

TMA947 / MMG621 — Nonlinear optimization

Lecture 3 — Introduction to optimality conditions

Emil Gustavsson, Zuzana Nedělková

October 20, 2017


Local and global optimality

We consider an optimization problem which is that to

$$\text{minimize } f(\mathbf{x}), \quad (1a)$$

$$\text{subject to } \mathbf{x} \in S, \quad (1b)$$

where $S \subseteq \mathbb{R}^n$ is a nonempty set, and $f : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ is a given function.



When $n = 1$, we know that an optimal solution will be found at

- boundary points of S ,
- stationary points, that is, where $f'(x) = 0$,
- discontinuities of f or f' .

Definition (global minimum). $\mathbf{x}^* \in S$ is a global minimum of f over S if

$$f(\mathbf{x}^*) \leq f(\mathbf{x}), \quad \mathbf{x} \in S$$

The next important definition is the one of a local minimum.

Definition (local minimum). $\mathbf{x}^* \in S$ is a local minimum of f over S if

$$\exists \varepsilon > 0 \text{ such that } f(\mathbf{x}^*) \leq f(\mathbf{x}), \quad \mathbf{x} \in S \cap B_\varepsilon(\mathbf{x}^*),$$

where $B_\varepsilon(\mathbf{x}^*) := \{ \mathbf{y} \in \mathbb{R}^n \mid \|\mathbf{y} - \mathbf{x}^*\| < \varepsilon \}$ is the Euclidean ball with radius ε centered at \mathbf{x}^* .

We denote $\mathbf{x}^* \in S$ as a *strict local minimum* of f over S if $f(\mathbf{x}^*) < f(\mathbf{x})$ holds above for $\mathbf{x} \neq \mathbf{x}^*$.

One of the most important theorems in optimization follows.

Theorem (Fundamental Theorem of global optimality). Consider the problem (1), where S is a convex set and f is convex on S . Then, every local minimum of f over S is also a global minimum

Proof. See Theorem 4.3 in the book. □

Existence of optimal solutions

First some basic notations.

- We say that a set $S \subseteq \mathbb{R}^n$ is *open* if for every $\mathbf{x} \in S$ there exists an $\varepsilon > 0$ such that $B_\varepsilon(\mathbf{x}) := \{ \mathbf{y} \in \mathbb{R}^n \mid \|\mathbf{y} - \mathbf{x}\| < \varepsilon \} \subset S$.
- We say that a set $S \subseteq \mathbb{R}^n$ is *closed* if $\mathbb{R}^n \setminus S$ is open.
- A *limit point* of a set $S \subseteq \mathbb{R}^n$ is a point \mathbf{x} such that there exists a sequence $\{\mathbf{x}_k\}_{k=1}^\infty \subset S$ fulfilling $\mathbf{x}_k \rightarrow \mathbf{x}$.
- We can then define a closed set as a set which contains all its limit points.
- We say that a set $S \subseteq \mathbb{R}^n$ is *bounded* if there exists a constant $C > 0$ such that $\|\mathbf{x}\| \leq C$ for all $\mathbf{x} \in S$.
- If a set is both closed and bounded, we call it *compact*.

Two important definitions needed to formulate Weierstrass' Theorem are the following.

Definition (weakly coercive function). A function f is said to be weakly coercive with respect to the set S if either S is bounded or

$$\lim_{\substack{\|\mathbf{x}\| \rightarrow \infty \\ \mathbf{x} \in S}} f(\mathbf{x}) = \infty$$

Definition (lower semi-continuity). A function f is said to be lower semi-continuous at \mathbf{x} if the value $f(\mathbf{x})$ is less than or equal to every limit of f as $\mathbf{x}_k \rightarrow \mathbf{x}$

In other words, f is lower semi-continuous at $\mathbf{x} \in S$ if

$$\mathbf{x}_k \rightarrow \mathbf{x} \quad \implies \quad f(\mathbf{x}) \leq \liminf_{k \rightarrow \infty} f(\mathbf{x}_k)$$




Figure 1: A lower semi-continuous function in one variable

Now we can formulate Weierstrass' Theorem which guarantees the existence of optimal solutions to an optimization problem as long as a few assumptions are satisfied.

Theorem (Weierstrass' Theorem). Consider the problem (1), where S is a nonempty and closed set and f is lower semi-continuous on S . If f is weakly coercive with respect to S , then there exists a nonempty, closed and bounded (thus compact) set of optimal solutions to the problem (1).

Proof. See Theorem 4.6 in the book. □

One way to remember the assumptions in Weierstrass' Theorem is to imagine what can go wrong, i.e., when does a problem not have an optimal solution. One example of an optimization problem where the solution set is empty is when $f(x) = 1/x$ and $S = [1, \infty)$.

Optimality conditions when $S = \mathbb{R}^n$

When $S = \mathbb{R}^n$, i.e., the problem is an unconstrained optimization problem, then the following theorem holds.

Theorem (necessary condition for optimality, C^1). If $f \in C^1$ on \mathbb{R}^n , then

$$\mathbf{x}^* \text{ is a local minimum of } f \text{ on } \mathbb{R}^n \quad \implies \quad \nabla f(\mathbf{x}^*) = \mathbf{0}$$

Proof. See Theorem 4.13. □

Note that $\nabla f(\mathbf{x}) = \left(\frac{\partial f(\mathbf{x})}{\partial x_1}, \dots, \frac{\partial f(\mathbf{x})}{\partial x_n} \right)^T$. The opposite of the theorem is, however, not true. Take $f(x) = x^3$ and $x = 0$. We can strengthen the theorem by assuming that f is also in C^2 .

Theorem (necessary condition for optimality, C^2). If $f \in C^2$ on \mathbb{R}^n , then

$$\mathbf{x}^* \text{ is a local minimum of } f \text{ on } \mathbb{R}^n \quad \implies \quad \begin{cases} \nabla f(\mathbf{x}^*) = \mathbf{0} \\ \nabla^2 f(\mathbf{x}^*) \succeq 0 \end{cases}$$

Proof. See Theorem 4.16. □

Remember that for a matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, the notion $\mathbf{A} \succeq 0$ (\mathbf{A} positive semidefinite) means that $\mathbf{x}^T \mathbf{A} \mathbf{x} \geq 0$, for all $\mathbf{x} \in \mathbb{R}^n$. Now once again, the opposite direction in the theorem is not true. However, by assuming positive definiteness of the Hessian of f , we can obtain a sufficient condition.

Theorem (sufficient condition for optimality, C^2). *If $f \in C^2$ on \mathbb{R}^n , then*

$$\begin{cases} \nabla f(\mathbf{x}^*) = \mathbf{0} \\ \nabla^2 f(\mathbf{x}^*) \succ 0 \end{cases} \implies \mathbf{x}^* \text{ is a strict local minimum of } f \text{ on } \mathbb{R}^n$$

Proof. See Theorem 4.17. □

To get sufficient conditions for a point to be a local minimum, we need to assume convexity of the function f .

Theorem (necessary and sufficient condition for optimality, C^1). *If $f \in C^1$ is convex on \mathbb{R}^n , then*

$$\mathbf{x}^* \text{ is a global minimum of } f \text{ on } \mathbb{R}^n \iff \nabla f(\mathbf{x}^*) = \mathbf{0}$$

Proof. See Theorem 4.18. □

Optimality conditions for $S \subseteq \mathbb{R}^n$

When $S = \mathbb{R}^n$ the directions we could move from a point \mathbf{x} and still stay feasible were \mathbb{R}^n itself. When we consider cases where $S \subset \mathbb{R}^n$ this might not hold.

Definition (feasible direction). *Let $\mathbf{x} \in S$. A vector $\mathbf{p} \in \mathbb{R}^n$ defines a feasible direction at \mathbf{x} if*

$$\exists \delta > 0 : \mathbf{x} + \alpha \mathbf{p} \in S, \text{ for all } \alpha \in [0, \delta].$$

So the feasible directions at a point $\mathbf{x} \in S$ describes the directions in which we can "move" without becoming infeasible.

Definition (descent direction). *Let $\mathbf{x} \in \mathbb{R}^n$. A vector \mathbf{p} defines a descent direction with respect to f at \mathbf{x} if*

$$\exists \delta > 0 : f(\mathbf{x} + \alpha \mathbf{p}) < f(\mathbf{x}), \text{ for all } \alpha \in (0, \delta].$$

Suppose that $f \in C^1$ around a point $\mathbf{x} \in \mathbb{R}^n$, and that $\mathbf{p} \in \mathbb{R}^n$. If $\nabla f(\mathbf{x})^T \mathbf{p} < 0$ then the vector \mathbf{p} defines a direction of descent with respect to f at \mathbf{x} . We can now state necessary optimality conditions for cases when $S \neq \mathbb{R}^n$.

Theorem (necessary optimality conditions). *Suppose that $S \subseteq \mathbb{R}^n$ and that $f \in C^1$ on S .*

a) *If $\mathbf{x}^* \in S$ is a local minimum of f over S , then*

$$\nabla f(\mathbf{x}^*)^T \mathbf{p} \geq 0$$

holds for all feasible directions \mathbf{p} at \mathbf{x}^ .*

b) *Suppose that S is convex. If \mathbf{x}^* is a local minimum of f over S , then*

$$\nabla f(\mathbf{x}^*)^T (\mathbf{x} - \mathbf{x}^*) \geq 0, \quad \mathbf{x} \in S \tag{2}$$

Proof. See Proposition 4.22 in the book. □

We refer to (2) as a *variational inequality* and we can now extend the notion of *stationary points* by denoting them as points fulfilling (2). This is first out of four definitions of a stationary point.

Now the necessary and sufficient conditions for optimality can be stated in the following theorem.

Theorem (necessary and sufficient optimality conditions). *Suppose $S \subseteq \mathbb{R}^n$ is a convex nonempty set and that $f \in C^1$ is a convex function on S . Then*

$$\mathbf{x}^* \text{ is a global minimum of } f \text{ over } S \iff \nabla f(\mathbf{x}^*)^\top (\mathbf{x} - \mathbf{x}^*) \geq 0, \quad \mathbf{x} \in S.$$

Proof. See Theorem 4.23 in the book. □

Note that when $S = \mathbb{R}^n$, the expression to the right just becomes $\nabla f(\mathbf{x}^*) = \mathbf{0}$. Why?

We will now present three additional definitions of a stationary point which are all equivalent to (2). The first one we get by taking the minimum of the left-hand-side of (2) and then realizing that the optimal value must be zero, i.e.,

$$\min_{\mathbf{x} \in S} \nabla f(\mathbf{x}^*)^\top (\mathbf{x} - \mathbf{x}^*) = 0. \quad (3)$$

Convince yourself that (2) and (3) are equivalent! Now we claim that (2) and (3) are also equivalent with

$$\mathbf{x}^* = \text{Proj}_S [\mathbf{x}^* - \nabla f(\mathbf{x}^*)]. \quad (4)$$

The equation (4) states that if you stand in a stationary point and take a step in the direction of the negative gradient and then project back to the feasible set, you should end up in the same point. The details for showing this equivalence can be found in the book (pp. 94–95). See also Figure 2.

For the last equivalent definition of a stationary point, we need to introduce the normal cone

Definition (normal cone). *Suppose the set S is closed and convex. Let $\mathbf{x} \in S$. Then the normal cone to S at \mathbf{x} is the set*

$$N_S(\mathbf{x}) := \{\mathbf{p} \in \mathbb{R}^n \mid \mathbf{p}^\top (\mathbf{y} - \mathbf{x}) \leq 0, \quad \mathbf{y} \in S\}.$$

Think of the normal cone at a point \mathbf{x} as all direction pointing "straight out" from the set. See Figure 2. Now the fourth equivalent definition of a stationary point is that

$$-\nabla f(\mathbf{x}^*) \in N_S(\mathbf{x}^*). \quad (5)$$

That (5) is equivalent to (2) is trivial.




Figure 2: Illustration of the normal cone and the projection operator.

Summary of optimality condition for convex $S \subseteq \mathbb{R}^n$

Definition (stationary point). Suppose that S is convex and that $f \in C^1$. A point $x \in S$ fulfilling the four equivalent statements a)–d) are called a stationary point.

a)

$$\nabla f(x^*)^T(x - x^*) \geq 0, \quad x \in S,$$

b)

$$\min_{x \in S} \nabla f(x^*)^T(x - x^*) = 0,$$

c)

$$x^* = \text{Proj}_S [x^* - \nabla f(x^*)],$$

d)

$$-\nabla f(x^*) \in N_S(x^*).$$

The two important theorems which will be utilized throughout the whole course are the following.

Theorem (necessary optimality conditions). Suppose that S is convex and that $f \in C^1$. Then

$$x^* \text{ is a local minimum of } f \text{ over } S \implies x^* \text{ is stationary}$$

Theorem (necessary and sufficient optimality conditions). Suppose that S is convex and that $f \in C^1$ is convex. Then

$$x^* \text{ is a global minimum of } f \text{ over } S \iff x^* \text{ is stationary}$$

As we will see later in the course, the last definition (the inclusion (5)) is the only one that can be extended to the case of non-convex sets S .

The separation theorem

Now we will present a very useful theorem for convex sets which says that: "If a point \mathbf{y} does not lie in a closed convex set S , then there exist a hyperplane separating \mathbf{y} from S ". Mathematically, this amounts to the following.

Theorem (the separation theorem). *Suppose that $S \subseteq \mathbb{R}^n$ is closed and convex, and that the point \mathbf{y} does not lie in S . Then there exists a vector $\boldsymbol{\pi} \neq \mathbf{0}$ and a scalar $\alpha \in \mathbb{R}$ such that $\boldsymbol{\pi}^T \mathbf{y} > \alpha$, and $\boldsymbol{\pi}^T \mathbf{x} \leq \alpha$ for all $\mathbf{x} \in S$.*

Proof. See Theorem 4.29. □




Figure 3: Illustration of the separation theorem.

The separation theorem can be used to prove Farkas' Lemma efficiently, see Theorem 4.33.