask what kind of behavior one should expect for values of t that are "relatively large" with respect to the support K. An answer to this question is given in the next lemma.

Lemma 24 *Assume that there exists $\bar{v} \in K$ such that $\|\bar{v}\| = R$. If $t \geq \sqrt{n}/R$, then*

$$\max_{s \in S_*^{n-1}} P\left(\left|\langle s, X_t \rangle\right| > \frac{1}{\sqrt{n}(4et)}\right) > \frac{1}{3}.$$

Proof Since $R = \|\bar{v}\| = \max\{\langle s, \bar{v}\rangle : s \in B_*(0,1)\}$, we can always find $v \in S_*^{n-1}$ such that $\langle v, \bar{v}\rangle = R$.

By obtaining a lower bound for v we automatically obtain a lower bound for the maximum over the dual sphere. So,

$$\max_{s \in S_*^{n-1}} P\left(|\langle s, X_t\rangle| > \frac{1}{\sqrt{n}(4et)}\right) \geq P\left(|\langle v, X_t\rangle| > \frac{1}{\sqrt{n}(4et)}\right).$$

It will be convenient to define the following sets

$$A = \left\{x \in K : |\langle v, x\rangle| < \frac{1}{\sqrt{n}(4et)}\right\} \quad \text{and} \quad B = \left\{(1-\alpha)x + \alpha\bar{v} \ : \ x \in A\right\},$$

where $\alpha = \dfrac{1}{2etR\sqrt{n}} < \dfrac{1}{2en}$.

Note that for any $y \in B$, $y \in K$ since x and $\bar{v} \in K$ and $\alpha \in (0,1)$. Moreover,

$$|\langle v, y\rangle| \geq \langle v, y\rangle = (1-\alpha)\langle v, x\rangle + \alpha\langle v, \bar{v}\rangle > -\frac{1}{4et\sqrt{n}} + \frac{1}{2etR\sqrt{n}}R = \frac{1}{4et\sqrt{n}}.$$

Thus, $y \notin A$.

Since $\|(1-\alpha)x + \alpha\bar{v}\| \leq (1-\alpha)\|x\| + \alpha\|\bar{v}\| \leq \|x\| + \alpha\|\bar{v}\|$, we have

$$\begin{aligned}
P(X_t \in K \setminus A) \ &\geq \ P(X_t \in B) = \int_B \frac{e^{-t\|y\|}dy}{\int_K e^{-t\|z\|}dz} \\
&= \ \frac{1}{\int_K e^{-t\|z\|}dz}\int_A e^{-t\|(1-\alpha)x+\alpha\bar{v}\|}(1-\alpha)^n dx \\
&\geq \ \frac{e^{-t\alpha\|\bar{v}\|}(1-\alpha)^n}{\int_K e^{-t\|z\|}dz}\int_A e^{-t\|x\|}dx = e^{-t\alpha\|\bar{v}\|}(1-\alpha)^n P(X_t \in A) \\
&\geq \ \tfrac{1}{2}P(X_t \in A) = \tfrac{1}{2}\left(1 - P(X_t \in K \setminus A)\right)
\end{aligned}$$

and we have $P(X_t \in K \setminus A) \geq 1/3$.

Note that we have used that under our definitions

$$e^{-t\alpha\|\bar{v}\|}(1-\alpha)^n > 1/2$$

for any $n \geq 1$. Indeed this is the case since

$$\begin{aligned} e^{-t\alpha\|\bar{v}\|}(1-\alpha)^n &= e^{-1/(2e\sqrt{n})}\left(1-\frac{1}{2etR\sqrt{n}}\right)^n \geq e^{-1/2e}\left(1-\frac{1}{2en}\right)^n \\ &\geq e^{-1/2e}\cdot e^{-1/(2e-1)} \geq 0.664 \geq 1/2. \end{aligned}$$

□

This result allows us to construct several useful corollaries.

Corollary 10 *Assume that there exists $\bar{v} \in K$ such that $\|\bar{v}\| = R$. If $t \geq \sqrt{n}/R$, then*

$$\max_{s \in S_*^{n-1}} E\left[\langle s, X_t\rangle^2\right] \geq \frac{1}{n(7et)^2}.$$

Proof Let $v \in S_*^{n-1}$ as in Lemma 24. Thus,

$$E\left[\langle v, X_t\rangle^2\right] \geq \frac{1}{n(4et)^2} P\left(\left|\langle v, X_t\rangle\right| > \frac{1}{\sqrt{n}(4et)}\right) \geq \frac{1}{n(4et)^2}\frac{1}{3} \geq \frac{1}{n(7et)^2}.$$

□

Corollary 11 *Suppose that K is unbounded. For every $t > 0$, we must have*

$$\max_{v \in S_*^{n-1}} E\left[\langle v, X_t\rangle^2\right] \geq \frac{1}{n(7et)^2}.$$

Proof For every $t > 0$, let $R = \frac{\sqrt{n}}{t}$. Since K is unbounded, there exists $\bar{v} \in K$ with $\|\bar{v}\| = R$ and $t \geq \frac{\sqrt{n}}{R}$. Thus, Corollary 10 can be applied and the result follows. □

As a byproduct, we will be able to provide a bound on R in the case of bounded K.

Corollary 12 *Suppose that for some $t > 0$ we have that*

$$\max_{s \in S_*^{n-1}} E\left[\langle s, X_t \rangle^2\right] < \frac{1}{n(7et)^2}.$$

Then K is bounded and $K \subset B(0, R)$ for $R = \frac{\sqrt{n}}{t}$.

Proof Suppose there exists $\bar{v} \in K$, with $\|\bar{v}\| \geq \frac{\sqrt{n}}{t}$. Let $R := \|\bar{v}\| \geq \sqrt{n}/t$. Then $t \geq \sqrt{n}R$ and $\|\bar{v}\| = R$, so by Lemma 24 we have a contradiction. □

Corollary 13 *If $\|\cdot\| = \|\cdot\|_2$ and K is unbounded, for every $t > 0$*

$$\lambda_{\max}\left(E\left[X_t X_t'\right]\right) \geq \frac{1}{n(7et)^2}.$$

Proof If one is using the Euclidean norm, we have $\|\cdot\|_2 = \|\cdot\| = \|\cdot\|_*$. Moreover, we have that

$$\lambda_{\max}\left(E\left[X_t X_t'\right]\right) = \max_{s \in S^{n-1}} \langle s, E\left[X_t X_t'\right] s \rangle = \max_{s \in S^{n-1}} E\left[\langle s, X_t \rangle^2\right],$$

and the result follows from Corollary 10. □

A similar result can be established for small values at t. Intuitively we will recover results known for the uniform density as we let t goes to zero.

Lemma 25 *Assume that there exists $\bar{v} \in K$ such that $\|\bar{v}\| = R$. If $t < n/R$,*

$$\max_{s \in S_*^{n-1}} P\left(\left|\langle s, X_t \rangle\right| > \frac{R}{4en}\right) > \frac{1}{3}.$$

Proof The proof is similar to the proof of Lemma 24 if one defines

$$A = \left\{x \in K : \left|\langle v, x \rangle\right| < \frac{R}{4en}\right\}.$$

□

Corollary 14 *Assume that K has diameter D and $t < n/D$. Then,*

$$E\left[\|X_t\|\right] \geq \frac{D}{12en}.$$

Proof If K has diameter D, there exist two points $\bar{v}, \bar{w} \in K$ such that $\|\bar{v} - \bar{w}\| = D$. We can assume $\bar{w}$ to be the origin, which implies that $\|\bar{v}\| = D$. The result follows from Lemma 25. □

We close this section with a simple observation with respect to the *entropy* of a density function

$$Ent(f) = -\int_{\mathbf{R}^n} f(x) \ln f(x) dx.$$

Corollary 15 *If f_t is a norm induced density function, we have that*

$$Ent(f_t) = tE[\|X_t\|]$$

and Lemmas 24 and 25 can be used to bound the entropy.

Proof The result follows by noting that $-f(x) \ln f(x) = t\|x\|e^{-t\|x\|}$. □

5.3 Testing the Boundedness of a Convex Set: a Density Homotopy

The algorithm we propose is a homotopy procedure to simulate a random variable which has desirable properties with respect to K. Motivated by the geometry of unbounded convex sets, the uniform density over K would be an interesting candidate. Unfortunately, as opposed to most frameworks in the literature, a random variable which is uniformly distributed over K will not be *proper* in the unbounded case and cannot be used. Instead, we will work with a parameterized family of densities, $\mathcal{F} = \{f_t : t \in (0, t_0]\}$, such that f_t is a proper density for every t. In addition, for any fixed compact subset of K the parameterized density uniformly converges to the uniform density over that compact set as $t \to 0$. As mentioned earlier, the algorithm

must provide us with a certificate of boundedness or unboundedness. Any nonzero element of the recession cone of K is a valid certificate of unboundedness. We will assume that a membership oracle for the recession cone of K itself is available.

On the other hand, the certificate of boundedness is more thought-provoking. If K is described by a set of linear inequalities, $K = \{x \in \mathbb{R}^n : Ax \leq b\}$, K will be bounded if and only if positive combinations of the rows of A span $\mathbb{R}^n$. More generally, if K is represented by a separation oracle, a valid certificate of boundedness would be a set of normal vectors associated with hyperplanes returned by the separation oracle whose positive combinations span $\mathbb{R}^n$. Note that a membership oracle provides much less information and we cannot sensibly extend the previous concept to our framework. Instead, our certificate of boundedness will be given by the eigenvalues of the second moment matrix associated with the random variables induced by the family $\mathcal{F}$. In contrast with the previous certificates, it will be a "probabilistic certificate of boundedness" since the true second moment matrix is unknown and must be estimated via a probabilistic method.

5.3.1 Assumptions and condition measures

In addition to Assumption 5, we make the following assumptions on the set K:

Assumption 6 *K is a closed convex set given by a membership oracle.*

Assumption 7 *There exists $r > 0$ such that $B(0, r) \subseteq K$.*

Assumption 8 *A membership oracle is available for C_K, the recession cone of K.*

We point out that the closedness of K could be relaxed with minor adjustments on the implementation of the random walk. Assumptions 5 and 8 specify how K is represented.

Assumption 7 is stronger than Assumption 4. It requires that we are given a point in the interior of K, which is assumed to be the origin without loss of generality. That is standard in the membership oracle framework, since the problem of finding a point in a convex set given only by a membership oracle is hard in general. Finally, we

emphasize that only a lower bound on r is required to implement that algorithm. Section 5.4.3 gives a simple procedure to obtain an approximation of r within a factor of $\sqrt{n}$.

In our analysis, besides the dimension of K, there are three geometric quantities that naturally arise: r, R, and τ_{C_K}. Not surprisingly, the dependence of the computational complexity of our algorithm on these geometric quantities differs if K is bounded or unbounded (recall that the case of $\tau_{C_K} = 0$ is fundamentally different if K is bounded or unbounded). Nonetheless, in either case the dependence on these quantities will be only logarithmic. An instance of the problem is said to be *ill-conditioned* if $\tau_{C_K} = 0$ and K is unbounded, otherwise the instance is said to be *well-conditioned.*

5.3.2 The algorithm

In order to define the algorithm, let f_t be defined as (5.1) with $\|\cdot\| = \|\cdot\|_2$ (the Euclidean norm). Let $\psi \in (0,1)$ and let *hit-and-run* be a geometric random walk which will simulate the next random variable (see Section 5.5 for details). This yields the following "exact" method to test the boundedness of K:

Density Homotopy Algorithm (Exact):

Input: r such that $B(0,r) \subset K$, define $t_0 = t_{\text{initial}}$, $\psi \in (0,1)$, and set $k \leftarrow 0$.

Step 1. (Initialization) Simulate $X_{t_0} \sim f_{t_0}(x)$.

Step 2. (Variance and Mean) Compute the covariance and mean: z_k and V_k.

Step 3. (Testing Unboundedness) If $X_{t_k} \in C_K \setminus \{0\}$, **stop**.

Step 4. (Testing Boundedness) If $\lambda_{\max}\left(V_k + z_k z_k^T\right) < \frac{1}{n(7et_k)^2}$, **stop**.

Step 5. (Update Density) Update the parameter: $t_{k+1} = (1-\psi)\cdot t_k$.

Step 6. (Random Point) Draw $X_{t_{k+1}} \sim$ *hit-and-run*$(\, f_{t_{k+1}},\ X_{t_k},\ V_k)$

Step 7. (Loop) Set $k \leftarrow k+1$, goto **Step 2**.

This (exact) method requires r, the exact simulation of $X_{t_{k+1}}$, and the exact computation of the covariance matrix V_{k+1} and mean z_{k+1} of the random variable $X_{t_{k+1}}$. In order to obtain a implementable method, we can use only approximations of these objects.

Detailed bounds on the computational complexity of the hit-and-run procedure,

on the estimation of $\hat{z}_{k+1}$ and $\hat{V}_{k+1}$, and $\hat{r}$ are provided in Sections 5.5, 5.6, and 5.4.3. Moreover, the use of an approximated covariance matrix and approximated mean must be taken into account on the test of boundedness (Step 4), which is done in Theorem 24. That is done by showing that the maximum eigenvalue of the second moment matrix, $\Omega_k = V_k + z_k z_k^T$, will be estimated up to a factor of three.

Each loop of the algorithm (Steps 2-7) is called an iteration throughout this Chapter. Thus, the work per iteration consists of (i) performing the hit-and-run random walk, (ii) computing an approximation for the covariance matrix, (iii) testing if the current point belongs to the recession cone, and (iv) computing the largest eigenvalue of a definite positive matrix. Although a highly accurate approximation for the covariance matrix is not needed, the probabilistic method used to estimate such matrix requires at least $O^*(n)$ samples. Such estimation will dictate the computational complexity per iteration.

Letting t_{final} denote the final value of the parameter t when the algorithm stops, the total number of iterations of the algorithm will be

$$\left\lceil \frac{1}{\psi} \cdot \ln\left(\frac{t_{\text{initial}}}{t_{\text{final}}}\right) \right\rceil.$$

Next we state the main theorem of the section.

Theorem 24 *Let K be a convex set satisfying Assumptions 5, 7, 8, and 6 and consider the homotopy algorithm using a family of densities $\mathcal{F} = \{f_t(x) \sim 1_K(x) \cdot e^{-t\|x\|_2} : t \in (0, t_0]\}$. Then:*

(i) If K is unbounded, the algorithm will compute a valid certificate of unboundedness in at most

$$O\left(\sqrt{n} \ln\left(\frac{n}{\delta} \frac{1}{\tau_{C_K}} \frac{R}{r}\right)\right) \quad \textit{iterations with probability } 1-\delta,$$

where each iteration makes at most $O^*\left(n^4 \ln\left(\frac{n}{\delta} \ln\left(\frac{1}{\delta} \frac{1}{\tau_{C_K}} \frac{R}{r}\right)\right)\right)$ calls to the membership oracle.

(ii) If K is bounded, the algorithm will compute a valid certificate of boundedness in

at most

$$O\left(\sqrt{n}\ln\left(n\frac{R}{r}\right)\right) \quad \textit{iterations with probability } 1-\delta,$$

where each iteration makes at most $O^*\left(n^4 \ln\left(\frac{n}{\delta}\ln\left(\frac{R}{r}\right)\right)\right)$ *calls to the membership oracle.*

The proof of Theorem 24, which is provided in Section 5.9, is built upon the analysis of the next sections.

5.4 Analysis of the homothopy algorithm

5.4.1 Stopping criteria: unbounded case

An appropriate certificate of unboundedness for a convex set is to exhibit a non-zero element of the recession cone of K. Assumption 8 allows us to correctly verify if a point is an element of C_K. For example, in the case of linear conic systems (5.3), any membership oracle for C itself can be used to construct a membership oracle for K and C_K.

We point out that if the algorithm terminates indicating that K is unbounded, a nonzero element of the recession cone was found (a certificate of unboundedness). Thus, the algorithm always terminates correctly in this case. The following corollary of Theorem 23 ensures that we can find such certificate, which provides a desirable stopping criteria in the case of K being unbounded.

Corollary 16 *Suppose* K *is unbounded. After*

$$T = \frac{1}{\psi}\ln\left(t_0\frac{96enR}{\delta^2\tau_{C_K}}\right)$$

iterations of the exact algorithm, we have $P\left(X_{t_T} \in C_K\right) \geq 1-\delta$.

Proof We start the algorithm with $t = t_0$, and after T iterations, we obtain $t_T = \frac{\delta^2\tau_{C_K}}{96enR}$. The result follows by applying Theorem 23 to X_{t_T}. □

5.4.2 Stopping criteria: bounded case

In contrast with the unbounded case, we lack a straightforward certificate for the case of K being bounded. In addition, unbounded sets whose recession cone has zero width should not be wrongly classified as bounded. That is, our analysis should cover those ill-conditioned cases as well.

On the search for an appropriate certificate, the mean of the random variable X_t appears as a natural candidate. Assuming that the set K is unbounded and line-free, its norm should increase as the parameter t decreases. On the other hand, if K is bounded, the mean will eventually be bounded no matter how much t decreases. Unfortunately, that analysis breaks down for sets that contain lines. For example, if K is symmetric the mean of X_t is zero for every $t > 0$, whether K is bounded or not.

In order to overcome that we consider the second moment matrix Ω_t of the random variable X_t. The matrix Ω_t will be large, in the positive definite sense, if either the covariance matrix or the norm of the mean of X_t is large. Again, if K is unbounded the maximum eigenvalue of Ω_t increases as the parameter t decreases. Otherwise, K being bounded, the maximum eigenvalue will eventually be bounded. That provide us with a nice criterion which is robust to instances where K contains lines and/or $\tau_{C_K} = 0$. We emphasize that the second order information is readily available, since we are required to compute the covariance matrix and the mean of X_t to implement the hit-and-run random walk for f_t (see Section 5.5 for details on the sampling theory and the reasons why we need to compute the covariance matrix to keep f_t in near isotropic position). Next corollary provides the desirable stopping criteria.

Corollary 17 *Suppose K is bounded. The exact algorithm will detect boundedness for all $t < \frac{1}{7eR\sqrt{n}}$, and will bound R by $\frac{\sqrt{n}}{t}$. Moreover, this will happen in at most $T = \frac{1}{\psi} \ln (t_0\ 7eR\sqrt{n})$ iterations.*

Proof If $t < \frac{1}{7eR\sqrt{n}}$, then $\lambda_{\max} (E[X_t X_t']) < \frac{1}{(7et)^2 n}$ from Corollary 13. From Corollary 12, this means that K is bounded and $R < \frac{\sqrt{n}}{t}$.

Moreover, starting with $t = t_0$, after $T = \frac{1}{\psi} \ln (t_0\ 7eR\sqrt{n})$ iterations we obtain $t_T \leq \frac{1}{7eR\sqrt{n}}$. □

To generate a certificate of boundedness, the maximum eigenvalue of the second moment matrix must be estimated (Section 5.6 covers the necessary theory for that). Since it will be estimated via a probabilistic method, there is a probability of failure on each iteration which must be properly controlled (see Section 5.8 for details). Thus, in contrast to the unbounded case, if the algorithm terminates indicating that K is bounded, there is a probability of failure associated with that decision which can be made smaller at the expense of computational time.

5.4.3 Initialization of the algorithm: unknown r

This section clarifies how to start the algorithm. As is usual under the membership oracle framework, it is assumed that we know a point in the interior of K, which is taken to be the origin for convenience (Assumption 7). In some applications of interest such points are readily available, for example $0 \in \mathbf{int}\ K$ in the conic system (5.3).

The implementation of Step 1 will be done by a simple accept-reject method, see [8]. Note that it is simple to draw a random variable X_{t_0} whose density is proportional to $f_{t_0}(x) \sim e^{-t_0\|x\|_2}$ on $\mathbb{R}^n$ instead of only on K (pick a point uniformly on S_2^{n-1} and then scale using a $\Gamma(n, t_0)$ random variable[1]). If it is the case that $X_{t_0} \in K$, we accept that point; otherwise, we draw another point according to f_{t_0} and repeat the process until a point in K is obtained.

Now, we need a constructive approach to bound r from below. Fortunately, a simple procedure is available for estimating r up to a factor of $\sqrt{n}$. That will be satisfactory since the final dependence on r is only logarithmic. Consider the following quantity:

$$\hat{r} = \min_{i=1,\ldots,n} \max\{\ t\ :\ \ t\mathbf{e}_i \in K,\ \ -t\mathbf{e}_i \in K\}, \tag{5.14}$$

where $\mathbf{e}_i \in \mathbb{R}^n$ denotes the i^{th} unit vector of $\mathbb{R}^n$. It is clear that $\hat{r}$ can be approximated in $O\left(n|\ln\frac{1}{r}|\right)$ operations (via a simple binary search) and will not increase the order of the computational complexity of the algorithm. The next lemma provides a guarantee

[1] A $\Gamma(\alpha, \beta)$ random variable is characterized by the density $f(x) = (\beta^\alpha/\Gamma(\alpha))x^{\alpha-1}e^{-\beta x}$ for $x \geq 0$, and zero otherwise.

of the quality of the approximation provided by $\hat{r}$.

Lemma 26 *Let r be the radius of the largest ball centered at the origin contained in K and let $\hat{r}$ be as defined in (5.14). Then $\hat{r} \geq r \geq \hat{r}/\sqrt{n}$.*

Proof Clearly, $\hat{r} \geq r$. Note that $\hat{r}\mathbf{e}_i \in K$ for every i. Thus, the convex hull of these points contains a ball of radius $\hat{r}/\sqrt{n}$ which is contained in K. Therefore, $\hat{r} \geq r \geq \hat{r}/\sqrt{n}$. □

We also need to ensure that the probability of the event $\{X_{t_0} \in K\}$ is reasonably large. The next lemma achieves this by a suitable choice of the initial parameter t_0 based on the radius r of the largest ball centered at the origin contained in K.

Lemma 27 *Assume that the ball centered at the origin with radius r is contained in K. Let X_{t_0} be a random variable whose density is proportional to $e^{-t_0\|x\|_2}$ for any $x \in \mathbb{R}^n$. Then if $t_0 \geq 2\frac{(n-1)}{r}$, we have*

$$P(X_{t_0} \in K) \geq 1 - e^{n-(t_0 r/2)}.$$

Proof For any $y \notin K$, $f_{t_0}(y) \leq f_{t_0}(0)e^{-t_0 r} = f_{t_0}(0)e^{-8n}$, since $\|y\|_2 > r$. Using Lemma 5.16 of [41] (since $t_0 r/(n-1) \geq 2$), we have

$$\begin{aligned} P(X_{t_0} \notin K) &\leq P(X_{t_0} \notin B(0,r)) = P\big(f_{t_0}(X_{t_0}) \leq f_{t_0}(0)e^{-t_0 r}\big) \leq \left(e^{1-\frac{t_0 r}{n-1}}\frac{t_0 r}{n-1}\right)^{n-1} \\ &\leq e^{n-1}\left(e^{-\frac{t_0 r}{n-1}}\frac{t_0 r}{n-1}\right)^{n-1} \leq e^{n-(t_0 r/2)} \end{aligned}$$

since $e^{-c}c \leq e^{-c/2}$ for every $c \geq 0$.

□

This allows us to efficiently implement the accept-reject method in Step 1 of the algorithm.

Corollary 18 *After computing $\hat{r}$ as in (5.14), it suffices to initialize the algorithm with*

$$t_0 = \frac{8n^{3/2}}{\hat{r}}$$

to obtain that $P(X_{t_0} \in K) \geq 1 - e^{-3n}$.

Proof Recall that Lemma 26 implies that$\hat{r} \geq r/\sqrt{n}$. By Lemma 27, $P(X_{t_0} \in K) \geq 1 - e^{n-4n^{3/2}\frac{r}{\hat{r}}} \geq 1 - e^{-3n}$. □

5.5 Sampling f_t via a geometric random walk

The ability to sample according to any density of the family $\mathcal{F}$ is the driving force of our algorithm. Although a variety of densities can be perfectly simulated with negligible computational effort, that is no longer the case if we restrict the support of the density to be an arbitrary convex set K given by a membership oracle. In fact, even to generate a random point distributed uniformly over a convex set is an interesting problem with many remarkable applications (linear programming, computing the volume, etc., see [31],[43],[40]).

Important tools to generate random points proportional to a density function restricted to a high dimensional convex set K are the so-called geometric random walks. Starting with a point in K, on each step the random walk moves to a point according to some distribution that depends only (i) on the current point, and (ii) on the desired density f to be simulated. Thus, the sequence of points of the random walk is a Markov Chain whose state space is K. Moreover, there are simple choices of transition kernels (which is the continuous state space analog for the transition matrix for a finite state Markov Chain) that make f the unique stationary distribution of this Markov Chain (for example, the celebrated Metropolis Filter), which automatically ensures several asymptotic results for arbitrary Markov Chains [8]. Going one step further, we are interested in the convergence rate to the stationary distribution, which is a much more challenging question (which could be arbitrarily slow in general). So we can bound the necessary number of steps required by the random walk to generate a random point whose density is approximately f.

By choosing $\mathcal{F}$ to be a family of logconcave densities, we will be able to invoke several results from a recent literature which proves the efficiency of one particular random walk called hit-and-run, see [37, 41, 42]. Roughly speaking, these results

show that if (i) a relatively good approximation of the covariance matrix is available, and (ii) the density of the current point is close to the desired density, then only $O^*(n^3)$ steps of the random walk are necessary to generate a good approximation for the desired random point.

In our context, recall that the distribution of interest f_{t_k} is changing at each iteration. The current approximation of the covariance matrix $\hat{V}_k$ will be used as the approximation to the true covariance matrix of the next iteration V_{k+1}, which in turn will be estimated by $\hat{V}_{k+1}$. In a similar way, the current point X_{t_k}, distributed approximately according to f_{t_k}, will be used as the starting point for the random walk to approximate a random variable distributed according to $f_{t_{k+1}}$. The parameter ψ, which dictates the factor by which t is decreased at each iteration, will be the largest value such that these approximations are valid from a theoretical perspective.

5.5.1 A geometric random walk: hit-and-run

There are several possible choices of geometric random walks. We refer to [59] for a recent survey. Here we use the so-called hit-and-run random walk. The implementation of this random walk requires the following as input: a density function f (known up to a constant), a starting point X^0, a covariance matrix V, and a number of steps m.

Subroutine: *hit-and-run*$(f,\ X^0,\ V,\ m)$

Step 0 Set $k \leftarrow 0$.

Step 1 Pick a random vector $d \sim N(0, V)$.

Step 2 Define the line $\ell(X^k, d) = \{X^k + td :\ t \in \mathbb{R}\}$.

Step 3 Move to a point X^{k+1} chosen according to f restricted to $\ell(X^k, d)$.

Step 4 Set $k \leftarrow k + 1$. If $k \leq m$, goto **Step 1**.

Step 5 Report X^m.

Although hit-and-run can be implemented for arbitrary densities, we will restrict ourselves to the case of logconcave densities. In such case, the implementation of

hit-and-run can be done efficiently, and we refer to the Appendix for a complete description for the case in which f_t is defined as in (5.1) with the Euclidean norm.

5.5.2 Sampling complexity

Here we state without proof results of the literature of sampling random points according to logconcave densities. We start with a complexity result for the mixing time of the hit-and-run random walk.

Theorem 25 *([40] Theorem 1.1) Let f be a logconcave density such that (i) the level set of probability $1/8$ contains a ball of radius s, (ii) $E_f[\|x - z_f\|^2] \leq S^2$, and (iii) the L_2 norm of the starting distribution σ with respect to the stationary distribution π_f is at most M. Let σ^m be the distribution of the current point after m steps of the hit-and-run applied to f with $V = I$. Then, after*

$$m = O\left(\frac{n^2 S^2}{s^2} \ln^5\left(\frac{nM}{\varepsilon}\right)\right) \quad \textit{steps,}$$

the total variation distance of σ^m and π_f is at most ε.

Theorem 25 bounds the rate of convergence of the geometric random walk not only on the dimension but also on the L^2 norm of the starting density with respect to the stationary density f_t, and on how "well-rounded" is f_t via the ratio S/s. The notion of "well-rounded" is deeply connected with the concept of near isotropic position. Next lemma quantifies that connection.

Lemma 28 *([41] Lemma 5.13) Let f be a density in C-isotropic position. Then $E_f[\|X - z_f\|^2] \leq Cn$, and any upper level set U of f contains a ball of radius $\pi_f(U)/e\sqrt{C}$.*

We point out that any (full dimensional) density can be put in near-isotropic position by a suitable linear transformation. By using an approximation of the covariance matrix $\hat{V}$ to implement the hit-and-run random walk such that all eigenvalues of $\hat{V}^{-1}V$ are between $1/C$ and C, f_t is in C-isotropic position. Thus, the ratio S/s

can be bounded by $8eC\sqrt{n}$ (in this case, note that only $m = O^*(n^3)$ steps of the random walk will be necessary to generate one random sample). Next section summarizes how to compute an approximation $\hat{V}_k$ which puts f_{t_k} in 2-isotropic position. Moreover, it will be shown that all densities simulated by the algorithm will be at most C-isotropic for a constant C (independent of the dimension) as we decrease the homotopy parameter t.

Remark 15 *In our analysis we will be assuming independence among different samples for simplicity (recall that they are separated by $m = O^*(n^3)$ steps of the random walk). Although this is not the case, independence can be approximated at the cost of an additional constant factor in the number m of steps of the random walk. Here we have chosen exposition over formalism since no additional insight is gained if we work out all the details.*

5.6 Estimating the covariance matrix, the mean, and the second moment matrix

In this section, we recall estimation results for the mean covariance matrix of a log-concave random variable. Moreover, we show that these estimates can be used to approximate the second moment matrix with a desired relative accuracy. Herein it will be assumed that independent identically distributed samples $\{X^i\}$ are available. We emphasize that these estimations depend only on the samples and not on the isotropic position of the density function. As stated before, the isotropic position plays an important role to bound the number of steps of the chain required to obtain each sample.

First we recall a result for estimating the mean and covariance matrix.

Lemma 29 *([41] Lemma 5.17) Let z and V denote respectively the mean and covariance matrix of a logconcave random variable. For any $\xi > 0$ and $\delta \in (0,1)$,*

using

$$N > \frac{4}{\xi^2} n \ln^2 \frac{1}{\delta} \quad \textit{samples and} \quad \hat{X} = \frac{1}{N} \sum_{i=1}^{N} X^i,$$

we have that

$$P\left(\left\| V^{-1/2} \left(\hat{X} - z \right) \right\|_2 > \xi \right) \leq \delta.$$

Lemma 30 *([31] Corollary A.2) Let V denote the covariance matrix of a logconcave random variable. Using $N > O(\ln^3 \frac{1}{\delta} n \ln^2 n)$ samples, where $\delta < 1/n$,*

$$\hat{V} = \frac{1}{N} \sum_{i=1}^{N} (X^i - \hat{X})(X^i - \hat{X})^T$$

we have that all eigenvalues of the matrix $\hat{V}^{-1}V$ are between $1/2$ and 2 with probability at least $1 - \delta$.

These results yield a useful estimation procedure for the second moment matrix. In particular the maximum eigenvalue of Ω will be estimated up to a (known) constant factor.

Lemma 31 *Let Ω denote the second moment matrix associated with a logconcave random variable. Then for $\xi < 1/4$ and using $N > O(\ln^3 \frac{1}{\delta} n \ln^2 n + \frac{1}{\xi^2} n \ln^2 \frac{1}{\delta})$, for $\delta < 1/n$, the matrix*

$$\hat{\Omega} = \hat{V} + \hat{X}\hat{X}^T$$

is such that all eigenvalues of $\hat{\Omega}^{-1}\Omega$ are within $(1/2 - 2\xi)$ and $(2 + 2\xi + \xi^2)$ with probability at least $1 - \delta$.

Proof In this proof let $\| \cdot \| = \| \cdot \|_2$. Lemma 29 yields that $\|V^{-1/2}(\hat{X} - z)\| \leq \xi$ with probability greater than $1 - \delta/2$. In this event, there exists $d \in \mathbb{R}^n$ with $\|d\| \leq 1$ satisfying

$$\hat{X} = z + \xi V^{1/2} d.$$

For any $w \in \mathbb{R}^n$, we have that

$$\begin{aligned}\left\langle w, \hat{\Omega} w \right\rangle &= \left\langle w, \hat{V} w \right\rangle + \left\langle \hat{X}, w \right\rangle^2 = \left\langle w, \hat{V} w \right\rangle + \left\langle z + \xi V^{1/2} d, w \right\rangle^2 \\ &= \left\langle w, \hat{V} w \right\rangle + \langle z, w \rangle^2 + 2\xi \langle z, w \rangle \left\langle d, V^{1/2} w \right\rangle + \xi^2 \left\langle d, V^{1/2} w \right\rangle^2 \qquad (5.15) \\ &\geq \tfrac{1}{2} \langle w, V w \rangle + \tfrac{1}{2} \langle z, w \rangle^2 - 2\xi \langle w, \Omega w \rangle = (\tfrac{1}{2} - 2\xi) \langle w, \Omega w \rangle .\end{aligned}$$

□

We will use the following corollary in our algorithm.

Corollary 19 *Using $N = O(n \ln^2 n \ln^3 \frac{1}{\delta})$ samples to estimate the second moment matrix $\hat{\Omega}$, with probability at least $1 - \delta$ we have that all eigenvalues of $\hat{\Omega}^{-1}\Omega$ are between $1/3$ and 3.*

Proof Set $\xi = 1/24$ and apply Lemma 31. □

Thus, with $O^*(n)$ samples per iteration one can properly estimate the mean, covariance matrix, and second moment matrix to conduct the algorithm.

5.7 Updating the parameter t: warm-start

Since the parameter ψ controls how fast the sequence $\{t_k\}$ decreases, its value is critical to the computational complexity of the algorithm. Although we are tempted to decrease t as fast as possible, we still need to use the current density as a "warm-start" to approximate the next density. That is, the L^2 norm of f_{t_k} with respect to $f_{t_{k+1}}$ needs to be controlled. Moreover, the covariance matrix associated with f_{t_k} should also be close to the covariance matrix of the next iterate $f_{t_{k+1}}$. The trade-off among these quantities will decide how fast one can decrease the parameter t. Kalai and Vempala show how to relate the L^2 norm of two logconcave densities and their covariance matrices. To the reader's convenience we state their results here.

Lemma 32 *([31] Lemma 3.8) Let f and g be logconcave densities over K with mean $z_f = E_f[x]$ and $z_g = E_g[x]$, respectively. Then, for any $c \in \mathbb{R}^n$*

$$E_f\left[\langle c, x - z_f \rangle^2\right] \leq 16 \, \|f/g\| \, E_g\left[\langle c, x - z_g \rangle^2\right]$$

This result implies that it is enough to control the L^2 norm in order to guarantee that the covariance matrices of the densities are relatively close. A simple extension of Lemma 3.11 of [31] yields the desire result.

Corollary 20 *Consider the densities f_{t_k} and $f_{t_{k+1}}$ such that $t_{k+1} = (1-\psi)t_k = (1-\frac{1}{\sqrt{n}})t_k$. For $n \geq 5$,*

$$\|f_{t_k}/f_{t_{k+1}}\| \leq e^{n/(n-1)} \quad \text{and} \quad \|f_{t_{k+1}}/f_{t_k}\| \leq e^{n/(n-2\sqrt{n})}.$$

Next lemma combines the previous results to ensure that all the densities used in the algorithm will be in 2^8-isotropic position.

Lemma 33 *Assuming that $n \geq 16$, and using $N = O^*(\kappa^3 n)$ samples in any iteration, the distribution encountered by the sampling algorithm in the next iteration is 2^9-isotropic with probability $1 - 2^{-\kappa}$.*

Proof By using Lemma 30, we have that f_{t_k} is 2-isotropic after we estimated its covariance matrix in Step 2 of the algorithm with probability $1 - 2^{-\kappa}$. Using Lemma 32 and Corollary 20, for any $v \in S^{n-1}$,

$$E_{f_{t_{k+1}}}\left[\left\langle v, x - z_{f_{t_{k+1}}}\right\rangle^2\right] \leq 16\, e^{n/(n-2\sqrt{n})} E_{f_{t_k}}\left[\left\langle v, x - z_{f_{t_k}}\right\rangle^2\right] \leq 32\, e^{n/(n-2\sqrt{n})} \leq 2^8,$$

since $n \geq 16$. □

Lemma 33 ensures that $O^*(n)$ samples suffice to estimate the current covariance matrix accurately enough to be used as an approximation to the covariance matrix associated with the next iteration.

5.8 Controlling the overall probability of failure

Before we proceed to the proof of Theorem 24, we prove a technical lemma which allows us to properly control the overall probability of failure of the algorithm. First

note that if the algorithm finds an element of C_K it will always stop correctly. Thus, the algorithm can fail only if it (wrongly) declares that K is a bounded set or if it does not terminate. The latter concerns with the stopping criteria of the algorithm and it was already analyzed in Sections 5.4.1 and 5.4.2. The first issue can occur only if the estimated second moment matrix $\hat{\Omega}_k$ differs significantly from the true matrix Ω_{t_k}. In turn their difference is controlled by the number of samples used to estimate $\hat{\Omega}_{t_k}$, which depends on the probability of failure used for the current iteration. Recall that we do not know the total number of iterations T a priori (since we do not have any estimates for R or τ_{C_K}), and so we cannot simply set the probability of failure to δ/T for each iteration. Instead of using a constant probability of failure for all iterations, we slightly increase the number of samples used to estimate $\hat{\Omega}_{t_k}$ at each iteration to ensure that the probability of failure will not accumulate indefinitely.

Lemma 34 *Let T be the total number of iterations of the algorithm and $1-\delta_f$ the desired probability of success. Initializing the algorithm with $\delta = \delta_f/4$ and setting the probability of failure of the i^{th} iteration to $\delta_i = \frac{1}{i^2}\delta_f/4$, we obtain:*
(i) the probability of success is at least $1-\delta_f/2$;
(ii) The smallest δ used by the algorithm satisfies $\delta \geq \frac{\delta_f}{4}\frac{1}{T^2}$.

Proof (i) By setting a probability of failure of the i^{th} iteration of the algorithm to $(1/i^2)\delta_f/4$, we have that the total probability of failure is bounded by

$$\sum_{i=1}^{T}\delta_i = \sum_{i=1}^{T}\frac{\delta_f}{4i^2} \leq \frac{\delta_f}{4}\sum_{i=1}^{\infty}\frac{1}{i^2} \leq \delta_f/2$$

since $\sum \frac{1}{i^2} = \pi^2/6 < 2$.

(*ii*) follows since the algorithm terminates after T iterations.

□

5.9 Proof of Theorem 24

Let $\psi = \frac{1}{\sqrt{n}}$ and let $1 - \delta_f$ be the desired probability of success. The algorithm terminates in $T = \sqrt{n} \ln \left(\frac{t_{initial}}{t_{final}} \right)$ iterations. Next we need to: (i) properly bound $t_{initial}$, (ii) ensure that the algorithm terminates, i.e., properly bound t_{final} (which also bounds T), and (iii) control the probability that it wrongly terminates.

(i) Lemma 27 yields $t_{initial} = t_0 \geq 8n/\hat{r} \geq 8n^{3/2}/r$.

(ii) We will use Corollaries 17 and 16 with $\delta = \delta_f/2$. Thus, the algorithm will correctly terminate with probability at least $1 - \delta_f/2$ after

$$\bar{T} = \sqrt{n} \ln \left((8n^{3/2}/r)(96enR)/((\delta_f/2)^2 \tau_{C_K}) \right)$$

iterations if K is unbounded, and after $\bar{T} = \sqrt{n} \ln \left((8n^{3/2}/r) 21eR\sqrt{n}) \right)$ iteration if K is bounded (note that we use 21 instead of 7 to take into account a factor of 3 in the approximation of the maximum eigenvalue of V_k).

(iii) Finally, we ensure that the probability of (wrongly) terminating before reaching iteration $\bar{T}$ is at most $\delta_f/2$. This is achieved by slightly reducing the probability of failure as described in Lemma 34.

Now we analyze the impact of reducing the probability of failure on the computational complexity of each iteration. Lemma 34(ii) ensures that it is sufficient to use

$$\delta = \frac{\delta_f/2}{4\bar{T}^2} = \frac{\delta_f}{32n \ln^2 (t_{initial}/t_{final})}.$$

This will dictate the number of samples needed to estimate the covariance matrix. In Lemma 33 we will need at most

$$\kappa = \frac{1}{\ln 2} \ln \frac{32\bar{T}^2}{\delta_f}$$

on any given iteration (given that we terminate correctly). That leads to a total of at most

$$O\left(n \ln^3 \left(\frac{\bar{T}}{\delta_f} \right) \right)$$

samples per iteration of the algorithm.

Since each sample is computed by using $O^*(n^3)$ steps, and we have $\bar{T}$ iterations, the overall complexity is

$$O^*\left(\bar{T}\, n^3\, n \ln^3\left(\frac{\bar{T}}{\delta_f}\right)\right) = O^*\left(n^{4.5} \ln\left(\frac{\bar{T}}{\delta_f}\right)\right).$$

The results of Theorem 24 follow by using the appropriate $\bar{T}$ on each case as defined in (ii).

5.10 Conclusions

In this work we study probability densities f_t induced by an arbitrary norm with a convex support K as defined by (5.1). Our goal is to relate geometric properties of K with analytical properties of f_t.

Using these properties, we also develop a test to decide whether a convex set, given only by a membership oracle, is bounded or unbounded. The computational complexity of the probabilistic method proposed is polynomial in the dimension of the set and only logarithmic on other condition measures (such as the width of the recession cone and the diameter of the set, respectively, for the unbounded and bounded cases).

Exploiting the geometry of unbounded convex sets was key in developing a probabilistic method to implement that test. A geometric phenomenon guarantees that most points of an unbounded convex set will be in its recession cone if that cone has positive width. In such cases, random points with large norms are likely to belong to the recession cone.

In contrast with probabilistic methods over convex sets in the literature, we need to explicitly deal with unbounded sets, and additional effort is needed to ensure that all the densities are well defined. Moreover, if K is unbounded our analysis shows

that an element of C_K^ϵ can be computed in

$$O^*\left(\sqrt{n}\ln\left(\frac{R}{r(\tau_{C_K}+\epsilon)}\right)\right) \quad \text{iterations.}$$

That is particularly relevant for the ill-conditioned case of $\tau_{C_K} = 0$. Although we cannot find an element of C_K (which has zero volume), the algorithm will generate a direction d, $\|d\| = 1$, such that $\mathbf{dist}(d, C_K) < \epsilon$ in at most

$$O^*\left(\sqrt{n}\ln\left(\frac{R}{r}\frac{1}{\epsilon}\right)\right)$$

iterations of the homotopy algorithm.

Chapter 6

Conclusions

In this thesis we explored connections between geometry, probability, and optimization under convexity. We demonstrated how geometric quantities can be used to create efficient algorithms for a variety of interesting problems. In particular, probabilistic tools, such as sampling random points, lend themselves very naturally in the analysis and implementation of these algorithms. Specifically,

Symmetry Function of a Convex Set: we established many new properties of this function first proposed by Minkowski [44]. In particular, the log-concavity of the symmetry function, probabilistic bounds on the distribution of the symmetry function over the set, etc. Moreover, we characterize its maximum points in general and provide efficient methods to compute/maximize this function in the case of polyhedral sets.

Projective Pre-conditioners for IPM: based upon new results that we obtained for the symmetry function, a family of pre-conditioner for interior point methods applied to homogeneous linear conic systems is proposed. In particular, we showed that there exists a pre-conditioner that makes the IPM strongly polynomial. We provided a constructive procedure to obtain such pre-conditioner with high probability.

Efficiency of a Re-scaled Perceptron Algorithm: we developed the analysis of the re-scaled perceptron algorithm, originally proposed in [11] for a system

of inequalities, for a linear conic system. In the process, we also developed a new complexity bound for the problem E.1.

Norm-induced Densities and Boundedness of a Convex Set: we investigated properties of random variables whose densities are induced by an arbitrary norm. We connect the average behavior of these random variable with geometric properties of its own support. These results motivates a method to test if a convex set, given only by membership oracle, is bounded or not. Moreover, this method efficiently provides a certificate of boundedness (associated with the covariance matrix) or unboundedness (an element of the recession cone) with high probability.

In addition, we emphasize that in many cases geometric properties can be seen as condition measures of the problems (e.g., width of a cone, $\mathbf{sym}(0, H)$, $B(0, r) \subset K \subset B(0, R)$). In turn these quantities can bound the behavior of solutions and help establish robustness properties.

6.1 Future research

There are numerous research questions associated with the themes of this thesis.

The symmetry function could be used in robust optimization to relax the symmetry assumption which is usually made in this context [28]. Another research question is to extend the symmetry concept from convex sets to log-concave densities [60]. Finally, the efficient computation of the symmetry function for non-polyhedra sets is still open.

As indicated in Chapter 3, there are many open questions concerning the projective pre-conditioners. The counterpart of the projective pre-conditioner for optimization problems is currently under investigation. A simple conversion to a feasibility problem would not exploit the underlying structure of the problem. Moreover, the use of θ^* to bound the number of iterations from below is also of theoretical interest.

With respect to the re-scaled perceptron algorithm, it is interesting to analyze the

performance of the method in the alternate system. In the case in which the original system is infeasible, the perceptron algorithm is generating an infeasibility certificate. It would be interesting to investigate the corresponding infeasibility certificates generated by the re-scaled perceptron algorithm.

The ability to sample from log-concave densities has already motivated many interesting algorithms. Although the log-concave case is now well understood, little is known outside this context. The ability of sampling from different classes of functions is promising for further developments. A near log-concave case was investigated by Applegate and Kannan in [2]. Under statistical assumptions, the central limit theorem can be used to enforce near log-concavity and thus the efficiency of the sampling method in [5]. An interesting extension of the current literature would be the ability to sample from the boundary of convex sets, or more generally, from sets with non-negative Ricci curvature [61].

As shown in this thesis we believe that the combination of probabilistic and deterministic methods will lead to better algorithms. Perhaps the most important issue that has not been addressed here is to map the frontier of problems that are solvable efficiently by probabilistic methods.

Appendix A

Hahn-Banach Theorem

Theorem 26 (Hahn-Banach for real vector spaces) *Let X be a Banach space and p a convex functional on X, then there exists a linear functional λ on X such that*

$$p(x) \geq \lambda(x) \qquad \forall x \in X$$

Proof Take $0 \in X$, and define $\tilde{X} = \{0\}$, $\tilde{\lambda}(0) = p(0)$.
If $\exists z \in X$, $z \notin \tilde{X}$, extend $\tilde{\lambda}$ from $\tilde{X}$ to the subspace generated by $\tilde{X}$ and z,

$$\tilde{\lambda}(tz + \tilde{x}) = t\tilde{\lambda}(z) + \tilde{\lambda}(\tilde{x}).$$

Suppose $x_1, x_2 \in \tilde{X}$ and $\alpha > 0, \beta > 0$.

$$\begin{aligned}
\beta\lambda(x_1) + \alpha\lambda(x_2) &= \lambda(\beta x_1 + \alpha x_2) \\
&= (\alpha+\beta)\lambda\left(\tfrac{\beta}{\alpha+\beta}x_1 + \tfrac{\alpha}{\alpha+\beta}x_2\right) \\
&\leq (\alpha+\beta)p\left(\tfrac{\beta}{\alpha+\beta}x_1 + \tfrac{\alpha}{\alpha+\beta}x_2\right) \\
&= (\alpha+\beta)p\left(\left[\tfrac{\beta}{\alpha+\beta}\right](x_1 - \alpha z) + \left[\tfrac{\alpha}{\alpha+\beta}\right](x_2 + \beta z)\right) \\
&\leq \beta p(x_1 - \alpha z) + \alpha p(x_2 + \beta z)
\end{aligned}$$

Thus,

$$\beta\left[-p(x_1 - \alpha z) + \lambda(x_1)\right] \leq \alpha\left[p(x_2 + \beta z) - \lambda(x_2)\right]$$

$$\frac{1}{\alpha}\left[-p(x_1-\alpha z)+\lambda(x_1)\right] \leq \frac{1}{\beta}\left[p(x_2+\beta z)-\lambda(x_2)\right]$$

$$\sup_{x_1\in X,\alpha>0}\frac{1}{\alpha}\left[-p(x_1-\alpha z)+\lambda(x_1)\right] \leq a \leq \inf_{x_2\in X,\beta>0}\frac{1}{\beta}\left[p(x_2+\beta z)-\lambda(x_2)\right]$$

Define $\tilde{\lambda}(z)=a$, then assuming $t>0$,

$$\begin{aligned}\tilde{\lambda}(tz+\tilde{x}) &= t\tilde{\lambda}(z)+\tilde{\lambda}(\tilde{x})=ta+\tilde{\lambda}(tx)\\ &\leq t\left(\tfrac{1}{t}p(\tilde{x}+tz)-\tfrac{\lambda(\tilde{x})}{t}\right)+\lambda(\tilde{x})\\ &\leq p(tz+\tilde{x})\end{aligned}$$

the case of $t<0$ is similar. To extend for X, we will use the Zorn's lemma.
Let $\mathcal{E}$ be a collection of all extensions e of λ, $e(x)\leq p(x)$ for all $x\in X_e$. Where $e_1 \prec e_2$ if $X_{e_1}\subseteq X_{e_2}$ and $e_1(x)=e_2(x)$ in X_{e_1}. Thus, $\mathcal{E}$ is a partially ordered and $\mathcal{E}\neq\emptyset$.

If $\{e_s\}_{s\in J}$ is a totally ordered subset of $\mathcal{E}$, $\Lambda=\cup_{s\in J}X_{e_s}$ is a subspace (monotone union of subspaces) and

$$e:\Lambda\to\mathbb{R},\quad e(x)=e_s(x),\quad \text{if } x\in X_{e_s}.$$

e is well defined and $e_s\prec e$ for all $s\in J$. Then e is maximal for J. Thus, $\mathcal{E}$ must have a maximal element (Zorn's Lemma) $\tilde{\lambda}$.

So, $\tilde{\lambda}$ must be defined on X, otherwise we could extend it contradicting the fact that it is maximal.

□

Appendix B

Symmetry Function

B.1 sym(x, S) and sym(S) under Relaxed Assumptions

All of the results in this paper are based on Assumption A, namely that S is a closed, bounded, convex set with an interior. Herein we discuss the implications of relaxing this set of assumptions.

As mentioned earlier, the assumption that S has an interior is a matter of convenience, as we could instead work with the relative interior of S on the affine hull of S, at considerable notational and expository expense.

The assumption that S is closed is also a matter of convenience, as most of the statements contained in the body of the paper would still remain valid by replacing inf $\leftarrow$ min and sup $\leftarrow$ max and/or by working with the closure of S, etc.

Suppose that we relax the assumption that S is bounded. If S is unbounded then S has a non-empty recession cone. In the case when the recession cone of S is not a subspace, then $\mathbf{sym}(S) = 0$. However, the case when the recession is a subspace is a bit more interesting:

Lemma 35 *Suppose that $S = P + H$, where H is a subspace and P is a bounded convex set in $H^{\perp}$, and $x \in S$; then* $\mathbf{sym}(x, S)$ *is completely defined by P, i.e.,* $\mathbf{sym}(x, S) = \mathbf{sym}(w, P)$ *where $x = w + h$ and $(w, h) \in H^{\perp} \times H$.*

Proof Without loss of generality, we can assume that $x = 0$ since symmetry is invariant under translation. Trivially, $-\alpha S \subseteq S$ if and only if $-\alpha(P+H) \subseteq (P+H)$. Since P and H lie in orthogonal spaces, for each $x \in S$, there exist a unique $(w,h) \in P \times H$ such that $x = w + h$. Since $-\alpha H = H$, $-\alpha x \in S$ if and only if $-\alpha w \in P$. □

B.2 Standard Interior-Point Method for Linear Programming

Consider the following linear programming problem in "dual" form, where M is an $m \times k$ matrix:

$$\begin{aligned} P: \quad \text{VAL} := \max_{x,s} \quad & c^T x \\ \text{s.t.} \quad & Mx + s = f \\ & s \geq 0 \\ & x \in \mathbb{R}^k, \quad s \in \mathbb{R}^m \ . \end{aligned} \tag{B.1}$$

For $\beta \in (0,1)$, a β-approximate analytic center of the primal feasibility inequalities $Mx \leq f$ is a feasible solution x^a of P (together with its slack vector $s^a = f - Mx^a$) for which there exists dual multipliers z^a that satisfy:

$$\begin{aligned} & Mx^a + s^a = f, \ s^a > 0 \\ & M^T z^a = 0 \\ & \|S^a z^a - e\| \leq \beta \ , \end{aligned} \tag{B.2}$$

where S is the diagonal matrix whose diagonal entries correspond to the components of s. Following [54] or [57], one can prove the following result about the efficiency of a standard primal interior-point method for approximately solving P.

Theorem 27 *Suppose that $\beta = 1/8$ and that (x^a, s^a, z^a) is a given β-approximate analytic center of the feasibility inequalities of P, and that $c = M^T\lambda$ for some $\lambda \in \mathbb{R}^m$. Then (x^a, s^a, z^a) can be used to start a standard interior-point method that will compute a feasible solution of P with duality gap at most $\bar{g}$ in at most*

$$\left\lceil (2+4\sqrt{m}) \ln \left(\frac{10m\|S^a\lambda\|}{\bar{g}} \right)^+ \right\rceil$$

iterations of Newton method.

Now consider the following linear programming problem format:

$$\begin{array}{lll} P' : & \text{VAL} := \max_{x,\theta} & \theta \\ & \text{s.t.} & Mx + d\theta + s = f \\ & & s \geq 0 \\ & & x \in \mathbb{R}^k, \quad \theta \in \mathbb{R}, \quad s \in \mathbb{R}^m \ . \end{array} \tag{B.3}$$

Again following [54] or [57], one can prove the following result about the efficiency of a standard primal interior-point method for approximately solving P.

Theorem 28 *Suppose that $\beta = 1/8$ and that (x^a, s^a, z^a) is a given β-approximate analytic center of the feasibility inequalities of P, and that $d^T z^a > 0$. Then (x^a, s^a, z^a) can be used to start a standard interior-point method that will compute a feasible solution of P' with duality gap at most $\bar{g}$ in at most*

$$\left\lceil (2+4\sqrt{m}) \ln \left(\frac{1.25m}{\bar{g} \cdot (d^T z^a)} \right)^+ \right\rceil$$

iterations of Newton method.

Appendix C

Analysis of the Computational Complexity to compute the symmetry of a polyhedra

Here we present the details of the computational complexity analysis of the methods for computing the symmetry of a polyhedra represented by a finite number of inequalities.

As mentioned in Section 2.4.2 we propose three different methods: a potential reduction, a primal-dual, and a dual method.

Here "$\Lambda \geq 0$" is componentwise for all m^2 components of Λ. This means that $\mathbf{sym}(\bar{x}, S) \geq \alpha$ if and only if (2.36) has a feasible solution. Thus, we can completely characterize the α-the level sets, $\Gamma(\alpha)$. This also implies that $\mathbf{sym}(S)$ is the optimal objective value of the following optimization problem:

$$\begin{array}{ll} \min\limits_{x,\Lambda,\alpha} & \alpha \\ \text{s.t.} & \Lambda A = -\alpha A \\ & \Lambda(b - Ax) \leq b - Ax \\ & \Lambda \geq 0 \ , \end{array} \tag{C.1}$$

and any solution $(x^*, \Lambda^*, \alpha^*)$ of (2.37) satisfies $\mathbf{sym}(S) = \alpha^*$ and $x^* \in S_{opt}(S)$.

Remark 16 *Notice that (2.37) is not a linear program.*

C.1 Potential reduction

In this section we will use $\gamma = \frac{1}{\alpha} \in [1, \infty)$. Therefore, we will be interested in reducing γ on each iteration, and we denote the optimal value by $\gamma^* = 1/\mathbf{sym}(S)$. Consider the upper level sets defined as

$$\Gamma(\gamma) = \{(\Pi, x, t) \in \mathbb{R}^{m \times m}_+ \times S \times \mathbb{R}^m_+ : \Pi A = -A, \Pi(b - Ax) + t = \gamma(b - Ax)\}.$$

The following quantities, associated with the set $S \subset \mathbb{R}^n$, are needed for our analysis:

$$L^{LO} = \min_{i=1,\ldots,m} \min_{x \in \Gamma(n)} (b - Ax)_i, \tag{C.2}$$

$$L^{UP} = \max_{x \in \Gamma(n)} \|b - Ax\|_1 \leq m \max_{i=1,\ldots,m} \max_{x \in \Gamma(n)} (b - Ax)_i. \tag{C.3}$$

These quantities are well defined and strictly positive if S is bounded and has non empty interior. We note that it is possible to relate L^{LO} and L^{UP} with the number of bits needed to represent the instance. We start with a simple lemma.

Lemma 36 *If S is bounded, $\Gamma(\gamma)$ is bounded and has a nonempty (relative) interior for every $n \geq \gamma > \gamma^*$.*

Proof Since S is bounded, all components of x are also bounded. Note that $(b - Ax) > L^{LO}e > 0$ and $(b - Ax) \leq L^{UP}e$. Using the fact $t \geq 0$ and $\Pi \geq 0$, we can trivially bound

$$\pi_{ik} \leq \gamma \frac{L^{UP}}{L^{LO}} \quad \text{and} \quad t_i \leq L^{UP}.$$

It is easy to see that the relative interior is non-empty for every $\gamma > \gamma^*$. □

Next we define the potential functional that will be used in our analysis.

$$
\begin{aligned}
P(\gamma) = \min_{x,\Pi,t} \quad & -\sum_{i=1}^{m}\sum_{k=1}^{m}\log(\pi_{ik}) - \sum_{i=1}^{m}\log(t_i) \\
s.t. \quad & \\
& \Pi A = -A \\
& \Pi(b - Ax) + t = \gamma(b - Ax).
\end{aligned} \tag{C.4}
$$

We denote by $P(x, \Pi, t) = -\sum_{i=1}^{m}\sum_{k=1}^{m}\log(\pi_{ik}) - \sum_{i=1}^{m}\log(t_i)$.

Lemma 37 *If $\gamma^1 > \gamma^2 > \gamma^*$, then $P(\gamma^1) < P(\gamma^2)$.*

Proof Let x^i, Π^i, t^i such that $P(x^i, \Pi^i, t^i) = P(\gamma^i)$. Then $x^2 \in \Gamma(\gamma^1)$ since $(\Pi^2, x^2, t^2 - (\gamma^1 - \gamma^2)(b - Ax^2))$ is feasible for the (C.4) with γ^1. Moreover, $\gamma^1 > \gamma^2$ implies that $(\gamma^1 - \gamma^2)(b - Ax^2) > 0$. Thus, $P(\gamma^2) > P(\gamma^1)$ follows directly. □

Since $\max_x \mathbf{sym}(x, S) \geq 1/dim(S)$, we know that $\Gamma(2dim(S))$ has a non-empty interior and $2dim(S)$ is a valid starting value for γ.

Theorem 29 *Given an $\beta(< 1/2)$-approximation (x, Π, t), for the analytical center with γ. It is possible to decrease γ by*

$$
\Delta\gamma := \frac{\beta}{2\|S(b - Ax)\|_2}
$$

where (U, S) are dual variables such that

$$
\|\Pi \otimes U - ee^T, St - e\|_2 < \beta
$$

Proof In order to show that $\gamma - \Delta\gamma \geq \gamma^*$ we will construct a feasible solution for $\Gamma(\gamma - \Delta\gamma)$. We can obtain a feasible solution by defining

$$
\Pi(b - Ax) + t - \Delta\gamma(b - Ax) = (\gamma - \Delta\gamma)(b - Ax)
$$

$$
\begin{aligned}
\widetilde{\Pi} &\leftarrow \Pi \\
\widetilde{x} &\leftarrow x \\
\widetilde{t} &\leftarrow t - \Delta\gamma(b - Ax).
\end{aligned}
$$

Note that by construction we have

$$\begin{aligned} \|\widetilde{\Pi} \otimes U - ee^T, S\widetilde{t} - e\|_2 &= \|\Pi \otimes U - ee^T, S(t - \Delta\gamma(b - Ax)) - e\|_2 \\ &\leq \|\Pi \otimes U - ee^T, St - e\|_2 + \|0, \Delta\gamma S(b - Ax)\|_2 \\ &\leq \beta + \Delta\gamma\|S(b - Ax)\|_2 \\ &< 2\beta, \end{aligned} \tag{C.5}$$

where the last inequality holds if we use

$$\Delta\gamma = \frac{\beta}{2\|S(b - Ax)\|_2} \geq \frac{\beta}{2\|S(b - Ax\|_1} = \frac{\beta}{2\left(\sum_{i=1}^{m} s_i(b - Ax)_i\right)}.$$

Moreover, since $2\beta < 1$, we still have $t_i > 0$. □

The next theorem ensures that the potential function must increase by a fixed amount if we decrease the value of γ by $\Delta\gamma$.

Theorem 30 *Using the previous definitions of* Δ *and* γ*, we have*

$$P(\gamma - \Delta\gamma) \geq P(\gamma) + \frac{1}{72}$$

Proof First we write the increment of the potential function as

$$P(\gamma - \Delta\gamma) - P(\gamma) = P(\gamma - \Delta\gamma) - \underbrace{P(\widetilde{\Pi}, \widetilde{x}, \widetilde{t})}_{(i)} + \underbrace{P(\widetilde{\Pi}, \widetilde{x}, \widetilde{t}) - P(\Pi, x, t)}_{(ii)} + \underbrace{P(\Pi, x, t) - P(\gamma)}_{(iii)}.$$

Next we will provide bounds over each of the three pairs.

(iii) we have that $P(\Pi, x, t) - P(\gamma) \geq 0$ since (Π, x, t) is feasible for $\Gamma(\gamma)$.

(ii) since $(\widetilde{\Pi}, \widetilde{x}, \widetilde{t})$ is a 2β approximation for $\gamma - \Delta\gamma$, we obtain

$$P(\widetilde{\Pi}, \widetilde{x}, \widetilde{t}) - \frac{(2\beta)^2}{2(1 - 2\beta)} \leq P(\gamma - \Delta\gamma) \leq P(\widetilde{\Pi}, \widetilde{x}, \widetilde{t}).$$

Thus, we have

$$P(\gamma - \Delta\gamma) - P(\widetilde{\Pi}, \widetilde{x}, \widetilde{t}) \geq -\frac{(2\beta)^2}{2(1 - 2\beta)}.$$

For (i), we have that (for the reader convenience we include an explanation on the left hand side in parenthesis)

$$\begin{aligned}
P(\widetilde{\Pi},\widetilde{x},\widetilde{t}) - P(\Pi,x,t) &= -\sum_{i=1,k=1}^{m} \log \widetilde{\pi}_{ik} - \sum_{i=1}^{m} \log \widetilde{t}_i + \sum_{i=1}^{m}\sum_{k=1}^{m} \log \pi_{ik} + \sum_{i=1}^{m} \log t_i \\
&= -\sum_{i=1}^{m}\sum_{k=1}^{m} \log \frac{\widetilde{\pi}_{ik}}{\pi_{ik}} - \sum_{i=1}^{m} \log \frac{\widetilde{t}_i}{t_i} \\
(\widetilde{\pi}_{ik} = \pi_{ik}) &= -\sum_{i=1}^{m}\sum_{k=1}^{m} \log \frac{\pi_{ik}}{\pi_{ik}} - \sum_{i=1}^{m} \log \frac{t_i - \Delta\gamma(b-Ax)_i}{t_i} \\
(\log 1 = 0) &= -\sum_{i=1}^{m} \log \left(1 - \frac{\Delta\gamma(b-Ax)_i}{t_i}\right) \\
(-\log(1-x) \geq x) &\geq \sum_{i=1}^{m} \Delta\gamma \frac{(b-Ax)_i}{t_i} \\
(\Delta\gamma \geq \tfrac{\beta}{2(\sum_{i=1}^{m} s_i(b-Ax)_i)}) &\geq \sum_{i=1}^{m} \frac{\beta}{2\|S(b-Ax)\|_1} \frac{(b-Ax)_i}{t_i} \\
(|s_i t_i - 1| < \beta \Rightarrow t < \tfrac{1+\beta}{s_i}) &\geq \sum_{i=1}^{m} \frac{\beta}{2\|S(b-Ax)\|_1} \frac{s_i(b-Ax)_i}{(1+\beta)} \\
&= \frac{\beta}{2(1+\beta)} \sum_{i=1}^{m} \frac{s_i(b-Ax)_i}{\|S(b-Ax)\|_1} \\
&= \frac{\beta}{2(1+\beta)} \frac{\|S(b-Ax)\|_1}{\|S(b-Ax)\|_1} \\
&= \frac{\beta}{2(1+\beta)}
\end{aligned} \tag{C.6}$$

Finally, setting $\beta = 1/8$, we have that

$$\begin{aligned}
P(\gamma - \Delta\gamma) - P(\gamma) &\geq -\frac{4\beta^2}{2(1-2\beta)} + \frac{\beta}{2(1+\beta)} + 0 \\
&= -\frac{1/16}{2(1-1/4)} + \frac{1/8}{2(1+1/8)} \\
&= 1/18 - 1/24 = 1/72
\end{aligned} \tag{C.7}$$

□

The previous theorem leads to the following complexity result.

Theorem 31 *If* $\gamma \geq \gamma^* + \varepsilon$, *then*

$$P(\gamma) \leq (m^2 + m)\log\left(\frac{1}{\varepsilon}\right) + m^2 log\left(\frac{2mL^{UP}}{L^{LO}}\right) + m\log\left(\frac{2}{L^{LO}}\right) \tag{C.8}$$

Proof To bound the value of $P(\gamma)$ we will construct a feasible solution based on the analytical center (Π, x, t) of $\Gamma(1/2(\gamma + \gamma^*)) = \Gamma(\gamma - 2\Delta\gamma)$, where $\Delta\gamma > \varepsilon/2$.

$$\begin{aligned} &\Pi A = -A \\ &\Pi(b - Ax) + t = (\gamma - 2\Delta\gamma)(b - Ax) \end{aligned} \tag{C.9}$$

The second equation gives us

$$\Pi(b - Ax) + \Delta\gamma(b - Ax) + t + \Delta\gamma(b - Ax) = \gamma(b - Ax)$$

We will introduce a positive perturbation on Π such that its contribution is bounded by $\Delta\gamma(b - Ax)$. Consider

$$\rho^T = \frac{L^{LO}}{mL^{UP}} e^T(\Pi + I) \tag{C.10}$$

There are two desired properties in this vector:

$$\rho^T A = \frac{L^{LO}}{mL^{UP}} e^T(\Pi + I)A = \frac{L^{LO}}{mL^{UP}} e^T(\Pi A + A) = \frac{L^{LO}}{mL^{UP}} e^T(-A + A) = 0$$

$$\rho_i \geq \frac{L^{LO}}{mL^{UP}}$$

Consider the following perturbation matrix $e\rho^T$, its contribution in the system is

$$
\begin{aligned}
e\rho^T(b-Ax) &= \tfrac{L^{LO}}{mL^{UP}} ee^T(\Pi+I)(b-Ax) \\
&= \tfrac{L^{LO}}{mL^{UP}} e(e^T\Pi+e^T)(b-Ax) \\
&= \tfrac{L^{LO}}{mL^{UP}} e(e^T\Pi(b-Ax)+\|b-Ax\|_1) \\
&< \tfrac{L^{LO}}{mL^{UP}} e(ne^T(b-Ax)+\|b-Ax\|_1) \\
&= \tfrac{L^{LO}}{mL^{UP}} e(n\|b-Ax\|_1+\|b-Ax\|_1) \\
&= \tfrac{L^{LO}}{mL^{UP}} (n+1)\|b-Ax\|_1 e \\
&\leq \tfrac{L^{LO}}{m}(n+1)e \\
&\leq \tfrac{n+1}{m}(b-Ax) \\
&\leq (b-Ax)
\end{aligned}
\tag{C.11}
$$

where we used that $n+1 \leq m$. Thus, by construction, we have that

$$
(\overline{\Pi}, x, \bar{t}) = (\Pi + \Delta\gamma e\rho^T, x, t + \Delta\gamma(b-Ax) + \Delta\gamma(b-Ax-e\rho^T(b-Ax))) \in \Gamma(\gamma).
$$

Next we can now bound the potential of $(\overline{\Pi}, x, \bar{t})$ as follows

$$
\begin{aligned}
P(\gamma) \leq P(\overline{\Pi}, x, \bar{t}) &= \sum_{i=1}^{m}\sum_{k=1}^{m} -\log(\pi_{ik}+\Delta\gamma\rho_k) - \\
&\quad - \sum_{i=1}^{m} \log(t_i + \Delta\gamma(b-Ax)_i + \Delta\gamma(b-Ax-\rho^T(b-Ax))) \\
&\leq -\sum_{i=1}^{m}\sum_{k=1}^{m} \log(\Delta\gamma\rho_k) - \sum_{i=1}^{m} \log(\Delta\gamma(b-Ax)_i) \\
&\leq -\sum_{i=1}^{m}\sum_{k=1}^{m} \log\left(\frac{\varepsilon}{2}\rho_k\right) - \sum_{i=1}^{m} \log\left(\frac{\varepsilon}{2}(b-Ax)_i\right) \\
&\leq -\sum_{i=1}^{m}\sum_{k=1}^{m} \log\left(\frac{\varepsilon L^{LO}}{2mL^{UP}}\right) - \sum_{i=1}^{m} \log\left(\frac{\varepsilon L^{LO}}{2}\right) \\
&= m^2 \log\left(\frac{2mL^{UP}}{\varepsilon L^{LO}}\right) + m\log\left(\frac{2}{\varepsilon L^{LO}}\right)
\end{aligned}
\tag{C.12}
$$

□

Potential Reduction Method:

Step 0. Let (Π^0, x^0, t^0) a 1/8-approx. of the center of $\Gamma(n)$, $k = 0$, $\gamma^k = n$

Step 1. Set $\gamma^{k+1} = \gamma^k - \Delta\gamma^k$

Step 2. Compute the new 1/8-approx. for γ^{k+1} using Centering Step

Step 3. If $P(\gamma^{k+1}) > m^2 \log\left(\frac{2mL^{UP}}{\varepsilon L^{LO}}\right) + m \log\left(\frac{2}{\varepsilon L^{LO}}\right)$, stop

Step 4. Set $k := k + 1$ and goto Step 1.

C.1.1 Primal-Dual Method

In this case, it is a standard linear programming problem. The only issue that needs to be considered is the relative precision. In this case, we know that the $\mathbf{sym}(S) \in [1/n, 1]$. Thus, it suffices to bound the duality gap by ε/n.

Theorem 32 *Let $\varepsilon \in (0, 1/10)$ be given. Suppose that $n \geq 2$ and x^a is a $\beta = \frac{1}{8}$-approximate analytic center of S. Then starting with x^a and using a standard feasible interior-point method to solve each of the linear program, we will compute an ε-approximate symmetry point of S in no more than*

$$O(m \ln\left(\frac{mn}{\varepsilon}\right))$$

total iterations of Newton's method.

C.1.2 Dual Method

As mentioned in Chapter 2, the dual method (Exact Method) is a two-stage method. From a complexity perspective, it is desirable to consider solving the $m+1$ linear programs of the Exact Method for a feasible and near-optimal solution. Ordinarily, this would be easy to analyze. But in this case, the approximately optimal solution to the m linear programs (2.40) will then yield imprecise *input* data for the linear program (2.42). Nevertheless, one can construct an inexact method with an appropriately good complexity bound. Below is a description of such a method.

Inexact Method:

Step 1 For $i = 1, \ldots, m$, approximately solve the linear program (2.40), stopping each linear program when a feasible solution $\bar{x}$ is computed for which the duality gap $\bar{g}$ satisfies $\bar{g} \leq \frac{\varepsilon(b_i - A_{i\cdot}\bar{x})}{4.1}$. Set $\bar{\delta}_i \leftarrow -A_{i\cdot}\bar{x}$.

Step 2 Let $\bar{\delta} := (\bar{\delta}_1, \ldots, \bar{\delta}_m)$. Approximately solve the linear program

$$\begin{array}{ll} \max\limits_{x,\theta} & \theta \\ \text{s.t.} & Ax + \theta(\bar{\delta} + b) \leq b, \end{array} \tag{C.13}$$

stopping when a feasible solution $(\bar{x}, \bar{\theta})$ is computed for which the duality gap $\bar{g}$ satisfies $\bar{\theta} \geq (\bar{\theta} + \bar{g})(1 - \frac{\varepsilon}{4.1})$. Then $\bar{x}$ is an ε-approximate symmetry point of S and $\frac{\bar{\theta}}{1-\bar{\theta}}(1 - \varepsilon/2) \leq \mathbf{sym}(S) \leq \frac{\bar{\theta}}{1-\bar{\theta}}(1 + 2\varepsilon/3)$.

Notice that this method requires that the LP solver computes primal and dual feasible points (or simply primal feasible points and the duality gap) at each of its iterations; such a requirement is satisfied, for example, by a standard feasible interior-point method, see Appendix B.2.

In order to prove a complexity bound for the Inexact Method, we will assume that S is bounded and has an interior, and that an approximate analytic center x^a of the system $Ax \leq b$ has already been computed; for details also see Appendix B.2.

Theorem 33 *Let $\varepsilon \in (0, 1/10)$ be given. Suppose that $n \geq 2$ and x^a is a $\beta = \frac{1}{8}$-approximate analytic center of S. Then starting with x^a and using a standard feasible interior-point method to solve each of the linear programs in Steps 1 and 2, the Inexact Method will compute an ε-approximate symmetry point of S in no more than*

$$\left\lceil 10m^{1.5} \ln\left(\frac{10m}{\varepsilon}\right)\right\rceil$$

total iterations of Newton's method.

The following proposition validates the assertions made at the end of Step 2 of the Inexact Method.

Proposition 10 *Let $\varepsilon \in (0, 1/10)$ be given, set $\tilde{\varepsilon} := \varepsilon/4.1$, and suppose that Steps 1 and 2 of the Inexact Method are executed, with output $(\bar{x}, \bar{\theta})$. Then*

(i) $\bar{\delta} = (\bar{\delta}_1, \ldots, \bar{\delta}_m)$ *satisfies* $(1-\tilde{\varepsilon})(b_i + \delta_i^*) \leq (b_i + \bar{\delta}_i) \leq (b_i + \delta_i^*)$ *for* $i = 1, \ldots, m$.

(ii) For any given $x \in S$, $\theta := \min_i \left\{ \frac{b_i - A_{i.}x}{\delta_i + b_i} \right\}$ *satisfies*

$$\mathbf{sym}(x, S) \in \left[\frac{\theta}{1-\theta} \left(1 - \frac{2\tilde{\varepsilon}}{1-\tilde{\varepsilon}} \right), \frac{\theta}{1-\theta} \right] ,$$

(iii) $\mathbf{sym}(\bar{x}, S) \geq (1-\varepsilon)\mathbf{sym}(S)$, *and*

(iv) $\frac{\bar{\theta}}{1-\bar{\theta}}(1 - \varepsilon/2) \leq \mathbf{sym}(S) \leq \frac{\bar{\theta}}{1-\bar{\theta}}(1 + 2\varepsilon/3)$.

Proof For a given $i = 1, \ldots, m$ let $\bar{g}$ denote the duality gap computed in the stopping criterion of Step 1 of the Inexact Method. Then $\delta_i^* \geq \bar{\delta}_i \geq \delta_i^* - \bar{g} \geq \delta_i^* - \tilde{\varepsilon}(b_i - A_{i.}\bar{x}) \geq \delta_i^* - \tilde{\varepsilon}(b_i + \delta_i^*)$, which implies

$$(1-\tilde{\varepsilon})(b_i + \delta_i^*) \leq (b_i + \bar{\delta}_i) \leq (b_i + \delta_i^*) ,$$

thus proving (i). To prove (ii), let $x \in S$ be given and let $\alpha := \mathbf{sym}(x, S)$ and $\check{\theta} := \min_i \left\{ \frac{b_i - A_{i.}x}{\delta_i^* + b_i} \right\}$. Then from Proposition 1 we have

$$\alpha = \min_i \left\{ \frac{b_i - A_{i.}x}{\delta_i^* + A_{i.}x} \right\} = \frac{\check{\theta}}{1-\check{\theta}} . \tag{C.14}$$

Notice that $\bar{\delta}_i \leq \delta_i^*$ for all i, whereby $\theta \geq \check{\theta}$, which implies that $\alpha = \frac{\check{\theta}}{1-\check{\theta}} \leq \frac{\theta}{1-\theta}$. We also see from (C.14) that $\check{\theta} \leq 1/2$. Next notice that (i) implies that

$$1/2 \geq \check{\theta} \geq \theta(1-\tilde{\varepsilon}) . \tag{C.15}$$

Therefore

$$\begin{aligned}
\alpha = \frac{\check{\theta}}{1-\check{\theta}} &\geq \frac{\theta(1-\tilde{\varepsilon})}{1-\check{\theta}} = \frac{\theta(1-\tilde{\varepsilon})}{1-\theta}\frac{1-\theta}{1-\check{\theta}} \\
&= \frac{\theta(1-\tilde{\varepsilon})}{1-\theta}\left(1+\frac{\check{\theta}-\theta}{1-\check{\theta}}\right) \geq \frac{\theta(1-\tilde{\varepsilon})}{1-\theta}\left(1+\frac{\check{\theta}-\frac{1}{1-\tilde{\varepsilon}}\check{\theta}}{1-\check{\theta}}\right) \\
&= \frac{\theta(1-\tilde{\varepsilon})}{1-\theta}\left(1+\frac{\check{\theta}(\frac{-\tilde{\varepsilon}}{1-\tilde{\varepsilon}})}{1-\check{\theta}}\right) \geq \frac{\theta(1-\tilde{\varepsilon})}{1-\theta}\left(1-\frac{\tilde{\varepsilon}}{1-\tilde{\varepsilon}}\right) \\
&\geq \frac{\theta}{1-\theta}\left(1-\frac{2\tilde{\varepsilon}}{1-\tilde{\varepsilon}}\right) ,
\end{aligned} \tag{C.16}$$

where the next-to-last inequality follows from $\check{\theta} \in [0, 1/2]$, thereby showing (ii).

Let θ^* denote the optimal objective value of (C.13), and notice that $\bar{\delta} \leq \delta^*$ implies that $\theta^* \geq \check{\theta}^*$. Now let $\bar{g}$ be the duality gap computed when the stopping criterion in Step 2 is met. Then

$$\theta \geq \bar{\theta} \geq (\bar{\theta}+\bar{g})(1-\tilde{\varepsilon}) \geq \theta^*(1-\tilde{\varepsilon}) \geq \check{\theta}^*(1-\tilde{\varepsilon}) . \tag{C.17}$$

From (ii) and (C.17) we have

$$\begin{aligned}
\mathbf{sym}(\bar{x}, S) &\geq \frac{\theta}{1-\theta}\left(1-\frac{2\tilde{\varepsilon}}{1-\tilde{\varepsilon}}\right) \geq \frac{\check{\theta}^*(1-\tilde{\varepsilon})}{1-\check{\theta}^*(1-\tilde{\varepsilon})}\left(1-\frac{2\tilde{\varepsilon}}{1-\tilde{\varepsilon}}\right) \\
&= \frac{\check{\theta}^*(1-\tilde{\varepsilon})}{1-\check{\theta}^*}\left(1-\frac{2\tilde{\varepsilon}}{1-\tilde{\varepsilon}}\right)\frac{1-\check{\theta}^*}{1-\check{\theta}^*(1-\tilde{\varepsilon})} \\
&\geq \mathbf{sym}(S)(1-\tilde{\varepsilon})\left(1-\frac{2\tilde{\varepsilon}}{1-\tilde{\varepsilon}}\right)\left(\frac{1/2}{1-1/2+(1/2)\tilde{\varepsilon}}\right) \\
&= \mathbf{sym}(S)(1-\tilde{\varepsilon})\left(1-\frac{2\tilde{\varepsilon}}{1-\tilde{\varepsilon}}\right)\left(1-\frac{\tilde{\varepsilon}}{1+\tilde{\varepsilon}}\right) \\
&\geq \mathbf{sym}(S)\left(1-\tfrac{4\tilde{\varepsilon}}{1-\tilde{\varepsilon}}\right) \geq \mathbf{sym}(S)(1-\varepsilon) ,
\end{aligned} \tag{C.18}$$

where the middle inequality uses the fact that $\check{\theta}^* \in [0, 1/2]$, and the final inequality uses the fact that $\epsilon \in (0, 1/10]$, thus showing (iii).

To prove (iv), note that

$$\mathbf{sym}(S) \geq \mathbf{sym}(\bar{x}, S) \geq \frac{\theta}{1-\theta}\left(1 - \frac{2\tilde{\varepsilon}}{1-\tilde{\varepsilon}}\right) \geq \frac{\theta}{1-\theta}\left(1 - \frac{\varepsilon}{2}\right) \geq \frac{\bar{\theta}}{1-\bar{\theta}}\left(1 - \frac{\varepsilon}{2}\right)$$

where the second inequality follows from part (ii), the third inequality follows since $\varepsilon \leq 1/10$, and the fourth inequality uses $\bar{\theta} \leq \theta$. Last of all, we have

$$\mathbf{sym}(S) = \frac{\check{\theta}^*}{1-\check{\theta}^*} \leq \frac{\frac{\bar{\theta}}{1-\tilde{\varepsilon}}}{1-\frac{\bar{\theta}}{1-\tilde{\varepsilon}}} = \frac{\bar{\theta}}{1-\bar{\theta}-\tilde{\varepsilon}} = \frac{\bar{\theta}}{1-\bar{\theta}}\left(\frac{1-\bar{\theta}}{1-\bar{\theta}-\tilde{\varepsilon}}\right) \leq \frac{\bar{\theta}}{1-\bar{\theta}}\left(1 + \frac{2\varepsilon}{3}\right),$$

where the first equality is from Proposition 7, the first inequality follows from (C.17), and the last inequality follows since $\varepsilon \leq 1/10$ and (C.15) implies that $\bar{\theta} \leq \theta \leq 41/80$. □

It remains to prove the complexity bound of Theorem 33, which will be accomplished with the help of the following two propositions.

Proposition 11 *Let $\varepsilon \in (0, 1/10)$ be given, and set $\tilde{\varepsilon} := \varepsilon/4.1$. Suppose that x^a is a $\beta = \frac{1}{8}$-approximate analytic center of S. Then starting with x^a, the stopping criterion of each linear program in Step 1 will be reached in no more than*

$$\left\lceil (2 + 4\sqrt{m}) \ln\left(\frac{42m}{\varepsilon}\right) \right\rceil$$

iterations of Newton's method.

Proof Step 1 is used to approximately solve each of the linear programs (2.40) for $i = 1, \ldots, m$. Let us fix a given i, and define $\lambda := -e^i$ where e^i is the i^{th} unit vector in $\mathbb{R}^m$. Then from Theorem 27 with $(M, f, c) = (A, b, -A_{i\cdot})$ we can bound the iterations used to solve (2.40) by

$$\left\lceil (2 + 4\sqrt{m}) \ln\left(\frac{10m\|S^a\lambda\|}{\bar{g}}\right)^+ \right\rceil . \tag{C.19}$$

Now notice that $\|S^a\lambda\| = s_i^a$. Let $(\bar{x}, \bar{s})$ denote the primal solution and slack vector computed in Step 1 when the stopping criterion is met. Also, to keep the analysis simple, we assume that the stopping criterion is met exactly. We have:

$$\bar{s}_i = b_i - A_{i\cdot}\bar{x} \geq b_i + \delta_i^* - \bar{g} \geq b_i + \delta_i^* - \tilde{\varepsilon}\bar{s}_i \geq b_i - A_{i\cdot}x^a - \tilde{\varepsilon}\bar{s}_i = s_i^a - \tilde{\varepsilon}\bar{s}_i \ ,$$

whereby $s_i^a \leq \bar{s}_i(1+\tilde{\varepsilon})$. Therefore

$$\frac{10m\|S^a\lambda\|}{\bar{g}} = \frac{10ms_i^a}{\tilde{\varepsilon}\bar{s}_i} \leq \frac{10m(1+\tilde{\varepsilon})}{\tilde{\varepsilon}} = \frac{41m(1+\varepsilon/4.1)}{\varepsilon} \leq \frac{42m}{\varepsilon} \ ,$$

since in particular $\varepsilon \in (0, 1/10)$. Substituting this inequality into (C.19) completes the proof. □

Proposition 12 *Let $\varepsilon \in (0, 1/10)$ be given, $m \geq 3$ and set $\tilde{\varepsilon} := \varepsilon/4.1$. Suppose that x^a is a $\beta = \frac{1}{8}$-approximate analytic center of S. Then starting with x^a, the stopping criterion of the linear program in Step 2 will be reached in no more than*

$$\left\lceil (2 + 4\sqrt{m}) \ln\left(\frac{6m}{\varepsilon}\right) \right\rceil$$

iterations of Newton's method.

Proof Let $s^a = b - Ax^a$ and let z^a denote the dual multipliers associated with (B.2) for $M = A$ and $f = b$. It follows from (B.2) and $m \geq 3$ that

$$(s^a)^T z^a = e^T(S^a z^a - e + e) \geq -\frac{1}{8}\sqrt{m} + m \geq \frac{9m}{10} \ . \tag{C.20}$$

Setting $(M, f, d) = (A, b, (\bar{\delta} + b))$ we see that (C.13) is an instance of (B.3), and from Theorem 28 we can bound the iterations used to solve (C.13) by

$$\left\lceil (2 + 4\sqrt{m}) \ln\left(\frac{1.25m}{\bar{g} \cdot (\bar{\delta} + b)^T z^a}\right)^+ \right\rceil \ . \tag{C.21}$$

We have

$$(\bar{\delta}+b)^T z^a \geq (b+\delta^*)^T z^a(1-\tilde{\varepsilon}) \geq (s^a)^T z^a(1-\tilde{\varepsilon}) \geq \frac{9m(1-\tilde{\varepsilon})}{10}$$

where the first inequality follows from part (i) of Proposition 10, the second inequality follows from $b+\delta^* \geq b - Ax^a = s^a$, and the third inequality follows form (C.20). We next bound $\bar{g}$. To keep things simple we again assume that the stopping criterion in Step 2 is satisfied exactly, whereby

$$\frac{1}{\bar{g}} = \frac{1-\tilde{\varepsilon}}{\tilde{\varepsilon}}\frac{1}{\bar{\theta}} \leq \frac{1}{\tilde{\varepsilon}\cdot\check{\theta}^*} = \frac{4.1}{\varepsilon}\left(1+\frac{1}{\mathbf{sym}(S)}\right) \leq \frac{4.1}{\varepsilon}(1+n) \leq \frac{4.1m}{\varepsilon}\ .$$

Here the first inequality follows from (C.17), the second equality follows from Proposition 7, the second inequality follows from Remark 3, and the last inequality follows since S is assumed to be bounded and so $m \geq n+1$. Combining the bounds on $(\bar{\delta}+b)^T z^a$ and $\bar{g}$ we then bound the logarithm term in the statement of the proposition as follows:

$$\frac{1.25m}{\bar{g}\cdot(\bar{\delta}+b)^T z^a} \leq \frac{1.25m\cdot 4.1m\cdot 10}{9m\varepsilon(1-\tilde{\varepsilon})} \leq \frac{6m}{\varepsilon}\ ,$$

since $\varepsilon \in (0, 1/10)$ implies that $\tilde{\varepsilon} \leq 1/41$. This completes the proof.□

Proof *of complexity bound of Theorem 33:* From Propositions 11 and 12 it follows that the total number of Newton steps computed by the Inexact Method is bounded from above by:

$$m\left\lceil(2+4\sqrt{m})\ln\left(\frac{42m}{\varepsilon}\right)\right\rceil + \left\lceil(2+4\sqrt{m})\ln\left(\frac{6m}{\varepsilon}\right)\right\rceil \leq \left\lceil 10m^{1.5}\ln\left(\frac{10m}{\varepsilon}\right)\right\rceil$$

since $m \geq n+1 \geq 3$ and $\varepsilon < 1/10$. □

Appendix D

Projective Preconditioner

D.1 A Primal-feasible Interior-Point Algorithm for $F_{\bar{s}}$ and its Complexity Analysis

Recall the parameterized barrier problem OP_η, and let $(x, \theta) = (x^k, \theta^k)$ be a feasible solution of OP_η for some $\eta > 0$. Then the locally quadratic model of OP_η at (x^k, θ^k) yields the following Newton equation system in variables (d, Δ, π, q):

$$\begin{array}{rrrcl} & & \bar{x}^T A^T \pi & = & \eta \\ H(x^k)d & & +A^T\pi \quad +\bar{s}\cdot q & = & -\nabla f(x^k) \\ Ad & +A\bar{x}\Delta & & = & 0 \\ \bar{s}^T d & & & = & 0\,. \end{array} \tag{D.1}$$

Here $\nabla f(x)$ and $H(x)$ denote the gradient and Hessian of $f(x)$, respectfully, (d, Δ) is the Newton step for the variables (x, θ), and (π, q) are the multipliers on the two linear equations of OP_η.

A primal-feasible interior-point algorithm for solving OP computes a sequence of approximate solutions (x^k, θ^k) of OP_{η^k} for an increasing sequence of values of η^k. Let $(\tilde{d}, \tilde{\Delta}, \tilde{\pi}, \tilde{q})$ solve the Newton system (D.1) using $\eta = \eta^k$, then the Newton step is $(\tilde{d}, \tilde{\Delta})$, and (x^k, θ^k) is defined to be a γ-approximate solution of OP^k_η if the norm of $\tilde{d}$ measured in the ellipsoidal norm induced by the Hessian of $f(\cdot)$ is not greater than

γ, i.e., $\sqrt{(\tilde{d})^T H(x^k)\tilde{d}} \leq \gamma$. The algorithm starts by setting $(x^0, \theta^0) \leftarrow (\bar{x}, -1)$ and by computing a value η^0 for which it is guaranteed that (x^0, θ^0) is a γ-approximate solution of OP_{η^0}. Inductively, if (x^k, θ^k) is defined to be a γ-approximate solution of OP_{η}^k, the algorithm then increases η^k to $\eta^{k+1} = \alpha \cdot \eta^k$ for some fixed value of $\alpha > 1$, and then computes the Newton step for (x^k, θ^k) for $\text{OP}_{\eta^{k+1}}$. The algorithm continues in this manner until $\theta^k \geq 0$ at some iteration k, at which point the algorithm stops and $\hat{x} := \frac{x^k + \theta^k \bar{x}}{1+\theta^k}$ is a feasible solution of $F_{\bar{s}}$. A formal statement of the algorithm is as follows:

Algorithm A

Step 1. (Initialization) Set $\bar{x} \leftarrow -\nabla f^*(\bar{s})/\vartheta$. If $A\bar{x} = 0$, STOP. Otherwise, set $k = 0$, $(x^0, \theta^0) \leftarrow (\bar{x}, -1)$, and define the following constants:

$$\gamma = \frac{1}{9}\,, \quad \beta = \frac{1}{4}\,, \quad \alpha = \frac{\sqrt{\vartheta} + \beta}{\sqrt{\vartheta} + \gamma}\,.$$

Temporarily set $\eta = 1$ and solve (D.1) for $(\tilde{d}, \tilde{\Delta}, \tilde{\pi}, \tilde{q})$ and set

$$\eta^0 = \frac{\gamma}{\sqrt{\tilde{d}^T H(\bar{x})\tilde{d}}}\,.$$

Step 2. (Increase η) $\eta^{k+1} \leftarrow \alpha \cdot \eta^k$

Step 3. (Compute and Take Newton Step) Set $\eta = \eta^{k+1}$, solve (D.1) for $(\tilde{d}, \tilde{\Delta}, \tilde{\pi}, \tilde{q})$. Set $(x^{k+1}, \theta^{k+1}) = (x^k, \theta^k) + (\tilde{d}, \tilde{\Delta})$

Step 4. (Test Solution) If $\theta^{k+1} \geq 0$, set $\tilde{x} := \frac{x^k + \theta^k \bar{x}}{1+\theta^k}$ and STOP. Otherwise set $k \leftarrow k + 1$ and go to **Step 2**.

In order to validate this algorithm, we will need to prove the following results. Let d^k denote the value of $\tilde{d}$ in (D.1) at (x^k, θ^k) using $\eta = \eta^k$. The norm of this Newton step in the norm induced by the Hessian $H(x^k)$ and is given by:

$$\|(d^k)\|_{x^k} := \sqrt{(d^k)^T H(x^k) d^k}\,.$$

Proposition 13 $\|(d^0)\|_{x^0} = \gamma$.

Proof Following Step 1 of Algorithm A, let $(\tilde{d}, \tilde{\Delta}, \tilde{\pi}, \tilde{q})$ solve (D.1) using $x^k = x^0 = \bar{x}$ and $\eta = 1$, and set η^0 by the prescribed formula. Using the fact that $\bar{s} = -\nabla f(\bar{x})/\vartheta$ (from (vi) of Remark 8 and Theorem 2.3.9 of [52]), it follows from direct substitution that $(d^0, \Delta^0, \pi^0, q^0) := (\eta^0\tilde{d}, \eta^0\tilde{\Delta}, \eta^0\tilde{\pi}, \eta^0\tilde{q} + \vartheta - \eta^0\vartheta)$ solves (D.1) using $x^k = x^0 = \bar{x}$ and $\eta = \eta^0$. Therefore $\|(d^0)\|_{x^0} = \sqrt{(d^0)^T H(x^0)(d^0)} = \eta^0\sqrt{\tilde{d}^T H(\bar{x})\tilde{d}} = \gamma$. □

Proposition 14 $\|(d^k)\|_{x^k} \leq \gamma$ *for all* $k = 1, 2, \ldots$.

Proof The proof is by induction, and for clarity we suppress the iteration counter k. Suppose that our current point is (x, θ). Let $d_{\bar{\eta}}$ denote the x-part of the Newton step for the parameter value $\eta = \bar{\eta}$. Then $d_{\bar{\eta}}$ can be decomposed as $d_{\bar{\eta}} = d^c + \bar{\eta}d^a$ where d^c is the centering direction and d^a is the affine scaling direction. It follows from the fact that $f(\cdot)$ has complexity value ϑ that $\|d^c\|_x \leq \sqrt{\vartheta}$. Furthermore, by assumption of the induction we have $\|d_{\bar{\eta}}\|_x \leq \gamma$. Then according to Step 2 of Algorithm A we increase η by the factor α. We write $d_{\alpha\bar{\eta}} = d^c + \alpha\bar{\eta}d^a = \alpha(d^c + \bar{\eta}d^a) + (1-\alpha)d^c = \alpha d_{\bar{\eta}} + (1-\alpha)d^c$, whereby:

$$\|d_{\alpha\bar{\eta}}\|_x \leq \alpha\|d_{\bar{\eta}}\|_x + (\alpha - 1)\|d^c\|_x \leq \alpha\gamma + (\alpha - 1)\sqrt{\vartheta} = \beta \ .$$

Letting $x^+ := x + d_{\alpha\bar{\eta}}$ denote the new value of x and letting $d^+_{\alpha\bar{\eta}}$ denote the Newton step at x^+ for the parameter value $\eta = \alpha\bar{\eta}$ it follows from Theorem 2.2.4 of [52] that

$$\|d^+_{\alpha\bar{\eta}}\|_{x^+} \leq \frac{\beta^2}{(1-\beta)^2} = \gamma \ ,$$

completing the inductive proof. □

Proposition 15 $\eta^0 \geq \dfrac{\gamma}{1 + \theta^*}$

Proof Since $(x, \theta) = (\bar{x}, -1)$ is feasible for (3.4) and from the self-concordance of

$f(\cdot)$ we have $\{\bar{x}\} + \{d : d^T H(\bar{x}) d \leq 1\} \subset C$, it follows that

$$\theta^* \geq \quad -1 + \quad \max_{d,\Delta} \quad \Delta$$

$$\begin{aligned} Ad \quad & +(A\bar{x})\Delta \;=\; 0 \\ \bar{s}^T d \quad & \qquad\quad\; = \; 0 \\ d^T H(\bar{x}) d \quad & \qquad\quad\; \leq \; 1 \,. \end{aligned}$$

Letting (d, Δ) solve the above maximization problem, to prove the proposition it therefore suffices to show that $\Delta = \gamma/\eta^0$, which we now do. The optimality conditions for the maximization problem above can be written as:

$$\begin{aligned} \bar{x}^T A^T \pi &= 1 \\ \rho H(\bar{x}) d + A^T \pi + \bar{s} \cdot q &= 0 \\ Ad + A\bar{x}\Delta &= 0 \\ \bar{s}^T d &= 0 \\ \rho &\geq 0 \\ d^T H(\bar{x}) d &\leq 1 \\ \rho \cdot (1 - d^T H(\bar{x}) d) &= 0 \,. \end{aligned} \tag{D.2}$$

Following Step 1 of Algorithm A, let $(\tilde{d}, \tilde{\Delta}, \tilde{\pi}, \tilde{q})$ solve (D.1) using $x^k = x^0 = \bar{x}$ and $\eta = 1$, and set η^0 by the prescribed formula. One then easily checks that

$$(d, \Delta, \pi, q, \rho) := \left(\frac{\tilde{d}}{\sqrt{\tilde{d}^T H(\bar{x}) \tilde{d}}}, \frac{\tilde{\Delta}}{\sqrt{\tilde{d}^T H(\bar{x}) \tilde{d}}}, \tilde{\pi}, \tilde{q} - \vartheta, \sqrt{\tilde{d}^T H(\bar{x}) \tilde{d}} \right)$$

satisfies (D.2) (again using the fact that $\bar{s} = -\nabla f(\bar{x})/\vartheta$ from (vi) of Remark 8 and Theorem 2.3.9 of [52]), and so (d, Δ) is an optimal solution of the maximization problem. Straightforward manipulation of the system (D.2) then shows that $\Delta = \sqrt{\tilde{d}^T H(\bar{x}) \tilde{d}}$, and so from the definition of η^0 we have $\Delta = \gamma/\eta^0$. $\square$

Before proving the next proposition, we state a result which is implicit in Renegar

[52], but is not stated explicitly. Rather than re-develop the notation and set-up of [52], we simply state the result and give a proof as if it appeared as part of the text of Chapter 2 of [52].

Lemma 38 (Essentially from Renegar [52]) *Under the notation and conditions of Chapter 2 of Renegar [52], suppose that y is an iterate of the barrier method for the barrier parameter η and the Newton step $n(y)$ for the function $f_\eta(x) := \eta\langle c, x\rangle + f(x)$ at y satisfies $\|n(y)\|_y \leq \gamma$ where $\gamma < 1/4$. Then*

$$c^T y \leq \mathrm{VAL} + \frac{\vartheta}{\eta}\left(\frac{1}{1-\delta}\right)$$

where $\delta = \gamma + \frac{3\gamma^2}{(1-\gamma)^3}$.

Proof Letting $z(\eta)$ denote the analytic center of the function $f_\eta(\cdot)$, it follows from Theorem 2.2.5 of [52] that $\|y - z(\eta)\|_y \leq \delta$, and hence from the self-concordance property that $\|y - z(\eta)\|_{z(\eta)} \leq \delta/(1-\delta)$. From inequality (2.14) of Section 2.4.1 of [52] we have:

$$c^T y \leq \mathrm{VAL} + \frac{\vartheta}{\eta}\left(1 + \|y - z(\eta)\|_{z(\eta)}\right) \leq \mathrm{VAL} + \frac{\vartheta}{\eta}\left(1 + \frac{\delta}{1-\delta}\right) = \mathrm{VAL} + \frac{\vartheta}{\eta}\left(\frac{1}{1-\delta}\right) .$$

□

Proposition 16 *Suppose (x, θ) is an iterate of Algorithm A for the parameter value η and $\eta \geq \vartheta/(\theta^*(1-\delta))$ where $\delta = \gamma + \frac{3\gamma^2}{(1-\gamma)^3}$. Then $\theta \geq 0$.*

Proof Converting Lemma 38 to the notation of problem (3.5) and Algorithm A, we have:

$$\theta \geq \theta^* - \frac{\vartheta}{\eta}\left(\frac{1}{1-\delta}\right) \geq \theta^* - \theta^* = 0 ,$$

where the last inequality follows from the supposition that $\eta \geq \vartheta/(\theta^*(1-\delta))$. □

Proof of Theorem 15: We first prove the iteration bound. Note that the parameter values set in Step 1 of Algorithm A imply:

$$1 - \frac{1}{\alpha} = \frac{1}{7.2\sqrt{\vartheta} + 1.8} \geq \frac{1}{9\sqrt{\vartheta}} . \tag{D.3}$$

Define δ as in the hypothesis of Proposition 16. Let

$$J := \left\lceil 9\sqrt{\vartheta} \ln\left(11\vartheta\left(1 + \frac{1}{\mathbf{sym}(0, H_{\bar{s}})}\right)\right)\right\rceil .$$

From Step 2 of Algorithm A we have $\eta^0/\eta^J = (1/\alpha)^J$ and taking logarithms we obtain $\ln(\eta^0) - \ln(\eta^J) = J\ln(1/\alpha) \leq J(1/\alpha - 1)$, and rearranging we obtain:

$$\begin{aligned}
\ln(\eta^J) &\geq \ln(\eta^0) + J\left(1 - \tfrac{1}{\alpha}\right) \\
&\geq \ln\left(\tfrac{\gamma}{1+\theta^*}\right) + \ln\left(11\vartheta\left(1 + \tfrac{1}{\mathbf{sym}(0,H_{\bar{s}})}\right)\right) && \text{(from Prop.15 and } (D.3)) \\
&\geq \ln\left(\tfrac{\gamma}{1+\theta^*}\right) + \ln\left(11\vartheta\left(\tfrac{\theta^*+1}{\theta^*}\right)\right) && \text{(from (3.6))} \\
&= \ln\left(\tfrac{\vartheta}{\theta^*(1-\delta)}\right) + \ln\left(11\gamma(1-\delta)\right) \geq \ln\left(\tfrac{\vartheta}{\theta^*(1-\delta)}\right),
\end{aligned}$$

where the last inequality follows since $11\gamma(1-\delta) \geq 1$ for the specific values of γ and δ given. Then from Proposition 16 it follows that $\theta^J \geq 0$, and so Algorithm A will stop.

It remains to prove the bound on $\mathrm{reldist}(\tilde{x}, \partial C)$. For $x \in \mathrm{int}C$, define the norm $\|\cdot\|_x$ by $\|v\|_x := \sqrt{v^T H(x) v}$ for $x \in \mathrm{int}C$, and for $x \in C$ define:

$$B_x(c, r) := \{y : A(y - c) = 0, \bar{s}^T(y - c) = 0, \|y - c\|_x \leq r\} .$$

Letting z denote the analytic center of $\mathcal{F}_{\bar{s}}$ it follows from Lemma 5 of [48] that $B_z(z, \vartheta) \supset \mathcal{F}_{\bar{s}}$. Assuming for simplicity that $\theta = 0$ for the final iterate of Algorithm A, we have $\|z - \tilde{x}\|_{\tilde{x}} \leq \delta := \gamma + 3\gamma^2/(1-\gamma)^3$ from Theorem 2.2.5 of [52], whereby $B_{\tilde{x}}(\tilde{x}, \vartheta/(1-\delta) + \delta/(1-\delta)) \supset \mathcal{F}_{\bar{s}}$ follows from the self-concordance of $f(\cdot)$. Since it is also true that $B_{\tilde{x}}(\tilde{x}, 1) \subset \mathcal{F}_{\bar{s}}$ it follows that $\sigma := \mathbf{sym}(\tilde{x}, \mathcal{F}_{\bar{s}}) \geq (1-\delta)/(\vartheta + \delta) \geq 1/(1.2\vartheta + 0.2)$ for the specific values of γ and δ herein. Finally, noting that there exists $\hat{x} \in \mathcal{F}_{\bar{s}}$ satisfying $B_{\bar{s}}(\hat{x}, \tau_F) \subset C$, where $B_{\bar{s}}(c, r)$ is the ball centered

at c of radius r in the $\bar{s}$-norm, it follows from the symmetry value of $\tilde{x}$ in $\mathcal{F}_{\bar{s}}$ that $B_{\bar{s}}(\tilde{x}-\sigma(\hat{x}-\tilde{x}),\sigma\tau_F) \subset C$, whereby taking convex combinations with $B_{\bar{s}}(\hat{x},\tau_F)$ yields

$$B_{\bar{s}}\left(\tilde{x},\frac{\tau_F}{1.2\vartheta+0.2}\right) \subset B_{\bar{s}}\left(\tilde{x},\sigma\tau_F\right) \subset B_{\bar{s}}\left(\tilde{x},\frac{2\sigma\tau_F}{1+\sigma}\right)$$

$$B_{\bar{s}}\left(\tilde{x},\frac{2\sigma\tau_F}{1+\sigma}\right) = \frac{1}{1+\sigma}B_{\bar{s}}(\tilde{x}-\sigma(\hat{x}-\tilde{x}),\sigma\tau_F) + \frac{\sigma}{1+\sigma}B_{\bar{s}}(\hat{x},\tau_F) \subset C\ ,$$

from which it follows that $\text{reldist}(\tilde{x},\partial C) \geq \tau_F/(1.2\vartheta+0.2)$.□

Appendix E

Binary Search and Newton's Method for solving (4.11)

Here we want analyze the optimization problem (4.11)

$$\begin{aligned} \theta^* = \min_v \quad & \langle v, A^*\bar{\lambda} \rangle + \bar{\mu}\|v\| \\ & v^T(A^*A + I)v \leq 1. \end{aligned}$$

to test if $\theta^* < -t$, where we usually want $t = \sigma/2 = 1/(64n)$.

To simplify our exposition we will cast (4.11) as follows

$$\begin{aligned} \min_v \quad & \langle c, v \rangle + \bar{\mu}\|v\| \\ & v^T M v \leq 1. \end{aligned} \tag{E.1}$$

where $\|c\| = \|A^*\bar{\lambda}\| > \bar{\mu}$, $\bar{\mu} \in \{-1, 0, 1\}$, and $M \succcurlyeq I$.

Assuming that $\bar{\mu} = 0$, (4.11) has a closed form solution which can be computed in $O(n^3)$ operations. We have that $\theta^* = \sqrt{c^T M^{-1} c}$ and $v^* = M^{-1}c/\sqrt{c^T M^{-1} c}$.

We focus our analysis on the case of $\bar{\mu} = 1$ since the other case can be addressed by the same methodology after a suitable preprocessing.

Before we proceed, we briefly review results on [55] and [62] which will be relevant in our analysis.

E.1 Combining Binary Search and Newton's Methods

To introduce the idea, we restrict our discussion to the univariate case, which will be the relevant case for solving (E.1), but many of the results hold for more general problems.

Consider the problem of solving the system $f(\xi) = 0$, where f is an analytic function. Following Newton's method, given a starting point z_0 we can define inductively

$$z_k = z_{k-1} - f(z_{k-1})/f'(z_{k-1}).$$

Definition 8 *We say that z_0 is an approximate zero of f if*

$$|z_k - z_{k-1}| \leq (1/2)^{2^{k-1}-1}|z_1 - z_0| \quad \text{for } k \geq 1.$$

That is, Newton method starting from z_0 converges quadratically from the very first iteration. In fact, we have

Proposition 17 *(Smale, [55]) If z_0 is an approximate zero and $z_n \to \xi$ and $n \to \infty$, then*

$$|z_n - \xi| \leq 2(1/2)^{2^{k-1}}|z_1 - z_0|$$

In [55] Smale provides sufficient conditions for a point z to be an approximate zero. The main theorem in [55] is as follows.

Theorem 34 *(Smale, [55]) Let f be an analytic function and assume that z is such that*

$$\sup_{k>1} \left| \frac{f^{(k)}(z)}{k! f'(z)} \right|^{1/(k-1)} \leq \frac{1}{8} \left| \frac{f'(z)}{f(z)} \right|.$$

Then z is an approximate zero for f.

The assumption in the theorem ensures that the higher order derivatives of f are small enough so that Newton's Method converges quadratically. Ye [62] builds upon that result and establishes the following theorem.

Theorem 35 *(Ye, [62])Let f be an analytic, convex, and monotonically decreasing in $[0, \infty)$. Furthermore, for every $z \geq 0$ and $k > 1$ let*

$$\left| \frac{f^{(k)}(z)}{k! f'(z)} \right|^{1/(k-1)} \leq \frac{\alpha}{z}$$

for some $\alpha > 0$. If the zero of f lies in the interval $\xi \in [\hat{z}, (1 + 1/8\alpha)\hat{z}]$, $\hat{z}$ is an approximate zero of f.

Moreover, it is also shown in [62] that $f(z) = c^T(Q + zI)^{-2}c - 1$ satisfies the conditions of Theorem 35 with $\alpha = 3/2$.

Next, assume that we have upper and lower bounds on the value of the root $0 < \underline{z} \leq z \leq \bar{z}$. That is, $f(\underline{z}) > 0$ and $f(\bar{z}) < 0$. We divide the interval $[\underline{z}, \bar{z}]$ into subintervals of the form

$$I^i = [(1 + 1/8\alpha)^i \underline{z}, (1 + 1/8\alpha)^{i+1} \underline{z}] \quad \text{for } i \geq 1.$$

From Theorem 35, if $\xi \in I^i$ we have that $(1 + 1/8\alpha)^i \underline{z}$ is an approximate root. Therefore, Newton's method can be started from this point and we can obtain a point $\hat{\xi}$ such that $|\hat{\xi} - \xi| < \varepsilon$ in at most $O(\ln \ln(\bar{z}/\varepsilon))$ iterations.

In order to find the correct subinterval, we first note that the total number of subintervals is $\lceil \ln(\bar{z}/\underline{z}) / \ln(1 + 1/8\alpha) \rceil$. Conducting a binary search on these intervals requires at most

$$O(\ln \ln(\bar{z}/\underline{z}) - \ln \ln(1 + 1/8\alpha))$$

evaluations of f where $-\ln \ln(1 + 1/8\alpha) \leq -\ln 1/16\alpha = \ln 16\alpha$.

E.2 Solving (E.1): $\bar{\mu} > 0$

Recall that we assume without loss of generality $\bar{\mu} = 1$, $\|c\| \geq 1 + \sigma$ (otherwise the optimal solution is zero, and in our case we have $\sigma = 1/32n$), and $M \succcurlyeq I$. Let $\alpha \geq 0$ denote the Lagrange multiplier associated with the inequality constraint in (E.1).

The gradient optimality conditions can be written as

$$c + \frac{v}{\|v\|} + 2\alpha M v = 0.$$

In fact, we have

$$v^* = -\|v^*\| \left(I + 2\alpha^* \|v^*\| M\right)^{-1} c,$$

which implies that

$$1 = c^T \left(I + 2\alpha^* \|v^*\| M\right)^{-2} c, \text{ or equivalently } 1 = c^T M^{-1} \left(M^{-1} + 2\alpha^* \|v^*\| I\right)^{-2} M^{-1} c.$$

Since $M \succcurlyeq I$, there exists a unique value $\gamma^* = 2\alpha^* \|v^*\|$ for the previous equality to hold. Once $\bar{\gamma}$ is known, we have that

$$1 = v^{*T} M v^* = \|v^*\|^2 c^T \left(I + \gamma^* M\right)^{-1} M \left(I + \bar{\gamma}^* M\right)^{-1} c$$

and the value of $\|v^*\|$ and α^* are easily computed.

In order to efficiently compute γ^*, we use the method proposed in Section E.1. For $z \geq 0$, define the function

$$f(z) = c^T M^{-1} \left(M^{-1} + zI\right)^{-2} M^{-1} c - 1.$$

In order to bound the value of γ^*, note that $f(0) = \|c\|^2 - 1 \geq \sigma > 0$ and $f(f(0)^3) > 0$ from Lemma 3 of [62]. Moreover,

$$f(\|c\|) = c^T M^{-1} \left(M^{-1} + \|c\| I\right)^{-2} M^{-1} c - 1 < \frac{c^T M^{-2} c}{\|c\|^2} - 1 \leq 0$$

since $M^{-1} \succ 0$ and $M \succ I$. Thus, we have

$$\sigma^3 \leq \gamma^* \leq \|c\|.$$

Theorem 36 *(Ye [62]) If $\bar{\mu} > 0$, a ε-solution for γ^* can be computed in at most*

$O(n^3 + n^2 \ln \ln 1/\varepsilon)$ *arithmetic operations.*

This give a new complexity bound for problems of the form (4.11).

The case of $\bar{\mu} < 0$ can be solved via the same method after a proper preprocessing via Renegar's algorithm as described in [62].

Theorem 37 *(Ye, [62]) A ε solution for γ^* can be computer in $O(n^4+n^2(\ln n) \ln \ln 1/\varepsilon)$ arithmetic operations.*

Appendix F

Hit-and-run for $f(x) \sim e^{-\|x\|_2}$

Here we specialize the method for implementing the hit-and-run method for logconcave densities proposed in [41] to the norm-induced density of the form $f(x) \sim e^{-\|x\|_2}$ if $x \in K$, and zero otherwise.

Given $x_0 \in K$, $d \in S^{n-1}$, define the line segment $\ell(x_0, d) = K \cap \{x_0 + td : t \in \mathbb{R}\}$ (note that if it is not a line segment, either d or $-d$ belongs to C_K). The next iterate should be chosen from $\ell(x_0, d)$ according to the density f. Let $M_\ell = \max_{y \in \ell(x_0,d)} f(y)$ denote the maximum density on $\ell(x_0, d)$ and, for $v \in (0, M_\ell)$, let $L_\ell(v)$ denote the corresponding level set restricted to $\ell(x_0, d)$.

It is a two step method:

Step 1. Choose v uniformly on $(0, M_\ell)$;

Step 2. Choose y uniformly on $L_\ell(v)$;

To implement this, we need to compute three quantities: M_ℓ, and the two end points of $L_\ell(v)$. It is possible to speed-up the process by using closed-form formulas since we know the functional form of f (although additional effort is necessary to adjust it for the support of f).

On $\ell(x_0, d)$, f can be written that $f(x_0 + td) \sim e^{-\sqrt{\|x_0\|^2 + 2t\langle x_0, d\rangle + t^2\|d\|^2}}$. It is convenient to note that

$$t^* = -\frac{\langle x_0, d\rangle}{\|d\|^2} \in \arg\max\{f(x_0 + td) : t \in \mathbb{R}\}$$

That is, for $x^* = x_0 + t^* d$, $f(x^*) \geq M_L$. If $x^* \in K$, we have found M_ℓ, otherwise, we need to make a binary search on $[x_0, x^*]$ to find it (note that in the second case we already have one of the endpoints of $L_\ell(v)$).

After drawing $v \in (0, M_\ell)$, again we can compute the explicit "unrestricted" point of where the endpoints should be,

$$\hat{t} \in \left\{ t : v = e^{-\sqrt{\|x_0\|^2 + 2t\langle x_0, d\rangle + t^2 \|d\|^2}} \right\}$$

or equivalently, the solution for the following quadratic equation

$$\hat{t}^2 \|d\|^2 + 2\hat{t}\,\langle x_0, d\rangle + \|x_0\|^2 + \ln^2 v = 0.$$

Again, if the solutions $\hat{t}$ lie in K we found an endpoint, otherwise, we need to conduct a binary search on $[x_0, x_0 + \hat{t}d]$.

Bibliography

[1] T. W. Anderson, *The integral of a symmetric unimodal function*, Proceedings of American Mathematical Society **6** (1955), 170–176.

[2] D. Applegate and R. Kannan, *Sampling and integration of near logconcave functions*, Proceedings 23th ACM STOC (1993), 156–163.

[3] K. Ball, *An elementary introduction to modern convex geometry*, Flavors of Geometry, MSRI Publications **31** (1997).

[4] E. Barnes, *private communication*, (1998).

[5] A. Belloni and V. Chernozhukov, *On the computational complexity of MCMC-based estimators under the central limit theorem framework*, MIT Operations Research Center Working paper (2006).

[6] D. Bertsimas and S. Vempala, *Solving convex programs by random walks*, Journal of the ACM **51** (2004), no. 4, 540–556.

[7] A. Blum, A. Frieze, R. Kannan, and S. Vempala, *A polynomial-time algorithm for learning noisy linear threashold functions*, Algorithmica **22** (1998), no. 1, 35–52.

[8] G. Casella and C. P. Robert, *Monte carlo statistical methods*, Springer, 1999.

[9] G. de Ghellink and J.-P. Vial, *A polynomial Newtow method for linear programming*, Algorithmica (1986), no. 1, 425–453.

[10] R. Dudley, *Uniform central limit thoerems*, Cambridge University Press, 2000.

[11] J. Dunagan and S. Vempala, *A simple polynomial-time rescaling algorithm for solving linear programs*, Proceedings of STOC'04 (2004).

[12] M. E. Dyer and A. M. Frieze, *On the complexity of computing the volume of a polyhedron*, SIAM Journal on Computing archive **17(5)** (1988), 967–974.

[13] M. Epelman and R. M. Freund, *A new condition measure, preconditioners, and relations between different measures of conditioning for conic linear systems*, SIAM Journal on Optimization **12** (2002), no. 3, 627–655.

[14] G. S. Fishman, *Choosing sample path length and number of sample paths when starting at steady state*, Operations Research Letters **16(4)** (1994), 209–220.

[15] R. M. Freund, *Combinatorial analogs of Brouwer's fixed-point theorem on a bounded polyhedron*, Journal of Combinatorial Theory, Series B **47(2)** (1989), 192–219.

[16] ———, *Projective transformation for interior-point algorithms, and a superlinearly convergent algorithm for the w-center problem*, Mathematical Programming **58** (1991), 203–222.

[17] ———, *On the primal-dual geometry of level sets in linear and conic optimization*, SIAM Journal on Optimization **13(4)** (2003), 1004–1013.

[18] R. M. Freund, *On the behavior of the homogeneous self-dual model for conic convex optimization*, MIT Operations Research Center Working paper OR-372-04, submitted to Mathematical Programming (2004).

[19] R. M. Freund and J. R. Vera, *Condition-based complexity of convex optimization in conic linear form via the ellipsoid algorithm*, SIAM Journal on Optimization **10** (1999), no. 1, 155–176.

[20] R. J. Gardner, *The Brunn-Minkowski inequality*, Bulletin of the American Mathematical Society **39** (2002), no. November 3, 355–405.

[21] C. J. Geyer, *Practical markov chain monte carlo*, Statistical Science **7(4)** (1992), 473–511.

[22] J. L. Goffin, *The relaxation method for solving systems of linear inequalities*, Mathematics of Operations Research **5** (1980), no. 3, 388–414.

[23] M. Grötschel, L. Lovász, and A. Schrijver, *Geometric algorithms and combiantorial optimization*, second ed., Springer-Verlag, Berlin, 1994.

[24] B. Grünbaum, *Partitions of mass-distributions and convex bodies by hyperplanes*, Pacific Journal of Mathematics **10** (1960), 1257–1261.

[25] ———, *Measures of symmetry for convex sets*, Proceedings of Symposia in Pure Mathematics **7** (1963), 233–270.

[26] B. Grünbaum, *Convex polytopes*, Wiley, New York, 1967.

[27] P. C. Hammer, *The centroid of a convex body*, Proc. Amer. Math. Soc. **5** (1951), 522–525.

[28] G. Iyengar, *private communication*, (2006).

[29] C. Lemaréchal J. F. Bonnans, J. Ch. Gilbert and C. Sagastizábal, *Numerical optimization: Theoretical and practical aspects*, Universitext, Springer-Verlag, 2003.

[30] F. John, *Extremum problems with inequalities as subsidiary conditions*, in Studies and Essays, Presented to R. Courant on His 60th Birthday, Interscience, New York **30** (1948), 187–204.

[31] A. Kalai and S. Vempala, *Simulating annealing for convex optimization*, to appear on Mathematic of Operations Research (2006).

[32] N. Karmarkar, *A new polynomial-time algorithm for linear programming*, Combinatorica **4** (1984), no. 4, 373–395.

[33] L. G. Khachiyan, *Rounding of polytopes in the real number model of computation*, Mathematics of Operations Research **21** (1996), no. 2, 307–320.

[34] L.G. Khachiyan, *A polynomial algorithm in linear programming*, Soviet Math. Dokl. **20** (1979), no. 1, 191–194.

[35] V. Klee, *The critical set of convex body*, Amer. J. Math. **75** (1953), 178–188.

[36] L. Leindler, *On a certain converse of Hölder's inequality ii*, Acta Sci. Math. Szeged **33** (1972), 217–223.

[37] L. Lovász, *Hit-and-run mixes fast*, Mathematical Programming **86** (1998), 443–461.

[38] L. Lovász and M. Simonovits, *On the randomized complexity of volume and diameter*, Proc. 33rd IEEE FOCS (1992), 482–491.

[39] L. Lovász and S. Vempala, *Hit-and-run is fast and fun*, Microsoft Technical Report.

[40] ———, *Fast algorithms for logconcave functions: Sampling, rounding, integration and optimization*, Submitted. (2006).

[41] ———, *The geometry of logconcave functions and sampling algorithms*, to appear in RSA (2006).

[42] ———, *Hit-and-run from a corner*, SIAM Journal of Computing **35** (2006), no. 4, 985–100.

[43] ———, *Simulated annealing in convex bodies and an $o^*(n^4)$ volume algorithm*, Journal of Computer and System Sciences **72** (2006), 392–417.

[44] H. Minkowski, *Allegemeine lehzätze über konvexe polyeder*, Ges. Abh., Leipzog-Berlin **1** (1911), 103–121.

[45] Y. Nesterov, *private communication*, (2005).

[46] Y. Nesterov and A Nemirovskii, *Interior-point polynomial algorithms in convex programming*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, 1993.

[47] ———, *Central path and Riemannian distances*, Tech. report, CORE Discussion Paper, CORE, Université Catholique de Louvain, Belgium, 2003.

[48] Y. Nesterov, M. Todd, and Y. Ye, *Infeasible-start primal-dual methods and infeasibility detectors for nonlinear programming problems*, Mathematical Programming **84** (1999), 227–267.

[49] G. Pataki, *On the closedness of the linear image of a closed convex cone*, Technical Report, University of North Carolina, TR-02-3 Department of Operations Research 1992.

[50] A. Prékopa, *Logarithmic concave measures and functions*, Acta Sci. Math. Szeged **34** (1973), 335–343.

[51] ———, *On logarithmic concave measures with applications to stochastic programming*, Acta Sci. Math. Szeged **32** (1973), 301–316.

[52] J. Renegar, *A mathematical view of interior-point methods in convex optimization*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, 2001.

[53] R. T. Rockafellar, *Convex analysis*, Princeton University Press, Princeton, New Jersey, 1970.

[54] C. Roos, T. Terlaky, and J.-Ph. Vial, *Theory and algorithms for linear optimization: An interior point approach*, Wiley, 1997.

[55] S. Smale, *Newton's method estimates from data at one point*, The Merging of Disciplines: New Directions in Pure, Applied and Computational Mathematics (1986), 185–196.

[56] M. J. Todd, *Potential-reduction methods in mathematical programming*, Mathematical Programming **76** (1997), no. 1, 3–45.

[57] P. Tseng, *Complexity analysis of a linear complementarity algorithm based on a Lyapunov function*, Mathematical Programming **53** (1992), 297–306.

[58] R. H. Tütüncü, K. C. Toh, and M. J. Todd, *SDPT3 — a MATLAB software package for semidefinite-quadratic-linear programming, version 3.0*, Tech. report, 2001, Available at http://www.math.nus.edu.sg/~mattohkc/sdpt3.html.

[59] S. Vempala, *Geometric random walks: A survey*, Combinatorial and Computational Geometry **52** (2005), 573–612.

[60] S. Vempala, *private communication*, (2005).

[61] ———, *private communication*, (2006).

[62] Y. Ye, *A new complexity result for minimizing a general quadratic function with a sphere constraint*, Recent advances in global optimization (1992), 19–31.

[63] Y. Ye, *On the von neumann economic growth problem*, Mathematics of Operations Research **20** (1995), no. 3, 617 – 633.

[64] Y. Zhang and L. Gao, *On numerical solution of the maximum volume ellipsoid problem*, SIAM Journal on Optimization **14(1)** (2003), 53–76.