Handout 1: Introduction, Convex Analysis

Instructor: Anthony Man–Cho So January 8, 2008

1 Overview

In this course we will consider a class of **mathematical programming** problems that can be expressed in the form:

$$(P) \qquad \begin{array}{ll} \text{minimize} & f(x) \\ \text{subject to} & x \in X \end{array}$$

Here, $f : \mathbb{R}^n \to \mathbb{R}$ is called the **objective function**, and $X \subset \mathbb{R}^n$ is called the **feasible region**. Thus, $x = (x_1, \ldots, x_n)$ is an n–dimensional vector, and we shall agree that it is represented in *column* form. The entries $x_1, \ldots, x_n$ are called the **decision variables** of (P). As the above formulation suggests, we are interested in a **global minimizer** of (P), which is defined as a point $x^* \in X$ such that $f(x^*) \le f(x)$ for all $x \in X$. We call $f(x^*)$ the **optimum value** of (P). A related notion is that of a **local minimizer**, which is defined as a point $x' \in X$ such that for some $\epsilon > 0$, we have $f(x') \le f(x)$ for all $x \in X \cap B(x', \epsilon)$. Here,

$$B(x', \epsilon) = \left\{ x \in \mathbb{R}^n : \|x - x'\|_2 \le \epsilon \right\}$$

is the **Euclidean ball** of radius $\epsilon > 0$ centered at x' (recall that for $x \in \mathbb{R}^n$, the 2–**norm** of x is defined as $\|x\|_2^2 = \sum_{i=1}^n x_i^2 \equiv x^T x$). Note that a global minimizer is automatically a local minimizer, but the converse is not necessarily true. In this course we shall devote a substantial amount of time to characterize the minimizers of (P) and to study how the structures of f and X affect our ability to solve (P). Before we do that, however, let us observe that problem (P) is quite general. For example, when $X = \mathbb{R}^n$, we have an **unconstrained optimization** problem; when X is *discrete* (e.g. $X = \{-1, +1\}^n \subset \mathbb{R}^n$), we have a **discrete optimization** problem. Other important classes of optimization problems include:

- **Linear Programming (LP) Problems:** Here, f is a **linear function**, i.e. a function of the form:

  $$f(x) = c_1 x_1 + c_2 x_2 + \cdots + c_n x_n \equiv c^T x$$

  with $c = (c_1, \ldots, c_n) \in \mathbb{R}^n$; and X is a set defined by **linear inequalities**, i.e. it takes the form:

  $$X = \left\{ x \in \mathbb{R}^n : a_i^T x \le b_i \quad \text{for } i = 1, \ldots, m \right\} \qquad (1)$$

  with $a_1, \ldots, a_m \in \mathbb{R}^n$ and $b_1, \ldots, b_m \in \mathbb{R}$. In more compact notation, we may write a linear programming problem as follows:

  $$\begin{array}{ll} \text{minimize} & c^T x \\ \text{subject to} & Ax \le b \end{array}$$

  where A is an $m \times n$ matrix whose i–th row is a_i^T, and $b = (b_1, \ldots, b_m)$ is an m–dimensional column vector (for any two vectors $u, v \in \mathbb{R}^n$, the inequality $u \le v$ means $u_i \le v_i$ for $i = 1, \ldots, n$). As we shall see later in the course, LP problems can be solved very efficiently.

- **Quadratic Programming (QP) Problems:** Here, X is as in (1), and f is a quadratic function, i.e. a function of the form:

$$f(x) = \sum_{i=1}^{n} \sum_{j=1}^{n} Q_{ij} x_i x_j \equiv x^T Q x$$

where $Q = [Q_{ij}]$ is an $n \times n$ symmetric matrix.

- **Semidefinite Programming (SDP) Problems:** Given an $n \times n$ symmetric matrix Q, we say that Q is **positive semidefinite** (denoted by $Q \succeq \mathbf{0}$) if $x^T Q x \geq 0$ for all $x \in \mathbb{R}^n$. Now, let $A_1, \ldots, A_m$ and C be given $n \times n$ symmetric matrices, and let $b_1, \ldots, b_m \in \mathbb{R}$ be given. Consider the optimization problem:

$$\begin{array}{lll} \text{minimize} & C \bullet Z \equiv \displaystyle\sum_{i=1}^{n} \sum_{j=1}^{n} C_{ij} Z_{ij} & \\ \text{subject to} & A_i \bullet Z = b_i & \text{for } i = 1, \ldots, m \\ & Z \succeq \mathbf{0} & \end{array} \tag{2}$$

Problem (2) is a so–called *semidefinite programming (SDP)* problem. Here, the feasible region can be expressed as:

$$X = \left\{ Z \in \mathbb{R}^{n(n+1)/2} : A_i \bullet Z = b_i \text{ for } i = 1, \ldots, m; \; Z \succeq \mathbf{0} \right\}$$

Similar to LP problems, SDP problems can also be solved efficiently.

The aforementioned classes of problems capture a wide range of applications. However, in order to convert a particular application into a problem of the form (P), we need to first identify the data and decision variables and then formulate the objective function and constraints. Let us now illustrate this process via some examples.

1.1 An Air Traffic Control Problem

Suppose that n airplanes are trying to land at the Hong Kong International Airport. Airplane i will arrive at the airport within the time interval $[a_i, b_i]$, where $i = 1, \ldots, n$. For simplicity, we shall assume that the airplanes arrive in the order $1, 2, \ldots, n$. Due to safety concerns, the control tower of the airport would like to maximize the so–called *shortest metering time*, i.e. the minimum over all inter–arrival times between two consecutive airplanes. How then should the airport assign the arrival time of each airplane?

Here, the decision variables are the arrival times of the airplanes, which we denote by $t_1, \ldots, t_n$. Then, we have the following optimization problem:

$$\begin{array}{lll} \text{maximize} & \min_{1 \leq j \leq n-1} (t_{j+1} - t_j) & \\ \text{subject to} & a_i \leq t_i \leq b_i & \text{for } i = 1, \ldots, n \\ & t_i \leq t_{i+1} & \text{for } i = 1, \ldots, n-1 \end{array} \tag{3}$$

It is not immediately clear that (3) can be formulated as an LP, but it can be done as follows. Let z be a new decision variable. Then, we may rewrite (3) as:

$$\begin{array}{lll} \text{maximize} & z & \\ \text{subject to} & t_{i+1} - t_i \geq z & \text{for } i = 1, \ldots, n-1 \\ & a_i \leq t_i \leq b_i & \text{for } i = 1, \ldots, n \\ & t_i \leq t_{i+1} & \text{for } i = 1, \ldots, n-1 \end{array}$$

which is an LP. We should point out that the above reformulation works only because we are *maximizing* instead of minimizing the quantity $\min_{1 \leq j \leq n-1}(t_{j+1} - t_j)$. In particular, the following problems:

$$\begin{array}{lll} \text{minimize} & \min_{1 \leq j \leq n-1}(t_{j+1} - t_j) & \\ \text{subject to} & a_i \leq t_i \leq b_i & \text{for } i = 1, \ldots, n \\ & t_i \leq t_{i+1} & \text{for } i = 1, \ldots, n-1 \end{array} \tag{4}$$

and

$$\begin{array}{lll} \text{minimize} & z & \\ \text{subject to} & t_{i+1} - t_i \geq z & \text{for } i = 1, \ldots, n-1 \\ & a_i \leq t_i \leq b_i & \text{for } i = 1, \ldots, n \\ & t_i \leq t_{i+1} & \text{for } i = 1, \ldots, n-1 \end{array} \tag{5}$$

are *not* equivalent, since the optimum value of (4) is finite, while the optimum value of (5) is $-\infty$.

1.2 A Data Fitting Problem

The previous example shows that sometimes one may be able to convert an optimization problem into an LP via some transformations. Here is another illustration of such possibility. Suppose that we are given m data pairs (a_i, b_i), where $a_i \in \mathbb{R}^n$ and $b_i \in \mathbb{R}$ for $i = 1, \ldots, m$, with $m \geq n$. We suspect that these pairs fall on a line, and so our goal is to determine the "slope" $x \in \mathbb{R}^n$ of a line that best fits the data. In other words, we would like to find an $x \in \mathbb{R}^n$ that minimizes some sort of error measure. One popular measure is the 1–*norm* of the *residual errors*, which is defined as:

$$\Delta_1 = \sum_{i=1}^{m} \left| b_i - a_i^T x \right| = \|b - Ax\|_1$$

where A is the $m \times n$ matrix whose i–th row is a_i^T. Thus, our optimization problem is simply:

$$\text{minimize} \quad \sum_{i=1}^{m} \left| b_i - a_i^T x \right| \tag{6}$$

Here, the objective function is nonlinear. However, we can turn problem (6) into an LP as follows. We first introduce m new decision variables $z_1, \ldots, z_m \in \mathbb{R}$. Then, it is not hard to see that (6) is equivalent to the following LP:

$$\begin{array}{lll} \text{minimize} & \displaystyle\sum_{i=1}^{m} z_i & \\ \text{subject to} & b_i - a_i^T x \leq z_i & \text{for } i = 1, \ldots, m \\ & -b_i + a_i^T x \leq z_i & \text{for } i = 1, \ldots, m \end{array}$$

Now, what if we want to minimize the 2–norm of the residual errors? In other words, we would like to solve the following problem:

$$\text{minimize} \quad \Delta_2 = \|b - Ax\|_2^2 = \sum_{i=1}^{m} \left(b_i - a_i^T x\right)^2 \tag{7}$$

It turns out that this is a particularly simple QP. In fact, since (7) is an unconstrained optimization problem with a *differentiable* objective function, we can solve it using calculus techniques. Indeed, the optimal solution to (7) is given by:

$$x^* = \left(A^T A\right)^{-1} A^T b$$

whenever A has full column rank (whence $A^T A$ is invertible).

At this point let us reflect a bit on the above examples. Intuitively, a linear problem (say, an LP) should be easier than a nonlinear problem, and a non–differentiable problem should be harder than a differentiable one. However, the above examples show that these need not be the case. Indeed, even though the 2–norm problem (7) is an QP, its optimal solution has a nice characterization, while the corresponding 1–norm problem (6) does not have such a feature. On the other hand, even though the objective function in (6) is non–differentiable, the problem can still be solved easily via LP.

From the above discussion, it is natural to ask what makes an optimization problem difficult. While it is hard to give an answer to such question without over–generalizing, let us at least identify a possible source of difficulty. Consider problem (7). Suppose that we impose the additional constraint that $x_i \in \{-1, +1\}$ for $i = 1, \ldots, n$ (such problem actually arises in engineering contexts). In other words, consider the following problem:

$$\begin{array}{ll} \text{minimize} & \|b - Ax\|_2^2 \\ \text{subject to} & x_i^2 = 1 \qquad \text{for } i = 1, \ldots, n \end{array} \tag{8}$$

Then, the resulting problem becomes computationally intractable, and one can show that an efficient algorithm for solving it is unlikely to exist. What distinguishes the seemingly very similar problems (7) and (8) is that problem (7) is a so–called **convex optimization problem**, while problem (8) is not. In the next section, we shall define the notion of convexity and study it in more detail.

2 Elements of Convex Analysis

2.1 Basic Definitions and Properties

We begin with some definitions.

Definition 1 *Let $S \subset \mathbb{R}^n$ be a set. We say that:*

1. *S is* **affine** *if $\alpha x + (1-\alpha)y \in S$ whenever $x, y \in S$ and $\alpha \in \mathbb{R}$;*

2. *S is* **convex** *if $\alpha x + (1-\alpha)y \in S$ whenever $x, y \in S$ and $\alpha \in [0, 1]$.*

Given $x, y \in \mathbb{R}^n$ and $\alpha \in \mathbb{R}$, the vector $z = \alpha x + (1-\alpha)y$ is called an **affine combination** of x and y. If $\alpha \in [0,1]$, then z is called a **convex combination** of x and y.

Geometrically, when x and y are distinct points in $\mathbb{R}^n$, the set:

$$L = \{z \in \mathbb{R}^n : z = \alpha x + (1-\alpha)y,\ \alpha \in \mathbb{R}\}$$

of all affine combinations of x and y is simply the *line* determined by x and y; and the set:

$$S = \{z \in \mathbb{R}^n : z = \alpha x + (1-\alpha)y,\ \alpha \in [0,1]\}$$

is the *line segment* between x and y. By convention, the empty set $\emptyset$ *is* convex. Note that the intersection of an arbitrary family of convex sets is again convex.

Here are some sets in Euclidean space whose convexity can be easily established by first principles:

Example 1 (Some Examples of Convex Sets)

1. **Non–Negative Orthant**: $\mathbb{R}^n_+ = \{x \in \mathbb{R}^n : x \geq \mathbf{0}\}$
2. **Euclidean Ball**: $B(\bar{x}, r) = \{x \in \mathbb{R}^n : \|x - \bar{x}\|_2 \leq r\}$
3. **Ellipsoid**: $E(\bar{x}, Q, r) = \left\{x \in \mathbb{R}^n : (x-\bar{x})^T Q (x - \bar{x}) \leq r^2\right\}$, *where Q is an $n \times n$ symmetric, positive definite matrix (i.e. $x^T Q x > 0$ for all $x \in \mathbb{R}^n \backslash \{\mathbf{0}\}$)*
4. **Simplex**: $\Delta = \{\sum_{i=0}^n \alpha_i x_i : \sum_{i=0}^n \alpha_i = 1,\ \alpha_i \geq 0 \text{ for } i = 0, 1, \ldots, n\}$, *where $x_0, x_1, \ldots, x_n$ are vectors in $\mathbb{R}^n$ such that the vectors $x_1 - x_0, x_2 - x_0, \ldots, x_n - x_0$ are linearly independent (equivalently, the vectors $x_0, x_1, \ldots, x_n$ are* affinely *independent)*
5. **Positive Semidefinite Cone**: $\mathcal{S}^n_+ = \{X \in \mathbb{R}^{n \times n} : X \text{ symmetric and positive semidefinite}\}$
6. **Convex Cone**: *A set $K \subset \mathbb{R}^n$ is called a* **cone** *if $\{\alpha x : \alpha > 0\} \subset K$ whenever $x \in K$. If K is also convex, then K is called a convex cone.*

We remark that one can generalize the notion of affine (resp. convex) combination of two points to any finite number of points. For instance, an affine combination of the points $x_1, \ldots, x_k \in \mathbb{R}^n$ is a point $z = \sum_{i=1}^k \alpha_i x_i$, where $\sum_{i=1}^k \alpha_i = 1$. This allows us to define the following notions:

Definition 2 *Let $S \subset \mathbb{R}^n$ be a non–empty set.*

1. *S is an* **affine manifold** *in $\mathbb{R}^n$ if it contains all its affine combinations. Equivalently, S is an affine manifold if it is the translation of some vector space V. Clearly, an intersection of affine manifolds is still an affine manifold.*
2. *The* **affine hull** *of S, denoted by aff(S), is the intersection of all affine manifolds containing S. In other words, aff(S) is the smallest affine manifold containing S.*
3. *The* **dimension** *of S, denoted by dim(S), is the dimension of the affine hull of S.*

Similarly, a convex combination of the points $x_1, \ldots, x_k \in \mathbb{R}^n$ is a point $z = \sum_{i=1}^k \alpha_i x_i$, where $\sum_{i=1}^k \alpha_i = 1$ *and* $\alpha_1, \ldots, \alpha_k \geq 0$. We may then define the following:

Definition 3 *The* **convex hull** *of a (not necessarily convex) set $S \subset \mathbb{R}^n$, denoted by conv(S), is the intersection of all convex sets containing S. In other words, the convex hull of S is the smallest convex set that contains S.*

Note that if S is convex, then $\text{conv}(S) = S$. The above definition can be viewed as an *exterior* characterization of the convex hull of a set. We also have the following *interior* characterization, whose proof is left to the reader as an exercise:

Proposition 1 *A set $S \subset \mathbb{R}^n$ is convex iff it contains every (finite) convex combination of points in S.*

Given Proposition 1, the following should come as no surprise:

Theorem 1 *Let $S \subset \mathbb{R}^n$ be an arbitrary set. Then, conv(S) is the set of all (finite) convex combinations of points in S.*

Now, suppose that S is convex. Then, by Theorem 1, any $x \in S$ can be represented as $x = \sum_{i=1}^k \alpha_i x_i$, where $x_1, \ldots, x_k \in S$, $\sum_{i=1}^k \alpha_i = 1$ and $\alpha_1, \ldots, \alpha_k \geq 0$. However, there is no a priori bound on the number k of points needed. The following theorem of Carathéodory remedies this situation:

Theorem 2 (Carathéodory's Theorem) *Let $S \subset \mathbb{R}^n$ be an arbitrary set. Then, any $x \in$ conv(S) can be represented as a convex combination of at most $n+1$ points in S.*

Proof Consider an arbitrary convex combination $x = \sum_{i=1}^k \alpha_i x_i$, with $x_1, \ldots, x_k \in S$ and $k \geq n+2$. The plan is to show that one of the coefficients α_i can be set to 0 without changing x. To begin, observe that since $k \geq n+2$, the vectors $x_2 - x_1, x_3 - x_1, \ldots, x_k - x_1$ must be *linearly dependent* in $\mathbb{R}^n$. In particular, there exist $\beta_1, \ldots, \beta_k \in \mathbb{R}$, not all zero, such that $\sum_{i=1}^k \beta_i x_i = \mathbf{0}$ and $\sum_{i=1}^k \beta_i = 0$. For $i = 1, \ldots, k$, define:

$$\alpha_i' = \alpha_i - t^* \beta_i \qquad \text{where } t^* = \min_{j:\beta_j > 0} \frac{\alpha_j}{\beta_j} = \max\left\{ t \geq 0 : \alpha_i - t\beta_i \geq 0 \text{ for } i = 1, \ldots, k \right\}$$

Note that $t^* < \infty$, since there exists at least one index j such that $\beta_j > 0$. Now, it is straightforward to verify that $x = \sum_{i=1}^k \alpha_i' x_i$, $\sum_{i=1}^k \alpha_i' = 1$, and $\alpha_1', \ldots, \alpha_k' \geq 0$, and that $|\{i : \alpha_i' > 0\}| \leq k-1$. Now, this process can be repeated until there are only at most $n+1$ non–zero coefficients. This completes the proof. □

We end this sub–section with some topological concepts. For an arbitrary set $S \subset \mathbb{R}^n$, its **interior** is defined to be the set:

$$\text{int}(S) = \{ x \in S : B(x, \epsilon) \subset S \text{ for some } \epsilon > 0 \}$$

The notion of interior is intimately related to the space in which the set S lies. For instance, consider the set $S = [0,1]$. When considered as a set in $\mathbb{R}$, then $\text{int}(S) = (0,1)$. However, if we consider S as a set in $\mathbb{R}^2$, then $\text{int}(S) = \emptyset$, since no 2–dimensional ball of positive radius is contained in S. Such ambiguity motivates the notion of **relative interior**. Specifically, a point $x \in S$ belongs to the relative interior of S (denoted by rel $\text{int}(S)$) if there exists an $\epsilon > 0$ such that $B(x, \epsilon) \cap \text{aff}(S) \subset S$. The advantage of this definition is the following:

Theorem 3 *Let $S \subset \mathbb{R}^n$ be a convex set with $\dim(S) \geq 1$. Then, S has a non–empty relative interior.*

Proof Let $k = \dim(S)$. Then, S contains $k+1$ affinely independent points $x_0, x_1, \ldots, x_k$. They generate the simplex $\Delta = \text{conv}(\{x_0, \ldots, x_k\})$. Clearly, we have $\Delta \subset S$. Moreover, since $\dim(\Delta) = k$, we have $\text{aff}(\Delta) = \text{aff}(S)$. Thus, it suffices to show that Δ has a non–empty relative interior. Towards that end, let

$$\bar{x} = \frac{1}{k+1} \sum_{i=0}^{k} x_i$$

Clearly, we have $\bar{x} \in \text{aff}(\Delta) = \text{aff}(S)$. Now, define $V^i = \text{aff}(\{x_0, x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_k\})$, and let $\epsilon_i = \min_{x \in V^i} \|x - \bar{x}\|_2$. Then, we have $\epsilon_i > 0$ for $i = 0, 1, \ldots, k$. Upon setting $\epsilon = \min_{0 \leq i \leq m} \epsilon_i$, we conclude that $B(\bar{x}, \epsilon) \cap \text{aff}(S) \subset S$, as required. $\square$

3 Further Reading

Convex analysis is a rich subject with many deep and beautiful results. For more details, one can consult the general references listed on the course website and/or the following books:

[1] A. Barvinok. *A Course in Convexity*, volume 54 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, Rhode Island, 2002.

[2] A. Brøndsted. *An Introduction to Convex Polytopes*, volume 90 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1983.

[3] J.-B. Hiriart-Urruty and C. Lemaréchal. *Fundamentals of Convex Analysis*. Grundlehren Text Editions. Springer–Verlag, Berlin/Heidelberg, 2001.

[4] A. Ruszczyński. *Nonlinear Optimization*. Princeton University Press, Princeton, New Jersey, 2006.

[5] G. M. Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer–Verlag, New York, revised first edition, 1995.