

1. THE LANGUAGE AND NOTATION OF NONLINEAR OPTIMIZATION

In finite dimensional optimization we are interested in locating solutions to the problem

$$\begin{aligned} \mathcal{P} : \quad & \underset{x \in X}{\text{minimize}} && f_0(x) \\ & \text{subject to} && x \in \Omega. \end{aligned}$$

where X is the variable space (or decision space), $f_0 : X \rightarrow \mathbb{R} \cup \{\pm\infty\}$ is called the objective function, and the set $\Omega \subset X$ is called the constraint region. The techniques that one employs in the study of \mathcal{P} are determined by the nature of the space X , the function f_0 , and the set Ω . The basic problem categories are as follows:

- (1) Variable Type:
 - (a) continuous variable: $X = \mathbb{R}^n$
 - (b) discrete variable: $X = \mathbb{Z}^n$
 - (c) mixed variable: $X = \mathbb{R}^s \times \mathbb{Z}^t$
- (2) Constraint Type:
 - (a) unconstrained: $\Omega = X$
 - (b) constrained: $\Omega \neq X$
- (3) Problem Type:
 - (a) Convex Programming:
 f_0 is a convex function and Ω is a convex set.

Definition 1.1. *The set $\Omega \subset \mathbb{R}^n$ is said to be convex if for every $x, y \in \Omega$ one has $[x, y] \subset \Omega$ where $[x, y]$ denotes the line segment connecting x and y :*

$$[x, y] = \{\lambda x + (1 - \lambda)y : 0 \leq \lambda \leq 1\}.$$

Definition 1.2. *The function $f : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\pm\infty\}$ is said to be convex if the set $\text{epi}(f) = \{(x, \mu) : f(x) \leq \mu\}$ is a convex set in \mathbb{R}^{n+1} . In particular,*

$$f(\lambda x + (1 - \lambda)y) \leq \lambda f(x) + (1 - \lambda)f(y)$$

for all $0 \leq \lambda \leq 1$ and points x, y for which not both $f(x)$ and $f(y)$ are infinite.

- (b) Linear Programming:

The minimization or maximization of a linear functional subject to a finite number of linear inequality and/or equality constraints. $f_0(x) := c^T x$ for some $c \in \mathbb{R}^n$ and

$$\Omega := \left\{ x \left| \begin{array}{ll} a_i^T x \leq b_i & i = 1, \dots, s \\ a_i^T x = b_i & i = s + 1, \dots, m \end{array} \right. \right\}.$$

Linear programming is a special case of convex programming. In this case the constraint region Ω is called a polyhedral convex set. Polyhedra have a very special geometric structure.

(c) Quadratic Programming:

The minimization or maximization of a quadratic objective functions over a convex polyhedron:

$$f_0(x) = \frac{1}{2}x^T Qx + b^T x + \alpha$$

EXERCISE: Show that $\nabla f_0(x) = \frac{1}{2}(Q + Q^T)x + b$ and $\nabla^2 f_0(x) = \frac{1}{2}(Q + Q^T)$.

FACT: f_0 is convex if and only if Q is positive semi-definite.

(d) Mini-Max:

$$f_0(x) = \max\{f_i(x) : i = 1, \dots, s\}.$$

(e) Nonlinear Programming

$$\Omega := \{x \in X : f_i(x) \leq 0, i = 1, \dots, s, f_i(x) = 0, i = j + 1, \dots, m\}$$

- (i) Differentiable: f_i is smooth $i = 0, \dots, m$
- (ii) Nonsmooth: at least one f_i is not smooth
- (iii) Semi-Infinite: $m = +\infty$.

(f) Box Constraints:

$$\begin{aligned} \Omega &:= \{x \in \mathbb{R}^n : l_i \leq x_i \leq u_i, i = 1, \dots, n\} \\ l_i &\in \mathbb{R} \cup \{-\infty\}, \quad u_i \in \mathbb{R} \cup \{+\infty\}, \quad l_i \leq u_i \end{aligned}$$

(g) Parameter Identification:

(i) Data Fitting:

Given data points $\{(x_i, y_i)\}_{i=1}^m \subset \mathbb{R}^t \times \mathbb{R}^s$, find the function $f(x) := m(p, x)$ from a parametrized class of functions

$$\mathcal{M} := \{m(p, \cdot) : p \in Y \subset \mathbb{R}^n\}$$

that “best” fits the data.

(A) Polynomial least-squares: $t = 1 = s$. $m(p, x) = p_0 + p_1x + \dots + p_nx^n$.

$\mathcal{M} = \mathcal{P}_n =$ set of polynomials of degree less than or equal to n .

Minimize the sum of squares

$$f_0(P) = \sum_{i=1}^m (m(p, x_i) - y_i)^2$$

over all choices of $p \in \mathbb{R}^{n+1}$.

Approach:

$$\begin{bmatrix} p_0 + p_1x_1 + p_2x_1^2 + \dots + p_nx_1^n \\ p_0 + p_1x_2 + p_2x_2^2 + \dots + p_nx_2^n \\ \vdots \\ p_0 + p_1x_m + p_2x_m^2 + \dots + p_nx_m^n \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_m \end{bmatrix}$$

$$\text{Vandermonde Matrix} \left\{ \underbrace{\begin{bmatrix} 1 & x_1 & x_1^2 & \cdots & x_1^n \\ 1 & x_2 & x_2^2 & \cdots & x_2^n \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_m & x_m^2 & \cdots & x_m^n \end{bmatrix}}_A \begin{bmatrix} p_0 \\ p_1 \\ \vdots \\ p_n \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_m \end{bmatrix} \right. \quad p = y$$

$$f_0(p) = \|Ap - y\|_2^2$$

$$\text{where } \|z\|_2 = \left(\sum_{i=1}^m z_i^2 \right)^{\frac{1}{2}}$$

EXERCISE: Show that $\nabla f_0(p) = A^T(Ap - y)$ and $\nabla^2 f_0(p) = A^T A$.

The focus of this course is the development of numerical methods for solving the general nonlinear programming problem under the assumption that all of the underlying functions are smooth. Nonetheless, on occasion we will also need to consider nondifferentiable functions. Our first objective is to come to some agreement about what we mean by a “solution” to \mathcal{P} .

Definition 1.3.

- i. A point $\bar{x} \in \Omega$ is said to be a **global solution** to the problem

$$\min\{f_0(x) : x \in \Omega\}$$

if $f_0(\bar{x}) \leq f_0(x)$ for all $x \in \Omega$. If $f_0(\bar{x}) < f_0(x)$ for all $x \in \Omega$, then \bar{x} is said to be a **strict global solution**.

- ii A point $\bar{x} \in \Omega$ is said to be a **local solution** to the problem $\min\{f_0(x) : x \in \Omega\}$ if there is an $\epsilon > 0$ such that

$$f_0(\bar{x}) \leq f_0(x) \quad \text{for all } x \in \Omega \text{ satisfying} \\ \|\bar{x} - x\| \leq \epsilon.$$

If $f_0(\bar{x}) < f_0(x)$ for all $x \in \Omega$ with $\|x - \bar{x}\| \leq \epsilon$, then \bar{x} is called a **strict local solution**. The solution \bar{x} is said to be **isolated** if \bar{x} is the only local solution in the set

$$\{x \in \Omega : \|x - \bar{x}\| \leq \epsilon\}.$$

EXAMPLES:

- (1) Global: (a) $f_0(x) = x, \quad \Omega = [0, 1]$ } strict
 (b) $f_0(x) = x^2, \quad \Omega = \mathbb{R}$ }
- (2) local: (a) $f_0(x) = x(x+1)(x-1)$ (no global)
 (b) $f_0(x) = (x-1)^2(x+1)^2$ (not strictly global)

- (3) strict global but not isolated

$$f_0(x) = x^2, \Omega := \{x \in [-1, 1] : f_1(x) = 0\}$$

$$f_1(x) = \begin{cases} x^2 \sin \frac{1}{x}, & x \neq 0 \\ 0, & x = 0 \end{cases}$$

The above definitions establish what is meant by local and global solutions to \mathcal{P} . However, since the definition requires us to test f_0 at potentially infinitely many points in order to verify optimality, the definition in itself is not a constructive test of optimality. From the algorithmic point of view we require a constructive test for optimality that can be used to indicate if a point is either optimal or nearly optimal. Such a test is referred to as an “optimality condition”. Any point that satisfies an optimality condition is said to be a critical point, or stationary point, for \mathcal{P} .

There are many types of optimality conditions. We will concentrate on those involving differential information. Recall the following result from the calculus.

Theorem 1.1. *If f_0 is differentiable at \bar{x} and \bar{x} is a local solution to the problem $\min\{f_0(x) : x \in \mathbb{R}^n\}$, then $\nabla f(\bar{x}) = 0$.*

We will recall the proof of this result shortly. But for the moment observe that the condition

$$\nabla f_0(x) = 0$$

is an optimality condition for the unconstrained problem $\min\{f_0(x) : x \in \mathbb{R}^n\}$. This condition is clearly not a sufficient condition for optimality (consider $f_0(x) = x^3$), but it is necessary. For this reason, we refer to it as a *first-order necessary condition for optimality*.

Note that the condition $\nabla f_0(x) = 0$ is testable. Simply evaluate ∇f_0 at \bar{x} . Moreover, it provides a ready means for testing “near” optimality, by considering the magnitude of the vector $\nabla f_0(\bar{x})$. Thus, for example, a reasonable stopping criteria for an algorithm would be

$$\|\nabla f(\bar{x})\| \leq \epsilon.$$

Hopefully, this would yield a finitely terminating procedure.

In what follows optimality conditions play a key role in both the design of our algorithms and our tests for termination. In general, algorithms are designed to locate points that satisfy some testable, or constructive, optimality condition and then terminate when the procedure “nearly” satisfies the condition. For this reason, our first-order of business is to derive workable optimality conditions. Since these conditions will be built on differential information, we begin by reviewing some facts from multi-variable calculus.