EE 381V: Large Scale Optimization **Fall 2012**

Lecture 2 — September 3

Lecturer: Caramanis & Sanghavi *Scribe: Hongbo Si, Qiaoyang Ye*

2.1 Overview of the last Lecture

The focus of the last lecture was to give some basic **key definitions**: convex sets; convex functions; and the basic setting for convex optimization. The **key ideas** introduced is the simple notion underlying the power of convexity: convexity $\Rightarrow$ local = global. Finally, the last lecture used these basic definitions to provide our **result**: a characterization of the optimal solution to an unconstrained optimization of a smooth convex objective function: x^* is an optimal solution iff $\nabla f(x^*) = 0$.

2.2 Overview of this Lecture

One of the main results of this lecture, is to extend the characterization above to the constrained optimization case. We see that this characterization is precisely what the (first order) KKT conditions express. Along the way, we define two important sets: the Tangent Cone and its polar, the Normal Cone, of a convex set at a point. Finally, we begin our discussion of separation, by talking about projection onto convex sets.

2.3 Convex Modeling

Before we proceed to the main results and definitions of this section, we continue introducing examples of utilizing convex optimization to model problems. We try to use the following examples to show how to understand, find and exploit a convex model for certain practical problems, and to practice "optimization as a way of thinking about a problem."

2.3.1 Max Flow and Min Cut

Consider a directed graph $G = (V, E)$, shown in Figure 2.1, where V denotes the set of the graph's vertices, and E denotes the set of edges. The capacity of link e is denoted by c_e. Given a source vertex, denoted as v_s, and a destination vertex, denoted as v_t, one may ask about the max throughput, or max flow from source to destination; alternatively, one might consider finding the bottleneck on the graph limiting throughput from the source to think sink. It turns out that the tightest bottleneck is equivalent to computing the maximum flow. We will understand this better in a future class when we discuss duality.

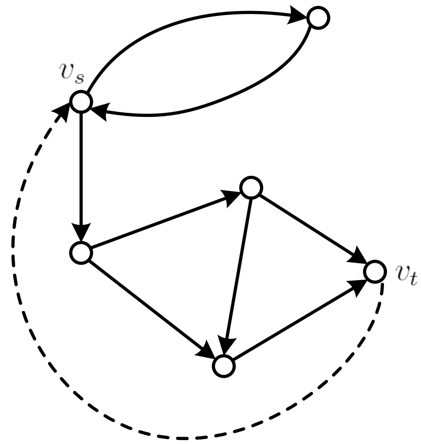

Figure 2.1. Max flow problem on a directed graph.

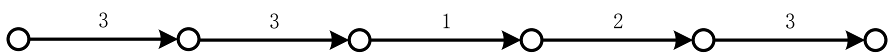

Figure 2.2. Simplest example showing max flow and bottleneck.

Let's begin with the simplest case, as shown in Figure 2.2. Obviously, the bottleneck is the middle edge with capacity 1, and it determines the max flow at the same time. This result is not limited to such simple graph. Indeed, *in general, on a directed graph, max flow = min cut.*

We will turn to the proof of this result later in this course, and now only concern ourselves with how to formulate the max flow problem as a convex optimization. We need to define the *decision variables*, the *constraints*, and the *objective function.*

- The decision variables are the amounts of flow that passes across edge $e \in E$, denoted by f_e.

- The constraints of the flow maximization problem are straightforward: each flow is nonnegative and cannot exceed the capacity. Moreover, flow-conservation dictates that the amount of flow into any node equals the amount of flow out of that node. Using $\mathrm{In}(v)$ to denote the set of edges flowing into node v, and similarly $\mathrm{Out}(v)$ the set of edges flowing out of a node v, we can write the constraints as follows:

$$0 \leq f_e \leq C_e, \ \forall e \in E, \tag{2.1}$$

$$\sum_{e \in \mathrm{In}(v)} f_e = \sum_{e \in \mathrm{Out}(v)} f_e, \ \forall v \in V, \tag{2.2}$$

$$\sum_{e \in \mathrm{Out}(v_s)} f_e = \sum_{e \in \mathrm{In}(v_t)} f_e, \tag{2.3}$$

where (2.1) is the natural condition of flows; (2.2) is the balance between inflow and outflow for a certain vertex; (2.3) is the flow balance over the whole graph, or also can be considered as adding a surplus edge from destination to source, as shown in Figure 2.1.

Note that all equality conditions are linear, therefore, we have the following optimization problem:

$$\begin{aligned} \max \ & \sum_{e \in \mathrm{Out}(v_s)} f_e \\ \text{s.t.} \ & A\boldsymbol{f} = 0 \\ & 0 \le f_e \le C_e, \ \forall e \in E. \end{aligned} \tag{2.4}$$

Here, we have used $\boldsymbol{f}$ to denote the vector of all flows. The matrix A describes the linear equations representing conservation of flow, i.e. (2.2) and (2.3). Thus, the columns of A correspond to the edges of the graph, and the rows to a node. An entry in A is '+1' if the edge is incident and an "in-node" for node v, it is '-1' if it is incident and an out-node, and it is '0' otherwise. The objective function is linear and hence convex.

2.3.2 Optimal Inequalities in Probability

Now consider a totally different problem. Assume X is a real valued random variable. Given the following moment constraints:

$$\mu_i = \mathbb{E}[X^i], \quad i = 1, 2, 3, 4, 5. \tag{2.5}$$

The question is how to find an lower bound and upper bound for $\mathbb{P}\{X \in [2, 3]\}$?

There may be many methods for finding the bounds. Here, we give an optimization view:

- The decision variable is the density of X, denoted by $f_X(x)$. Note that unlike the previous problem, now we have infinitely many decision variables.
- Combining with the given constraints, an optimization problem is formulated as follow:

$$\begin{aligned} \min \ & \int_2^3 f_X(x)dx \\ \text{s.t.} \quad & f_X(x) \ge 0, \quad \forall\, x \in \mathbb{R}, \\ & \int f_X(x)dx = 1, \\ & \int x^i f_X(x)dx = \mu_i, \quad i = 1, 2, 3, 4, 5. \end{aligned} \tag{2.6}$$

This is an optimization problem with infinitely many variables, but other than the non negativity constraints, only 6 other constraints. This is called a *semi-infinite optimization problem.* Facing this optimization problem, and in particular the infinite variables, one may ask when does this formulation make sense, and in particular, does it accurately describe the initial problem. Next, we have to also ask when can we solve it via reasonable computational techniques. Answers come from "duality" and duality comes from local optimality, as we will understand better in future lectures. In fact, it turns out that we can solve the univariate problem exactly and easily, and we can approximate multivariate problems to arbitrary precision, by solving convex optimization problems.

2.4 Characterizing Optimal Solutions

We now turn to obtaining a characterization of the optimal solution to a constrained convex optimization problem, with a smooth convex objective function. We will begin with some further concepts related to convex sets.

For almost all that we do we need some notion of distance. This is needed, in order to talk about closed and open sets, convergence, limits, etc.

Definition 1. *A* metric *is a mapping $d(\cdot,\cdot) : \mathcal{X} \times \mathcal{X} \to \mathbb{R}_+$ that satisfies: $d(x, y) \geq 0$ with equality iff $x = y$; $d(x, y) = d(y, x)$, and $d(x, z) + d(z, y) \geq d(x, y)$.*

Using this, we can give a simple definition of open and closed sets.

Definition 2. (Open Sets) *A set $C \subseteq \mathbb{R}^n$ is called* **open** *if $\forall x \in C$, $\exists \varepsilon > 0$ such that $B_\varepsilon(x) \subseteq C$.*

Example 1. *The unit sphere without boundary in $\mathbb{R}^n$ is an open set: $B_1(0) = \{x : \|x\| < 1\}$.*

Example 2. *The set of all $n \times n$ symmetric matrices with strictly positive eigenvalues, S^n_+, is an open set.*

Definition 3. (Closed Sets) *A set $C \subseteq \mathbb{R}^n$ is called* **closed** *if $x_n \in C, x_n \to \bar{x} \;\Rightarrow\; \bar{x} \in C$.*

It is useful to have basic familiarity with these definitions, and to be able to manipulate them. As an example, we prove the following.

Proposition 1. *A set is closed iff its complement is open.*

Proof: We first prove that if C is closed, then C^c must be open. To this end, assume C^c is not open. Then there exists $x \in C^c$ such that for any $\epsilon > 0$, $B_\epsilon(x) \nsubseteq C^c$. Let $\{\epsilon_n\}$ be a sequence such that $\epsilon \to 0$, then we can find a sequence of points $\{x_n\}$ in $B_\epsilon(x) \cap C$ such that $x_n \to x$. Note that $x_n \in C$ and C is closed, so $x \in C$, which contradicts our initial assumption $x \in C^c$.

Conversely, suppose C is open. We show C^c is closed. Assume C^c is not closed. Then there exist a sequence $x_n \in C^c$ such that $x_n \to x$ with $x \in C$. Since C is open, there exists $\epsilon > 0$ such that $B_\epsilon(x) \subseteq C$. Thus the convergence assumed is impossible. □

Definition 4. (Interior) *A point $x \in C$ is an* **interior** *point of C, denoted by $x \in \mathrm{Int}C$, if $\exists \varepsilon > 0$ such that $B_\varepsilon(x) \subseteq C$.*

The distinction between an interior point and a boundary point, is important for optimization, since no constraints are active at an interior point. The notion of *relative interior* is particularly important in convex analysis and optimization, because many feasible sets are often expressed using an intersection with an affine space that is not full dimensional. We do not go into the details here, but simply give the basic idea through a picture. For a point to be in the interior of a set, the set must contain a small ball around that point. Relative interior replaces this with the requirement that the set contain the intersection of a ball around the point and the affine hull of the set itself. Figure 2.3 illustrates this idea. The shaded circular shape lies in three dimensions, but its affine hull is only two dimensional, and hence it can have no interior. However, the notion of relative interior recovers the intuitive notion of "interior" of the shape, recognizing that it is really a two-dimensional object.

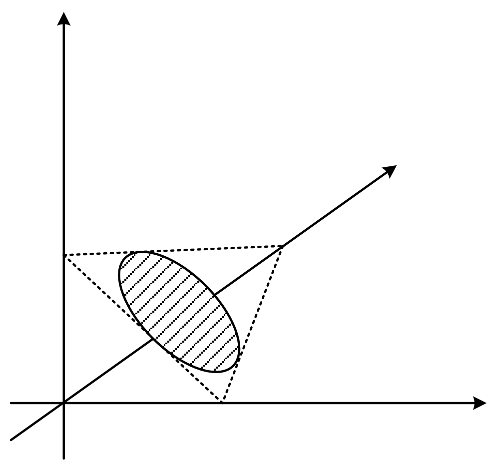

Figure 2.3. Interior and relative interior.

Recall the definition of convex hull. Figure 2.4 shows the convex hull of 8 points in 2D space. Note that some points are in the interior of the convex hull, while some others are at the "corner". We define these corner points as extreme points.

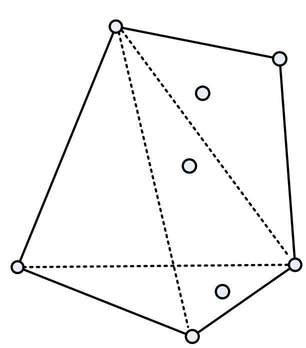

Figure 2.4. Extreme points.

Definition 5. (Extreme Points) *Let $C \in \mathbb{R}^n$ be a convex set. A point $x \in C$ is called an* **extreme point** *if $\nexists x_1, x_2 \in C, x_1 \neq x_2$ such that $x = (x_1 + x_2)/2$.*

Example 3. *The extreme points of the unit ball, $B_1(0) = \{x \in \mathbb{R}^n \,|\, \|x\|_2 \leq 1\}$, are the points of unit magnitude. For a more interesting example, consider the set of positive semidefinite (symmetric) matrices with spectral norm at most one. The set of extreme points is the set of rank one unit norm matrices.*

As shown in Figure 2.4, any point in the interior always lies in a triangle with vertices as extreme points. That is, any point in the convex hull of the extreme points can always be expressed as a convex hull of at most three extreme points. This is in general true for convex sets in n dimensions.

Theorem 2.1. (Carathéodory's Theorem) *Given a convex set $C \in \mathbb{R}^n$, every point $x \in C$ can be described as a convex combination of at most $(n+1)$ extreme points of C.*

Proof: Denote the convex hull of set C as conv(C). By the definition of conv(C), for $\forall x \in C, \exists \lambda_1, \cdots, \lambda_k \in \mathbb{R}_+$, and $x_1, \cdots, x_k \in C$ such that

$$x = \sum_{i=1}^{k} \lambda_i x_i, \tag{2.7}$$

where $\sum_{i=1}^k \lambda_i = 1$, and $\lambda_i \geq 0, \forall 1 \leq i \leq k$. Note that $k > n+1$ and without loss of generality $\lambda_i > 0$, otherwise there is nothing to prove.

Consider the $(n+1)$ homogeneous linear equations in $k > (n+1)$ variables $\{\mu_1, \cdots, \mu_k\}$:

$$\underbrace{\sum_{i=1}^{k} \mu_i x_i = 0,}_{n \text{ equations}} \quad \underbrace{\sum_{i=1}^{k} \mu_i = 0.}_{1 \text{ equations}} \tag{2.8}$$

Since $k > n+1$, there exists a solution $\{\mu_i\}$ to these equations, other than the trivial all-zeros solution. In particular, since $\sum_i \mu_i = 0$, there must be at least one strictly positive μ_i. Let $\alpha \in \mathbb{R}$ be

$$\alpha \triangleq \min_{1 \leq i \leq k} \left\{ \frac{\lambda_i}{\mu_i} : \mu_i > 0 \right\} = \frac{\lambda_j}{\mu_j}, \tag{2.9}$$

where $j = \arg\min_i \left\{ \frac{\lambda_i}{\mu_i} : \mu_i > 0 \right\}$ and then define $\hat{\lambda}_i = \lambda_i - \alpha\mu_i$. Then we have $\hat{\lambda}_i \geq 0$ with one $\hat{\lambda}_j = 0$. Thus, we get

$$\sum_{i=1}^{k} \hat{\lambda}_i x_i = \sum_{i=1}^{k} (\lambda_i - \alpha\mu_i) x_i = \sum_{i=1}^{k} \lambda_i x_i + 0 = x, \tag{2.10}$$

$$\sum_{i=1}^{k} \hat{\lambda}_i = \sum_{i=1}^{k} (\lambda_i - \alpha\mu_i) = \sum_{i=1}^{k} \lambda_i = 1, \tag{2.11}$$

and

$$\hat{\lambda}_j = \lambda_j - \alpha\mu_j = 0. \tag{2.12}$$

Therefore, x is described as an affine combination of at most $k-1$ points in C. The above procedure can be repeated until one obtains a representation for x in terms of a convex combination of at most $n+1$ points in C. $\square$

Remark 1. *Many interesting properties of convex sets related to discrete and geometry follow from similar proof ideas. For the interested reader, we refer to Helly and Radon's theorems.*

Last lecture we gave the definition of a convex cone. We repeat the definition here.

Definition 6. (Convex Cone) *A set $K \subseteq \mathbb{R}^n$ is called a* **convex cone**, *if $x_1, x_2 \in K$ implies that $\lambda_1 x_1 + \lambda_2 x_2 \in K$, $\forall \lambda_1, \lambda_2 \geq 0$.*

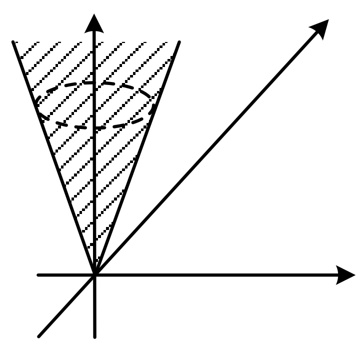

Figure 2.5. Example of convex cone in 3D.

Definition 7. (Polar Cone) *Let $K \subseteq \mathbb{R}^n$ be a cone. Then the* **polar cone** *of K, denoted by K°, is described by*

$$K^\circ = \{x \in \mathbb{R}^n : \langle x, v \rangle \leq 0, \ \forall v \in K\}. \tag{2.13}$$

Example 4. $(\mathbb{R}^2_+)^\circ = \mathbb{R}^2_-$.

Example 5. $(S^n_+)^\circ = -S^n_+$.

Exercise 1. *If K is a closed convex cone, then $K^{\circ\circ} = K$.*

Definition 8. (Feasible Directions) *Let $C \subseteq \mathbb{R}^n$ be a nonempty set, and let $x \in C$. Then the set of all* **feasible directions** *of C at x, denoted by $F_C(x)$, is defined as follows*

$$F_C(x) = \{d : \exists \varepsilon > 0, \text{such that } x + \varepsilon d \in C\}. \tag{2.14}$$

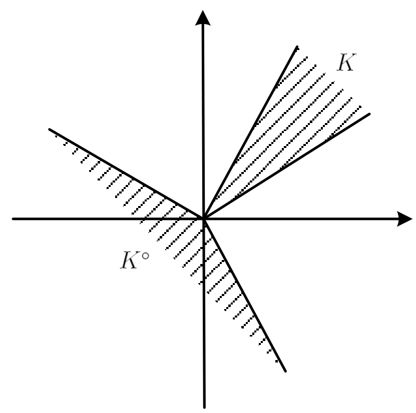

Figure 2.6. Example of cone and its polar cone in 2D.

Definition 9. (Tangent Cone) *Let $C \subseteq \mathbb{R}^n$ be a nonempty set, and let $x \in C$. Then the* **tangent cone** *of C at x, denoted by $T_C(x)$, is defined as follows*

$$T_C(x) = closure(F_C(x)). \tag{2.15}$$

Definition 10. (Normal Cone) *Let $C \subseteq \mathbb{R}^n$ be a nonempty, convex set, and let $x \in C$. Then the* **normal cone** *of C at x, denoted by $N_C(x)$, is defined as follows*

$$N_C(x) = \{s : \langle s, y - x \rangle \leq 0, \ \forall y \in C\}. \tag{2.16}$$

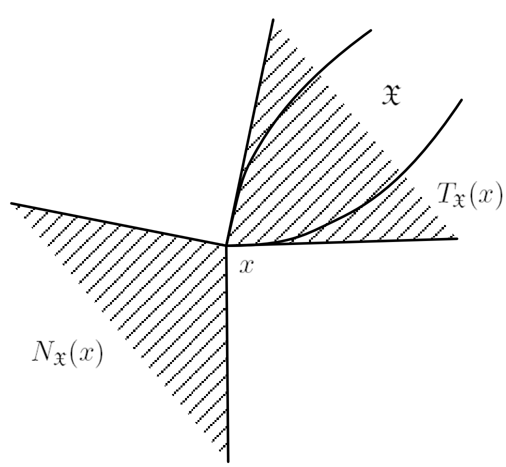

Figure 2.7. Example of tangent cone and normal cone.

Theorem 2.2. *Let $C \in \mathbb{R}^n$ be a nonempty, convex set, and let $x \in C$. Then the normal cone of C at x is the polar cone of the tangent cone of C at x. That is,*

$$N_C(x) = (T_C(x))^\circ. \tag{2.17}$$

Proof: Let $s \in N_C(x)$. Then for any $d \in f_C(x)$, there exists $\epsilon > 0$ such that $x + \epsilon d \in C$. Hence,

$$\langle s, d \rangle = \frac{1}{\epsilon} \langle s, x + \epsilon d - x \rangle \leq 0.$$

For any $\hat{d} \in T_C(x)$, there exists a sequence d_n such that $d_n \in F_C(x)$ and $d_n \to \hat{d}$. As $\langle s, d_n \rangle \leq 0$, we have $\langle s, \hat{d} \rangle \leq 0$, which means $s \in (T_C(x))^\circ$, i.e. $N_C(x) \subseteq (T_C(x))^\circ$.

On the other hand, let $s \in (T_C(x))^\circ$. Then $\langle s, d \rangle \leq 0$ for $\forall\ d \in T_C(x)$. For any $y \in C$, because C is convex, there exists $\hat{d} \in T_C(x)$ and $\alpha > 0$ such that $y = x + \alpha \hat{d}$. Hence, we have that

$$\langle s, y - x \rangle = \alpha \langle s, \hat{d} \rangle \leq 0.$$

Thus, $s \in N_c(x)$, i.e. $(T_C(x))^\circ \subseteq N_C(x)$. □

Having these concepts in hand, let's turn to how to relate these to the characterization of an optimal solution for an optimization problem. Consider a general convex optimization given by

$$\begin{aligned} \min\ & f(x) \\ \text{s.t.}\ & x \in \mathfrak{X}. \end{aligned} \tag{2.18}$$

By convexity, local and global optimality are equivalent. Thus, intuitively, x^* is optimal iff no descent directions are feasible, where descent direction v at x^* means $\langle \nabla f(x^*), v \rangle \leq 0$. Thus, if x^* is optimal, we have

$$\begin{aligned} & \langle \nabla f(x^*), x - x^* \rangle \geq 0,\ \forall x \in \mathfrak{X} \\ \Longleftrightarrow\ & \langle -\nabla f(x^*), x - x^* \rangle \leq 0,\ \forall x \in \mathfrak{X} \\ \Longleftrightarrow\ & -\nabla f(x^*) \in (T_{\mathfrak{X}}(x^*))^\circ \\ \Longleftrightarrow\ & -\nabla f(x^*) \in N_{\mathfrak{X}}(x^*) \\ \Longleftrightarrow\ & 0 \in \nabla f(x^*) + N_{\mathfrak{X}}(x^*) \end{aligned} \tag{2.19}$$

This result is a general optimal condition, we will see how it is related to our earlier result for unconstrained optimization. We also show through example, that this is precisely the geometric condition that the first-order KKT conditions are attempting to express.

- For an unconstrained optimization problem, min : $f(x)$, we can consider $\mathfrak{X} = \mathbb{R}^n$. Note that for any $x \in \mathfrak{X}$, $N_{\mathfrak{X}}(x) = \{0\}$. Therefore our optimality condition becomes $0 \in \nabla f(x^*) + \{0\}$, or simply $\nabla f(x^*) = 0$, which coincides with previous result.

- For constrained optimization problem, consider the following example, illustrated in Fig. 2.8:

$$\begin{aligned} \min\ & f(x) \\ \text{s.t.}\ & g_1(x) \leq 0 \\ & g_2(x) \leq 0. \end{aligned} \tag{2.20}$$

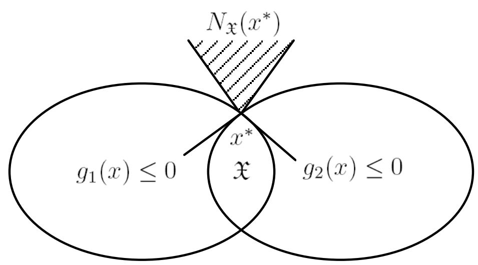

Figure 2.8. Example of optimal condition.

Note that in this case, $\mathfrak{X} = \{x : g_1(x) \le 0,\ g_2(x) \le 0\}$, and the normal cone at the point x^* is simply given by $N_{\mathfrak{X}}(x^*) = \text{cone}(\nabla g_1(x^*), \nabla g_2(x^*))$. So the optimal condition becomes

$$0 = \nabla f(x^*) + \lambda_1 \nabla g_1(x^*) + \lambda_2 \nabla g_2(x^*). \tag{2.21}$$

This is exactly the first order KKT condition for optimality. In order to derive the KKT condition, we relied on the fact that the normal cone at x^* could be expressed using the gradients of the constraint functions. As we will see later, while this is commonly the case, it is not always so. This is why some problems, even though convex, may fail to have Lagrange multipliers.

2.5 Projection and Separation

We now discuss some properties of projection and separation. Using these definitions, in the next lecture we will show that any closed convex set is equal to the intersection of all half-spaces that contain it. This means that if a point x does not belong to a convex set C, then there is a hyperplane with proves this, i.e., there is a half space that contains C but does not contain the point. Note that this is not the case for non-convex sets.

The first result that will be quite useful for this, is obtaining a *variational characterization* of projection. This is quite familiar for projection onto an affine manifold. Indeed, we have:

Proposition 2. *Let V be an affine set, and let x be some point not in V. Then the point $v^* \in V$ is the solution to*

$$\min_{v \in V} : \ \|x - v\|,$$

if and only if

$$\langle x - v^*, v - v^* \rangle = 0, \quad \forall v \in V.$$

Proof: This follows immediately from the Pythagorean Theorem: For any point $v \in V$, we have

$$\|v - x\|_2^2 = \|v - v^*\|_2^2 + \|v^* - x\|_2^2.$$

□

Something similar holds when we replace the affine set V with a general (closed) convex set C. First we define the notion of projection onto a convex set.

Definition 11. (Projection) *Let $C \subseteq \mathbb{R}^n$ be a closed convex set. For $x \in \mathbb{R}^n$, the* **projection** *of x on C, denoted by $Proj_C(x)$, is defined as follows*

$$Proj_C(x) \triangleq \arg\min_{y \in C} \|x - y\|. \tag{2.22}$$

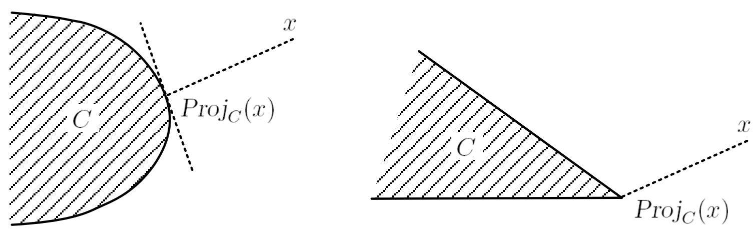

Figure 2.9. Examples of projection.

Proposition 3. (Uniqueness of Projection) *Let $C \subseteq \mathbb{R}^n$ be a closed convex set. For $x \in \mathbb{R}^n$, the projection $Proj_C(x)$ is* **unique**.

Proof: Let x_1, x_2 be projections of x on C, i.e., $\mathrm{Proj}_C(x) = x_1$ and $\mathrm{Proj}_C(x) = x_2$, with $x_1 \neq x_2$. We will use the following equation

$$\frac{\|a + b\|^2}{2} = \|a\|^2 + \|b\|^2 - \frac{\|a - b\|^2}{2} \tag{2.23}$$

to obtain a contradiction. Setting $a = x_1 - x$ and $b = x_2 - x$, we have

$$\frac{1}{2}\|x_1 - x + x_2 - x\|^2 = \|x_1 - x\|^2 + \|x_2 - x\|^2 - \frac{1}{2}\|x_1 - x_2\|^2. \tag{2.24}$$

This implies

$$\|\frac{x_1 - x + x_2 - x}{2}\|^2 < \frac{1}{2}\|x_1 - x\|^2 + \frac{1}{2}\|x_2 - x\|^2. \tag{2.25}$$

Using that $\|x_1 - x\| = \|x_2 - x\|$, we have

$$\|\frac{x_1 + x_2}{2} - x\|^2 < \|x_1 - x\|^2. \tag{2.26}$$

Since C is convex, the point $\frac{x_1+x_2}{2}$ lies in C and is not equal to x_1 or x_2. So the above inequality contradicts with x_1 and x_2 are both projections. □

Definition 12. (Strictly Convex) *A function f is said to be* **strictly convex** *if the following holds*

$$f(\lambda x_1 + (1 - \lambda)x_2) > \lambda f(x_1) + (1 - \lambda)f(x_2), \forall x_1 \neq x_2, \lambda \in (0, 1). \tag{2.27}$$

Proposition 4. (Unique Solution) *If function f is strictly convex, and $C \subseteq \mathbb{R}^n$ is a closed convex set, then the optimization problem*

$$\begin{aligned} \min \ & f(x) \\ s.t. \ & x \in C \end{aligned} \tag{2.28}$$

has a **unique** *solution if it has any solutions. It is guaranteed to have a solution as long as the convex function f has compact sub-level sets (or as long as the intersection of C with one sub-level set of f is compact).*

Proof: The last assertion follows by Weierstrass's theorem. Next, we show that if the problem has a solution, it is unique: assume to the contrary that x_1 and x_2 are both optimal solutions to (2.28). Then we have $\frac{x_1+x_2}{2} \in C$, due the convexity of set C. By strictly convexity of f, we have

$$f(\frac{x_1+x_2}{2}) < \frac{1}{2}f(x_1) + \frac{1}{2}f(x_2) = f(x_1) = f(x_2). \tag{2.29}$$

This contradicts that x_1 and x_2 are optimal solutions. $\square$

Therefore, this provides a more general proof of uniqess, since the objective function defining projection is indeed strictly convex.

In the next lecture we show that the solution to the projection optimization problem must satisfy a variational property very similar to the affine case.