Separation Theorems

Akshay Agrawal
akshayka@cs.stanford.edu

January 21, 2019

Abstract

The theory of duality in convex analysis, a subfield of mathematical analysis that studies convex sets and convex functions, rests upon a collection of theorems that give conditions under which two convex sets can be separated by a hyperplane. These theorems are known as *separation theorems*. In this note, we present and prove two important separation theorems. This note is self-contained: we begin by defining convex sets, placing special emphasis on affine sets, hyperplanes, and halfspaces. We then develop basic topological properties of convex sets, introducing the notion of relative interior. Upon establishing these preliminaries, we state and prove necessary and sufficient conditions for proper and strong separation of convex sets. We close by highlighting a few important consequences of the separation theorems. For simplicity, we limit our setting to Euclidean space; however, the approach taken here parallels one employed in the setting of functional analysis. This note is very much modeled after [Roc70].

1 Introduction

Convex analysis is a subfield of mathematical analysis that studies convex sets and convex functions. A convex set is a set that contains as subsets the line segment between any two points belonging to the set; a convex function is any real-valued function whose epigraph is a convex set. Convex analysis has played an outsized role in applied mathematics over the course of the past few decades, due to the fact that the problem of minimizing a convex function admits computationally tractable (often polynomial-time) algorithms [BV04; NN94]; as such, many students of science and engineering feel compelled to study convex analysis. What these students will find is that convex analysis is an elegant mathematical subject worth studying in its own right.

Central to the study of convex analysis is the duality between convex sets and hyperplanes. The fact that a closed convex set is the intersection of all closed halfspaces containing it is of particular importance; a primary purpose of this note is to build up to this fundamental result. To do so, we

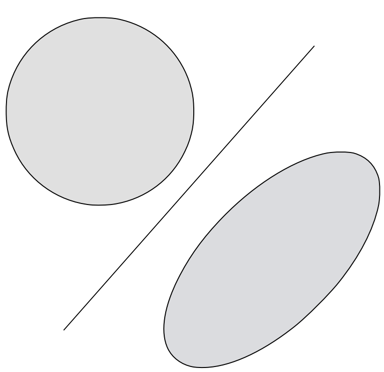

Figure 1: Two convex sets in $\mathbb{R}^2$ and a hyperplane separating them.

will develop two *separation theorems* which give necessary and sufficient conditions under which convex sets can be separated by a hyperplane. We will define separation precisely in §4, but the intuition should be clear: a hyperplane is said to separate two convex sets if one convex set lies on side of it, and the other set lies on the other side (see Fig. 1). For simplicity we limit our setting to Euclidean space, *i.e.*, $\mathbb{R}^n$ equipped with the usual inner product, denoted $\langle \cdot, \cdot \rangle$; however, our approach is topological and is modeled after the one in [Roc70] (see *e.g.* [HUL12] for a geometric approach that exploits the structure of $\mathbb{R}^n$). Readers need only know real analysis and some very basic linear algebra to proceed.

2 Convex sets

We have already stated informally stated the definition of a convex set. For completeness, we give a formal definition below.

Definition 2.1. A set $C \subseteq \mathbb{R}^n$ is said to be *convex* if for any points x and y belonging to C, and for all $\theta \in [0, 1]$,

$$\theta x + (1 - \theta)y \in C.$$

For any two points $x \in \mathbb{R}^n$, $y \in \mathbb{R}^n$, the set

$$\{\, \theta x + (1 - \theta)y \ :\ \theta \in [0, 1] \,\} \tag{2.1}$$

is called the *line segment* between x and y. It becomes obvious why this name is appropriate upon rewriting (2.1) as

$$\{\, y + \theta(x - y) \ :\ \theta \in [0, 1] \,\}.$$

A set containing every *line* through any two points belonging to the set is called an *affine set*.

Definition 2.2. A set $A \subseteq \mathbb{R}^n$ is *affine* if for any points x and y belonging to A, and for all $\theta \in \mathbb{R}$,

$$\theta x + (1 - \theta)y \in A.$$

Clearly, every affine set is also a convex set; also, the sets $\{0\}$ and $\mathbb{R}^n$ are trivially affine sets. An important example of an affine set is the *hyperplane*.

Example 2.3. Let a be a nonzero point in $\mathbb{R}^n$ and let β be a real number. The set

$$H = \{\, x \ :\ \langle a, x \rangle = \beta \,\}$$

is called a *hyperplane*.

The vector a in 2.3 is a *normal* to the hyperplane H, since if $_0 \in H$, then $H = \{\, x \ :\ \langle a, x - x_0 \rangle = 0 \,\}$. Note that a and $-a$ define the same hyperplane, and that scalar multiples of these two points are the only normals to H. For this reason, it is natural to say that every hyperplane has two sides, a notion we make precise in the following example.

Example 2.4. Let a be a nonzero point in $\mathbb{R}^n$ and β a real number. The sets

$$H_1 = \{\, x \ :\ \langle a, x \rangle \leq \beta \,\}, \quad H_2 = \quad \{\, x \ :\ \langle a, x \rangle \geq \beta \,\}$$

are called *closed halfspaces*. If the inequalities above are replaced with strict inequalities, the corresponding sets are called *open halfspaces*. The set $H = \{\, x \ :\ \langle a, x \rangle = \beta \,\}$ is said to be the hyperplane corresponding to the halfspaces H_1 and H_2, and H_1 and H_2 are said to be the closed halfspaces corresponding to H.

Every closed halfspace is also a closed set. To see this, consider a halfspace

$$H = \{\, x \,:\, \langle a, x\rangle \leq \beta \,\}$$

defined by nonzero $a \in \mathbb{R}^n$ and real β. The function f given by $x \mapsto \langle a, x\rangle$ is surjective because it is homogeneous and not identically zero, so the set $(-\infty, \beta]$, which is closed, is in its image. Because f is continuous, the pre-image $f^{-1}((-\infty, \beta]) = H$ is a closed set. Similarly, it can be shown that every open halfspace is an open set.

3 Relative interiors of convex sets

It is often the case that the interiors of convex sets in $\mathbb{R}^n$ are empty, even when the sets have a natural analog of an interior: for example, a disc or a triangle in $\mathbb{R}^3$ have empty interior with respect to the topology of $\mathbb{R}^3$, even though each of these sets each would have non-empty interior if it were projected into a affine set containing it. This motivates the concept of a *relative interior*, which is the the interior of a set relative to the smallest affine set containing it.

Before proceeding, we introduce some notation that will simplify the exposition. We use B to denote the open ball of radius 1 centered at $0 \in \mathbb{R}^n$, and $x + \epsilon B$ to denote the open ball of radius $\epsilon > 0$ centered at x. Similarly, for any set $S \subseteq \mathbb{R}^n$, we write $S + \epsilon B$ to denote the set of points whose distance from S is less than ϵ, *i.e.*,

$$S + \epsilon B = \{\, x \,:\, \|x - y\|_2 < \epsilon \text{ for some } y \in S \,\}.$$

We define scalar multiplication of sets in the natural way: for $\lambda \in \mathbb{R}$, $S \subseteq \mathbb{R}^n$, $\lambda S = \{\, \lambda x \,:\, x \in S \,\}$. Similarly, by $S + T$ we meant the Minkowski sum: $S + T = \{\, x + y \,:\, x \in S, y \in T \,\}$. Additionally, we write $\operatorname{cl} S$ to denote the closure of a set S, and $\operatorname{int} S$ to denote its interior.

Definition 3.1. The *affine hull* of a set $S \subseteq \mathbb{R}^n$ is the set of all affine combinations of points in S, and it is denoted $\operatorname{aff} S$:

$$\operatorname{aff} S = \{\, \theta_1 x_1 + \cdots + \theta_k x_k \,:\, x_1, \ldots, x_k \in C,\ \theta_1 + \cdots + \theta_k = 1 \,\}.$$

Definition 3.2. The *relative interior* of a convex set $C \subseteq \mathbb{R}^n$, which we denote $\operatorname{ri} C$, is the interior of C relative to $\operatorname{aff} C$, *i.e.*,

$$\operatorname{ri} C = \{\, x \in C \,:\, \exists \epsilon > 0,\ (x + \epsilon B) \cap \operatorname{aff} C \subseteq C \,\}$$

If C is open relative to $\operatorname{aff} C$ (equivalently, if $\operatorname{ri} C = C$), we say that C is *relatively open.*

Convex sets are topologically simple, in that their closures and relative interiors obey many algebraic properties. Below, we develop some of these properties that we will put to use in proving the separation theorems in §4. The proofs of these properties are not particularly enlightening, and, as such, we choose to omit a few. Readers who seek completeness should reference [Roc70, §6].

Theorem 3.3. *Let $C \subseteq \mathbb{R}^n$ be a convex set, and let $x \in \operatorname{ri} C$ and $y \in \operatorname{cl} C$. Then for each $\theta \in [0, 1)$,*

$$(1 - \theta)x + \theta y \in \operatorname{ri} C.$$

Proof. We will assume for simplicity that $\operatorname{ri} C = \operatorname{int} C$. In fact, this assumption does not sacrifice generality, due to linear algebraic reasons that are outside the scope of this note (see [Roc70, p. 45]). Fix $\theta \in [0, 1)$. In light of our assumption, it suffices to show that there exists $\epsilon > 0$ such that

$$(1-\theta)x + \theta y + \epsilon B \subseteq C.$$

Because y is a limit point of C, it follows that $y \in C + \epsilon B$ for every $\epsilon > 0$, and in particular that

$$\begin{aligned}(1-\theta)x + \theta y + \epsilon B &\subseteq (1-\theta)x + \theta(C + \epsilon B) + \epsilon B \\ &= (1-\theta)[x + \epsilon(1+\theta)(1-\theta)^{-1}B] + \theta C.\end{aligned}$$

Because $x \in \operatorname{int} C$, we can take ϵ to be so small that

$$(1-\theta)[x + \epsilon(1+\theta)(1-\theta)^{-1}B] + \theta C \subseteq C$$

This proves that

$$(1-\theta)x + \theta y + \epsilon B \subseteq C,$$

i.e., that $x \in \operatorname{int} C = \operatorname{ri} C$. □

One useful way to rephrase Theorem 3.3 is the following: if a line segment in a convex set C has one endpoint in $\operatorname{ri} C$ and another in its boundary (or more generally its closure), then every point in the segment except the boundary endpoint lies in $\operatorname{ri} C$. Many nice topological properties of convex sets are consequences of this fact, including the following one.

Theorem 3.4. *For any convex set $C \subseteq \mathbb{R}^n$, it is always the case that $\operatorname{cl} \operatorname{ri} C = \operatorname{cl} C$ and $\operatorname{ri} \operatorname{cl} C = \operatorname{ri} C$.*

Proof. We will prove the first claim in the theorem, that $\operatorname{cl} \operatorname{ri} C = \operatorname{cl} C$. Of course, $\operatorname{cl} \operatorname{ri} C \subseteq \operatorname{cl} C$, since $\operatorname{ri} C \subseteq C$. For the other direction, consider any $y \in \operatorname{cl} C$, and fix some $x \in \operatorname{ri} C$ (such an x must exist whenever $C \neq \emptyset$; see [Roc70, Theorem 6.2]). The line segment between x and y, excluding y, is a subset of $\operatorname{ri} C$ by Theorem 3.3. Hence y is a limit point of $\operatorname{ri} C$, *i.e.*, $y \in \operatorname{cl} \operatorname{ri} C$, proving that $\operatorname{ri} C \subseteq \operatorname{cl} \operatorname{ri} C$. □

Corollary 3.4.1. *Let C_1 and C_2 be convex subsets of $\mathbb{R}^n$. Then $\operatorname{cl} C_1 = \operatorname{cl} C_2$ if and only if $\operatorname{ri} C_1 = \operatorname{ri} C_2$.*

The following theorem will furnish for us two corollaries that will be useful in §4.

Theorem 3.5. *Let $C_1, C_2, \ldots, C_n$ be convex subsets of $\mathbb{R}^n$ such that $\cap_i \operatorname{ri} C_i \neq \emptyset$. Then*

$$\operatorname{cl} \cap_i^n C_i = \cap_i^n \operatorname{cl} C_i$$

and

$$\operatorname{ri} \cap_i^n C_i = \cap_i^n \operatorname{ri} C_i.$$

Proof. Let x be some point in $\cap_i C_i$. For any $y \in \cap_i \operatorname{cl} C_i$, the half-open line segment $\{\, (1-\theta)x + \theta y : \theta \in [0, 1) \,\}$ is contained in $\cap_i \operatorname{ri} C_i$, by Theorem 3.4, and moreover y is a limit point of this line segment. Hence $\cap_i \operatorname{cl} C_i \subseteq \operatorname{cl} \cap_i \operatorname{ri} C_i$, and clearly $\operatorname{cl} \cap_i \operatorname{ri} C_i \subseteq \operatorname{cl} \cap_i C_i \subseteq \cap_i \operatorname{cl} C_i$. This proves that $\operatorname{cl} \cap_i^n C_i = \cap_i^n \operatorname{cl} C_i$. This chain of inclusions also shows that $\cap_i \operatorname{ri} C_i$ and $\cap_i C_i$ have the same closures; by Corollary 3.4.1, we conclude that these two sets share the same relative interior, and in particular that

$$\operatorname{ri} \cap_i C_i \subseteq \cap_i \operatorname{ri} C_i.$$

Proving the reverse inclusion requires more work, which we omit here. We refer the reader to [Roc70, Theorem 6.5] for a complete proof. □

Corollary 3.5.1. *Let C be a convex set, M an affine which has non-empty intersection with* $\operatorname{ri} C$*. Then* $\operatorname{ri}(M \cap C) = M \cap \operatorname{ri} C$*, and* $\operatorname{cl}(M \cap C) = M \cap \operatorname{cl} C$*.*

Corollary 3.5.2. *Let C_1 and C_2 be convex sets such that $C_2 \subseteq \operatorname{cl} C_1$ but $C_2 \not\subseteq \operatorname{cl} C_1 \setminus \operatorname{ri} C_1$. Then* $\operatorname{ri} C_2 \subseteq \operatorname{ri} C_1$*.*

Finally, we state one more useful fact, without proof [Roc70, Corollary 6.6.2].

Theorem 3.6. *For any convex subsets C_1 and C_2 of $\mathbb{R}^n$,*

$$\operatorname{ri}(C_1 + C_2) = \operatorname{ri} C_1 + \operatorname{ri} C_2,$$

$$\operatorname{cl}(C_1 + C_2) \supseteq \operatorname{cl} C_1 + \operatorname{cl} C_2,$$

4 Separation theorems

We are now very nearly ready to state, and prove, two important separation theorems. Recall that every hyperplane has exactly two halfspaces that correspond to it; we defined this precisely in Example 2.4. For any two convex subsets C_1 and C_2 of $\mathbb{R}^n$, a hyperplane H is said to *separate* C_1 and C_2 if C_1 is contained in one of the halfspaces corresponding to H and C_2 is contained in the other halfspace. A hyperplane separates C_1 and C_2 *properly* if both are not wholly contained in the hyperplane. It separates the sets *strongly* if there exists an $\epsilon > 0$ such that $C_1 + \epsilon B$ is contained in one of the open halfspaces corresponding to H and $C_2 + \epsilon B$ is contained in the other open halfspace.

In this section, we will present two theorems: one giving a necessary and sufficient condition for proper separation, and the other giving a necessary and sufficient condition for strong separation. We begin by presenting a key linear algebraic lemma.

Lemma 4.1. *Let $C \subseteq \mathbb{R}^n$ be a non-empty relatively open convex set, and let $M \subseteq \mathbb{R}^n$ be a non-empty affine set such that $M \cap C = \emptyset$. Then there exists a hyperplane H such that $M \subseteq H$ and one of the two open halfspaces associated with H contains C.*

Proof. If M is a hyperplane, then the result is immediate; if one of the open halfspaces of M did not contain C, then M would have non-empty intersection with C (because C is convex), which would contradict our hypothesis. So assume that M is not a hyperplane, and without loss of generality assume $0 \in M$. We will construct a subspace M' of dimension one higher than that of M which also does not intersect C. Because every hyperplane has dimension one minus the ambient dimension, iterating this procedure finitely many times will yield a hyperplane that does not intersect M.

Because M is not a hyperplane, its orthogonal complement contains a two-dimensional subspace P. Consider the set $C' = P \cap (C - M)$, and note that $C - M \supseteq C$ (since $0 \in M$), but $0 \notin C - M$; hence, $0 \notin C'$. Additionally, C' is a relatively open convex subset of P, by Corollary 3.5.1 and Theorem 3.6. We seek a line $L \subseteq P$ through 0 that does not intersect C', because, for any such line L, $M' = M + L$ will be a subspace of one higher dimension than M not intersecting C. If C' is empty, then evidently such a line exists. Otherwise, if $\operatorname{aff} C$ is a line not containing 0, then we can simply take the line parallel to $\operatorname{aff} C$ through 0; if $\operatorname{aff} C$ is a line containing 0, we can take a line perpindicular to it containing 0. Else, if $\operatorname{aff} C$ is two-dimensional, then the set $K = \bigcup \{ \lambda C' : \lambda > 0 \}$ is a convex cone containing C' but not containing 0. In this case, we can just take L to be one of the two boundary rays of K, extended to a line through 0. $\square$

4.1 A necessary and sufficient condition for proper separation

The following characterization of proper separation is useful.

Lemma 4.2. *Let S and T be non-empty subsets of $\mathbb{R}^n$. There exists a hyperplane separating S and T properly if and only if there exists $a \in \mathbb{R}^n$ such that*

$$\inf\{\, \langle a, x\rangle \,:\, x \in S \,\} \geq \sup\{\, \langle a, x\rangle \,:\, x \in T \,\} \tag{4.1}$$

and

$$\sup\{\, \langle a, x\rangle \,:\, x \in S \,\} > \inf\{\, \langle a, x\rangle \,:\, x \in T \,\}. \tag{4.2}$$

Proof. Assume that $a \in \mathbb{R}^n$ satisfies (4.1) and (4.2), and fix any β such that

$$\inf_{x \in S}\{\langle a, x\rangle\} \geq \beta \geq \sup_{x \in T}\{\langle a, x\rangle\}.$$

Because a satisfies (4.2), $a \neq 0$. So $H = \{\, x \,:\, \langle a, x\rangle = \beta \,\}$ is a hyperplane. By our choice of β, the closed halfspace $\{\, x \,:\, \langle a, x\rangle \geq \beta \,\}$ contains C_1, and the other closed halfspace corresponding to H contains C_2. Furthermore, (4.2) implies that C_1 and C_2 are not both contained in H. Hence H separates C_1 and C_2 properly.

Now assume that C_1 and C_2 are separated properly by some hyperplane parameterized by $a \in \$R^n$ and $\beta \in \mathbb{R}$ such that $C_1 \subseteq \{\, x \,:\, \langle a, x\rangle \geq \beta \,\}$ and C_2 is contained in the other halfspace. Then $\langle a, x\rangle \geq \beta$ for every $x \in C_1$ and $\langle a, x\rangle \leq \beta$ for every $x \in C_2$, so in particular

$$\inf\{\, \langle a, x\rangle \,:\, x \in S \,\} \geq \sup\{\, \langle a, x\rangle \,:\, x \in T \,\}.$$

The fact that the hyperplane separates C_1 and C_2 properly means that for at least one $x \in C_1$, $\langle a, x\rangle > \beta$ and, similarly, for at least one $x \in C_2$, $\langle a, x\rangle < \beta$. This implies (4.2). $\square$

We are now ready to state and prove the first of our two separation theorems, which states that two convex sets can be separated properly if and only if their relative interiors do not meet.

Theorem 4.3. *Let C_1 and C_2 be non-empty convex subsets of $\mathbb{R}^n$. There exists a hyperplane properly separating C_1 and C_2 if and only if $\operatorname{ri} C_1 \cap \operatorname{ri} C_2 = \emptyset$.*

Proof. Let $C = C_1 - C_2$. By Theorem 3.6, $\operatorname{ri} C = \operatorname{ri} C_1 - \operatorname{ri} C_2$, which means that $0 \notin \operatorname{ri} C$ if and only if $\operatorname{ri} C_1 \cap \operatorname{ri} C_2 = \emptyset$. Because $0 \notin \operatorname{ri} C$, the affine set $M = \{0\}$ does not intersect $\operatorname{ri} C$; by Lemma 4.1, there exists a hyperplane containing M such that $\operatorname{ri} C$ is a subset of one of its corresponding open halfspaces. And, because $C \subseteq \operatorname{cl}(\operatorname{ri} C)$ (by Theorem 3.4), the closure of that halfspace contains C. Hence, there exists a (nonzero) $a \in \mathbb{R}^n$ such that

$$0 \leq \inf\{\, \langle a, x\rangle \,:\, x \in C \,\} = \inf\{\, \langle a, x\rangle \,:\, x \in C_1 \,\} - \sup\{\, \langle a, x\rangle \,:\, x \in C_2 \,\},$$

$$0 < \sup\{\, \langle a, x\rangle \,:\, x \in C \,\} = \sup\{\, \langle a, x\rangle \,:\, x \in C_1 \,\} - \inf\{\, \langle a, x\rangle \,:\, x \in C_2 \,\}$$

By Lemma 4.2, this implies that C_1 and C_2 and be separated properly. The above conditions actually imply that $0 \notin \operatorname{ri} C$: speficially, they imply that the halfspace $S = \{\, x \,:\, \langle a, x\rangle \geq 0 \,\}$ contains C such that $\operatorname{int} D = \operatorname{ri} D \not\ni 0$ intersects C. This means that $\operatorname{ri} C \subseteq \operatorname{ri} D$ (by Corollary 3.5.2). $\square$

Theorem 4.3 implies a fact that is very intuitive, at least in the setting of $\mathbb{R}^3$: non-empty disjoint convex sets can be properly separated.

Corollary 4.3.1. *If C_1 and C_2 are non-empty, disjoint convex subsets of $\mathbb{R}^n$, then there exists a hyperplane that properly separates them.*

Proof. If C_1 and C_2 do not meet, then their relative interiors do not meet. Hence by the previous theorem they can be properly separated. $\square$

Note that disjointness is not enough to imply strong separation, for the closures of C_1 and C_2 might meet.

4.2 A necessary and sufficient condition for strong separation

The necessary and sufficient condition for strong separation is quite intuitive: two non-empty convex sets can be separated strongly precisely when the distance between the two sets is positive.

Theorem 4.4. *Let C_1 and C_2 be non-empty convex subsets of $\mathbb{R}^n$. There exists a hyperplane strongly separating C_1 and C_2 if and only if*

$$\inf\{\, \|x_1 - x_2\|_2 \,:\, x_1 \in C_1, c_2 \in C_2 \,\} > 0,$$

i.e., if and only if 0 *is not in* $\mathrm{cl}(C_1 - C_2)$.

Proof. If C_1 and C_2 admit a strong separation, then there exists a positive ϵ such that the open ball $C_1 + \epsilon B$ does not intersect the ball $C_2 + \epsilon B$; this implies that the distance between C_1 and C_2 isp positive. Conversely if $C_1 + \epsilon B$ and $C_2 + \epsilon B$ do not meet, they can be properly separated (by Corrolary 4.3.1). It follows that the sets $C_1 + \epsilon/2B$ and $C_2 + \epsilon/2B$ lie in opposite open halfspaces, and in particular that C_1 and C_2 can be strongly separated. To summarize, we have that C_1 and C_2 can be strongly separated if and only if there exists an $\epsilon > 0$ such that $C_1 + \epsilon B$ does not meet $C_2 + \epsilon B$, *i.e.*,

$$0 \notin (C_1 + \epsilon B) - (C_2 + \epsilon B) = C_1 - C_2 - 2\epsilon B.$$

This in turn implies that $2\epsilon B \cap (C_1 - C_2) = \emptyset$, *i.e.*, $0 \notin \mathrm{cl}(C_1 - C_2)$.

$\square$

5 Applications

Separation theorems have many interesting consequences, both in convex analysis itself and in applications such as convex optimization and economics. In this section, we present two important consequences within convex analysis.

5.1 Duality between convex sets and halfspaces

At the heart of duality in convex analysis is the fact that a closed convex set is the intersection of all halfspaces containing it. Or, put another way, every closed convex set can be expressed as the set of solutions to a system of (non-strict) linear inequalities $\langle a_i, x\rangle \le \beta_i$ for i in an index set.

Theorem 5.1. *A closed convex set is the intersection of all the halfspaces that contain it.*

Proof. The theorem is trivial whenever $C = 0$ or $C = \mathbb{R}^n$, so let us assume that neither of these situations are the case. For any $a \notin C$, the convex sets $\{a\}$ and C admit strong separation by Theorem 4.4, *i.e.*, there exists a hyperplane separating them strongly. One of the corresponding closed halfspaces contains C but not a. Hence, the intersection of all closed halfspaces containing C is exactly equal to C. $\square$

5.2 Supporting hyperplanes

Let S be a subset of $\mathbb{R}^n$, and let x_0 be a point in its boundary. If $a \neq 0$ satisfies $\langle a, x\rangle \leq \langle a, x_0\rangle$ for all $x \in S$, then the hyperplane $\{\, x \,:\, \langle a, x\rangle = \langle a, x_0\rangle \,\}$ is called a *supporting hyperplane* to S at x_0. Equivalently, this hyperplane separates $\{x_0\}$ and S and contains $\{x_0\}$. We refer to the halfspace $\{\, x \,:\, \langle a, (x - x_0)\rangle \leq 0, x \in S \,\}$ corresponding to a supporting hyperplane $\{\, x \,:\, \langle a, x\rangle = \langle a, x_0\rangle \,\}$ of S as a *supporting halfspace* of S.

Theorem 5.2. *Let C be a non-empty convex subset of $\mathbb{R}^n$, and let x be any point in the boundary of C. Then there exists a supporting hyperplane to C at x.*

Proof. If the interior of C is non-empty, then the sets $\operatorname{ri}\{x\} = \{x\}$ and $\operatorname{int} C$ are disjoint. By Corollary 4.3.1, these sets can be properly separated, implying that $\{x\}$ and C can be separated (not necessarily properly). If on the other hand $\operatorname{int} C = \emptyset$, then C is a subset of an affine set with dimension strictly less than n. Hence it is entirely contained in some hyperplane, and this hyperplane trivially separates $\{x\}$ and C. $\square$

Theorems 5.1 and 5.2 immediately imply the following corollary.

Corollary 5.2.1. *A closed convex set C is equal to the intersection of its supporting halfspaces.*

Example 5.3. A *cone* is a set K with the property that if $x \in \mathcal{K}$, then θx is also in $\mathcal{K}$ for all $\theta \geq 0$. The *dual cone* $\mathcal{K}^*$ of $\mathcal{K}$ is the set $\{\, y \,:\, \langle y, x\rangle \geq 0,\, x \in \mathcal{K} \,\}$. Geometrically, $y \in \mathcal{K}$ if and only if $-y$ supports $\mathcal{K}$ at the origin, *i.e.*, the dual cone $\mathcal{K}^*$ is the set of supporting hyperplanes of $\mathcal{K}$ (up to a change in sign). Hence, the dual cone derives its name from the duality between halfspaces and convex sets.

Convex cones and their duals arise frequently in convex optimization; they play a role analogous to subspaces in linear algebra. While convex optimization is in general NP-hard, Nesterov and Nemirovski [NN94] showed that the problem of minimizing linear functions over cross-products of certain convex cones can be solved in polynomial time (up to an additive ϵ error).

Example 5.4. The *epigraph* of a real-valued function f is the set $\{\, (x, t) \,:\, f(x) \leq t \,\}$. A *convex function* is a real-valued function whose epigraph is a convex set; if the epigraph of a convex function is closed, the function is said to be *closed*. By theorem 5.1, the epigraph of a convex function is equal to the intersection of all the halfspaces containing it, and by corollary 5.2.1, the epigraph of a closed convex function can be described as the intersection of its supporting halfspaces.

References

[BV04] Boyd, S. and Vandenberghe, L. *Convex Optimization*. New York, NY, USA: Cambridge University Press, 2004.

[HUL12] Hiriart-Urruty, J.-B. and Lemaréchal, C. *Fundamentals of convex analysis*. Springer Science & Business Media, 2012.

[NN94] Nesterov, Y. and Nemirovski, A. *Interior-point Polynomial Algorithms in Convex Programming*. Society for Industrial and Applied Mathematics, 1994.

[Roc70] Rockafellar, R. T. *Convex Analysis*. Princeton University Press, 1970.