

CHAPTER IV: Convex Optimization

1 Basic definitions

Consider the optimization problem

$$\begin{aligned} & \underset{\mathbf{x} \in \mathbb{R}^n}{\text{minimize}} && f_0(\mathbf{x}) && \text{(P)} \\ & \text{s.t.} && f_i(\mathbf{x}) \leq 0 && (\forall i \in [m]). \end{aligned}$$

- The function f_0 is called the *objective function*.
- The inequalities $f_i(\mathbf{x}) \leq 0$ (for $i \geq 1$) are called *constraints*.
- A vector $\mathbf{x} \in \mathbb{R}^n$ is called *feasible*, or a *feasible solution*, if it satisfies all the constraints, i.e.,

$$f_i(\mathbf{x}) \leq 0 \quad (\forall i \in [m]).$$

- The set $\mathcal{F} = \{\mathbf{x} \in \mathbb{R}^n \mid f_1(\mathbf{x}) \leq 0, \dots, f_m(\mathbf{x}) \leq 0\}$ of all feasible solutions is called the *feasible set*.
- The value $p^* = \inf\{f_0(\mathbf{x}) \mid \mathbf{x} \in \mathcal{F}\} \in \mathbb{R} \cup \{-\infty, +\infty\}$ is the *optimal value* of the problem.
- $p^* = +\infty$ iff the feasible set \mathcal{F} is empty; then we say that Problem (P) is *infeasible*.
- $p^* = -\infty$ iff there is a sequence $\mathbf{x}_i \in \mathcal{F}$ with $f_0(\mathbf{x}_i) \rightarrow -\infty$; in this case, (P) is *unbounded from below*.
- When f_0 is a constant function, the problem asks to find any feasible vector $\mathbf{x} \in \mathcal{F}$. In this case (P) is a *feasibility problem*.
- If $\mathbf{x}^* \in \mathcal{F}$ satisfies $f_0(\mathbf{x}^*) = p^*$, we say that \mathbf{x}^* is an *optimal solution* of the problem, or that \mathbf{x}^* *solves* (P). Sometimes, we'll say that \mathbf{x} is a *global optimum*, to stress the difference with *local optima* (cf. definition below). The set of all optimal solutions of (P) is called the *optimal set*.
- The vector \mathbf{x} is called ϵ -*suboptimal* if it is feasible and $f_0(\mathbf{x}) \leq p^* + \epsilon$.

Remark: The constraints are understood in the sense of extended values of the functions, i.e., $\mathbf{x} \notin \text{dom } f_i \implies f_i(\mathbf{x}) = \infty$. Hence,

$$\mathcal{F} \subseteq \bigcap_{i=0}^m \text{dom } f_i.$$

Definition 1 (local optimum). The vector \mathbf{x} is called a *local optimum* for Problem (P) if it solves the problem

$$\begin{aligned} & \underset{\mathbf{z} \in \mathbb{R}^n}{\text{minimize}} && f_0(\mathbf{z}) && \text{(P}_R\text{)} \\ & \text{s.t.} && f_i(\mathbf{z}) \leq 0 && (\forall i \in [m]); \\ & && \|\mathbf{z} - \mathbf{x}\| \leq R \end{aligned}$$

for some $R > 0$. In other words, there is a neighbourhood of \mathbf{x} in \mathcal{F} in which f_0 is minimized at \mathbf{x} .

Note that it is possible to include an *equality constraint* in a problem of the form (P), by considering two reversed inequalities $f(\mathbf{x}) \leq 0$ and $-f(\mathbf{x}) \leq 0$.

If all the functions f_0, \dots, f_m are convex, then we say that (P) is a *convex optimization problem*. If there is an equality constraint $f_i(\mathbf{x}) = 0$ in a convex optimization problem, then f_i must be both concave and convex, hence f_i is affine. Therefore, it is often convenient to write the equality constraints separately in a convex optimization problem:

$$\begin{aligned} \underset{\mathbf{x} \in \mathbb{R}^n}{\text{minimize}} \quad & f_0(\mathbf{x}) && (\text{P}_{\text{Eq}}) \\ \text{s.t.} \quad & f_i(\mathbf{x}) \leq 0 && (\forall i \in [m]); \\ & A\mathbf{x} = \mathbf{b}. \end{aligned}$$

Remark: The above definitions can be extended in the obvious manner for maximization problems (since maximizing $f_0(\mathbf{x})$ is the same as minimizing $-f_0(\mathbf{x})$). Hence, the problem “**maximize** $_{\mathbf{x} \in \mathcal{F}} f_0(\mathbf{x})$ ” is called *convex* if f_0 is concave and the f_i s are convex ($\forall i \geq 1$).

The next result states an essential property of convex optimization problems, which contributes to make them efficiently solvable:

Theorem 1. *Let (P) be a convex optimization problem. Then, any local optimum \mathbf{x}^* is also a global optimum.*

Proof. Let \mathbf{x}^* be a local optimum. Then, there exist a $R > 0$ such that

$$\mathbf{z} \in \mathcal{F}, \|\mathbf{z} - \mathbf{x}^*\| \leq R \implies f_0(\mathbf{z}) \geq f_0(\mathbf{x}^*).$$

Now, we assume (by contradiction) that \mathbf{x}^* is not globally optimal, i.e., there exists an $\mathbf{y} \in \mathcal{F}$ such that $f_0(\mathbf{y}) < f_0(\mathbf{x}^*)$. Let $\mathbf{z} = \theta\mathbf{y} + (1 - \theta)\mathbf{x}^*$, where $\theta = \frac{R}{2\|\mathbf{y} - \mathbf{x}^*\|}$ is chosen so that $\|\mathbf{z} - \mathbf{x}^*\| = \frac{R}{2}$. We have $\mathbf{z} \in \mathcal{F}$ by convexity of \mathcal{F} ; it then follows that $f_0(\mathbf{z}) \geq f_0(\mathbf{x}^*)$, but this contradicts the convexity inequality of f_0 :

$$f_0(\mathbf{z}) \leq \theta f_0(\mathbf{y}) + (1 - \theta)f_0(\mathbf{x}^*) < f_0(\mathbf{x}^*).$$

□

2 Problem reformulations

In this section, we review a few techniques useful to reformulate a problem to another *equivalent problem*. Sometimes, this is necessary, because the new problem will have better properties (e.g. solution can be characterized more easily, or it can be solved more efficiently). This, of course, would require a proper definition of what it means for two problems to be equivalent. While this can be stated in a more rigorous manner, we’ll just say that Problems P_1 and P_2 are *equivalent* if an optimal solution of P_1 can be transformed to an optimal solution of P_2 , and vice versa. This will become clear with the following examples. We use the (nonstandard) notation $P_1 \sim P_2$ to say that P_1 and P_2 are equivalent.

- (a) **Eliminating (or adding) equality constraints:** When a problem contains equality constraints “ $A\mathbf{x} = \mathbf{b}$ ”, we can remove these constraints by using a simple change of variables. Indeed, $A\mathbf{x} = \mathbf{b}$ means that \mathbf{x} belongs to some affine set L , and we know that L admits an alternative representation of the form $L = \{C\mathbf{z} + \mathbf{d} : \mathbf{z} \in \mathbb{R}^m\}$. (To obtain such a representation, take for \mathbf{d} a particular solution¹

¹For example, one solution is given by $\mathbf{d} = A^\dagger \mathbf{b}$, where A^\dagger is the Moore-Penrose pseudo inverse of A ; when A has full column rank, $A^\dagger = (A^T A)^{-1} A^T$.

of the equation $A\mathbf{x} = \mathbf{b}$, and let $C \in \mathbb{R}^{m \times r}$ be a matrix whose columns form a basis of $\mathbf{Ker} A$.) Then, with the change of variables $\mathbf{x} = C\mathbf{z} + \mathbf{d}$, Problem (P_{Eq}) can be seen to be equivalent to

$$\begin{aligned} & \underset{\mathbf{z} \in \mathbb{R}^r}{\text{minimize}} && f_0(C\mathbf{z} + \mathbf{d}) \\ & \text{s.t.} && f_i(C\mathbf{z} + \mathbf{d}) \leq 0 \quad (\forall i \in [m]); \end{aligned}$$

If (P_{Eq}) is convex, then the new problem is convex, too (composition with an affine mapping). An optimal solution \mathbf{z}^* of the above problem readily gives an optimal solution $\mathbf{x}^* = C\mathbf{z}^* + \mathbf{d}$ of (P_{Eq}) , and conversely, an optimal solution \mathbf{x}^* can be transformed to an optimal \mathbf{z}^* by solving the equation $C\mathbf{z}^* = \mathbf{x}^* - \mathbf{d}$ for \mathbf{z}^* .

- (b) **Slack variables** One can replace linear inequalities by linear equalities by introducing a *slack variable*. Indeed,

$$\mathbf{a}_i^T \mathbf{x} \leq b \iff \exists s \geq 0 : \mathbf{a}_i^T \mathbf{x} + s = b.$$

This means that we obtain an equivalent problem when we replace the constraint “ $\mathbf{a}_i^T \mathbf{x} \leq b$ ” by the system of two constraints “ $\mathbf{a}_i^T \mathbf{x} + s = b, s \geq 0$ ”, which involves the new set of variables $(\mathbf{x}, s) \in \mathbb{R}^{n+1}$.

- (c) **Epigraph reformulation:** It is always possible to assume that the objective function of a problem is linear. Indeed, Problem (P) is equivalent to

$$\begin{aligned} & \underset{\mathbf{x} \in \mathbb{R}^n, t \in \mathbb{R}}{\text{minimize}} && t \\ & \text{s.t.} && f_i(\mathbf{x}) \leq 0 \quad (\forall i \in [m]); \\ & && f_0(\mathbf{x}) \leq t. \end{aligned}$$

- (d) **Partial minimization:** When the constraints of the problem involve different blocks of variables, it is possible to reduce the problem by solving it (partially) for some blocks of variables. For example, the following two problems are equivalent:

$$\begin{aligned} & \underset{\mathbf{x}_1 \in \mathbb{R}^{n_1}, \mathbf{x}_2 \in \mathbb{R}^{n_2}}{\text{minimize}} && f_0(\mathbf{x}_1, \mathbf{x}_2) && \sim && \underset{\mathbf{x}_1 \in \mathbb{R}^{n_1}}{\text{minimize}} && \tilde{f}_0(\mathbf{x}_1) \\ & \text{s.t.} && f_i(\mathbf{x}_1) \leq 0 \quad (\forall i \in [m_1]); && && \text{s.t.} && f_i(\mathbf{x}_1) \leq 0 \quad (\forall i \in [m_1]), \\ & && g_j(\mathbf{x}_2) \leq 0 \quad (\forall j \in [m_2]); && && && \end{aligned}$$

where we have defined $\tilde{f}_0(\mathbf{x}_1) := \inf \{f_0(\mathbf{x}_1, \mathbf{x}_2) \mid \mathbf{x}_1 \in \mathbb{R}^{n_1}, \mathbf{x}_2 \in \mathbb{R}^{n_2}, g_j(\mathbf{x}_2) \leq 0 \ (\forall j \in [m_2])\}$.

3 First order optimality conditions

When the objective function f_0 is differentiable and Problem (P) is convex, we can derive a simple optimality condition, which depends only on ∇f_0 and the feasibility set $\mathcal{F} = \{\mathbf{x} \in \mathbb{R}^n \mid f_1(\mathbf{x}) \leq 0, \dots, f_m(\mathbf{x}) \leq 0\}$.

Theorem 2. *Let f_0 be differentiable. Then, a vector \mathbf{x}^* is optimal for the convex problem (P) if and only if \mathbf{x}^* is feasible (i.e., $\mathbf{x}^* \in \mathcal{F}$), and*

$$\forall \mathbf{y} \in \mathcal{F}, \quad \nabla f_0(\mathbf{x}^*)^T (\mathbf{y} - \mathbf{x}^*) \geq 0.$$

Proof. \Leftarrow : We use the first order condition of convexity for f_0 :

$$\forall \mathbf{y} \in \mathcal{F}, \quad f_0(\mathbf{y}) \geq f_0(\mathbf{x}^*) + \underbrace{\nabla f_0(\mathbf{x}^*)^T (\mathbf{y} - \mathbf{x}^*)}_{\geq 0} \geq f_0(\mathbf{x}^*).$$

Hence, \mathbf{x}^* is optimal.

\implies : We will prove the contrapositive statement. Assume there exists $\mathbf{y} \in \mathcal{F}$ such that $\nabla f_0(\mathbf{x}^*)^T(\mathbf{y} - \mathbf{x}^*) < 0$. We recognize that $\nabla f_0(\mathbf{x}^*)^T(\mathbf{y} - \mathbf{x}^*)$ is the directional derivative of f_0 in the direction of $(\mathbf{y} - \mathbf{x}^*)$, that is,

$$\nabla f_0(\mathbf{x}^*)^T(\mathbf{y} - \mathbf{x}^*) = \left. \frac{d}{dt} (t \mapsto f_0(\mathbf{x}^* + t(\mathbf{y} - \mathbf{x}^*))) \right|_{t=0} = \lim_{t \rightarrow 0} \frac{f_0(\mathbf{x}^* + t(\mathbf{y} - \mathbf{x}^*)) - f_0(\mathbf{x}^*)}{t} < 0.$$

Hence, for $t > 0$ small enough, we have $f_0(\mathbf{x}^* + t(\mathbf{y} - \mathbf{x}^*)) < f_0(\mathbf{x}^*)$, and $\mathbf{x}^* + t(\mathbf{y} - \mathbf{x}^*) \in \mathcal{F}$ by convexity of \mathcal{F} . This shows that \mathbf{x}^* is not optimal. \square

Remark: This theorem has a simple geometric interpretation. It says that \mathbf{x}^* is optimal if and only if $\nabla f_0(\mathbf{x}^*) = 0$, or if $\nabla f_0(\mathbf{x}^*)$ defines a supporting hyperplane of \mathcal{F} at \mathbf{x}^* .

The theorem above already allows us to characterize optimal solutions in a number of simple situations. For example, when the problem is unconstrained (i.e., $\mathcal{F} = \mathbb{R}^n$), we get the well-known fact that \mathbf{x}^* is optimal iff it solves the equation $\nabla f_0(\mathbf{x}^*) = 0$. The next result characterizes optimal solutions of a problem with equality constraints only

Proposition 3. Consider the optimization problem

$$\begin{aligned} \underset{\mathbf{x} \in \mathbb{R}^n}{\text{minimize}} \quad & f_0(\mathbf{x}) \\ \text{s.t.} \quad & A\mathbf{x} = \mathbf{b}, \end{aligned}$$

where f_0 is convex and differentiable. Then, \mathbf{x}^* is optimal iff $\nabla f_0(\mathbf{x}^*) \in \mathbf{Im} A^T$.

Proof. For any \mathbf{x}^* in the affine set $\mathcal{F} = \{\mathbf{x} : A\mathbf{x} = \mathbf{b}\}$, it holds $(\mathbf{y} \in \mathcal{F} \iff \mathbf{y} - \mathbf{x}^* \in \mathbf{Ker} A)$. So, the first order optimality condition for \mathbf{x}^* is

$$\forall \mathbf{y} \in \mathcal{F}, \nabla f_0(\mathbf{x}^*)^T(\mathbf{y} - \mathbf{x}^*) \geq 0 \iff \forall \mathbf{v} \in \mathbf{Ker} A, \nabla f_0(\mathbf{x}^*)^T \mathbf{v} \geq 0.$$

A linear function which is nonnegative everywhere on a vector space is necessarily constant. So, since $\mathbf{0} \in \mathbf{Ker} A$, the optimality condition can be rewritten as $\forall \mathbf{v} \in \mathbf{Ker} A, \nabla f_0(\mathbf{x}^*)^T \mathbf{v} = 0$. This is equivalent to saying that $\nabla f_0(\mathbf{x}^*)$ is orthogonal to $\mathbf{Ker} A$, that is, $\nabla f_0(\mathbf{x}^*) \in \mathbf{Im} A^T$. \square

Another simple case, which occurs often in practice, is the case of optimization over the nonnegative orthant \mathbb{R}_+^n .

Proposition 4. Consider the optimization problem

$$\begin{aligned} \underset{\mathbf{x} \in \mathbb{R}^n}{\text{minimize}} \quad & f_0(\mathbf{x}) \\ \text{s.t.} \quad & \mathbf{x} \geq \mathbf{0}, \end{aligned}$$

where f_0 is convex and differentiable. Then, \mathbf{x}^* is optimal iff the following condition holds:

$$\mathbf{x}^* \geq \mathbf{0}, \quad \nabla f_0(\mathbf{x}^*) \geq \mathbf{0}, \quad \text{and} \quad \forall i \in [n], \left(x_i = 0 \text{ or } \frac{\partial f_0}{\partial x_i}(\mathbf{x}^*) = 0 \right).$$

Proof. The first order condition for $\mathbf{x}^* \in \mathcal{F}$ reads

$$\forall \mathbf{y} \geq \mathbf{0}, \nabla f_0(\mathbf{x}^*)^T(\mathbf{y} - \mathbf{x}^*) \geq 0.$$

This already implies $\nabla f_0(\mathbf{x}^*)^T \geq 0$ (Otherwise, if $\frac{\partial f_0}{\partial x_i}(\mathbf{x}^*) < 0$ for some i , the above condition would be violated when $y_i \rightarrow \infty$.) Now, for $\mathbf{y} = \mathbf{0}$, we obtain

$$\nabla f_0(\mathbf{x}^*)^T \mathbf{x}^* = \sum_i x_i^* \frac{\partial f_0}{\partial x_i}(\mathbf{x}^*) \leq 0.$$

Each term of this sum is the product of two nonnegative numbers, so it is nonnegative. It follows that the above inequality can only hold if for all $i \in [n]$, $x_i^* \cdot \frac{\partial f_0}{\partial x_i}(\mathbf{x}^*) = 0$. This yields the condition of the proposition.

Conversely, assume that the condition of the proposition holds, and let $\mathbf{y} \geq \mathbf{0}$. Then, for all i we have

$$\frac{\partial f_0}{\partial x_i}(\mathbf{x}^*) \cdot (y_i - x_i^*) \geq 0,$$

because this expression is either the product of two nonnegative terms, or it is equal to 0 whenever $\frac{\partial f_0}{\partial x_i}(\mathbf{x}^*) = 0$. Finally, summing over i yields the first order optimality condition for \mathbf{x}^* . \square