Lecture 10 : Theory of constrained optimization

This lecture presents the theory of constrained optimization. In the first section of the lecture, we will discuss the first-order necessary conditions for a solution of constrained optimization - the Karush-Kuhn-Tucker condition. In the remaining part, we investigate the Larangian functions and the dual function. We will prove several properties of the dual functions and the weak duality.

1 Notions for the feasible sets of constrained optimization

In this lecture, we consider in particular an optimization of the following form:

$$\begin{aligned} \underset{x \in \mathbb{R}^n}{\text{Minimize}} \quad & f(x) \\ \text{subject to:} \quad & c_i(x) = 0, \quad \forall i \in \mathcal{E}, \\ & c_i(x) \geq 0 \quad \forall i \in \mathcal{I}. \end{aligned} \tag{CO}$$

where f and $c_i, i \in \mathcal{I} \cup \mathcal{E}$ are C^1 functions.

We consider the feasible set $\mathcal{F} := \{x \mid c_i(x) = 0, \forall i \in \mathcal{E}, c_i(x) \geq 0, \forall i \in \mathcal{I}\}$. In the case, $\mathcal{F}$ and f are convex, we know that the necessary and sufficient condition for $x^\star \in \mathcal{F}$ is an optimal solution is:

$$\nabla f(x^\star)^\top (y - x^\star) \geq 0, \forall y \in \mathcal{F},$$

which can be translated literally as: for any direction $d \in \mathbb{R}^n$ such that $x + td \in \mathcal{F}, \forall t \in [0, a], a > 0$, we have: $\nabla f(x^\star)^\top d \geq 0$. For a convex set $\mathcal{F}$, the set of of such direction can be simply characterized as:

$$\{y - x^\star \mid y \in \mathcal{F}\}.$$

For general feasible $\mathcal{F}$, we need to consider the following (local) notion:

Definition 1.1 (Tangent cone)**.** The vector d is said to be tangent (or tangent vector) to $\mathcal{F}$ at a point $x \in \mathcal{F}$ if there is a sequence $\{z_k\}_{k \in \mathbb{N}}, z_k \in \mathcal{F}, \forall k \in \mathbb{N}$, and a sequence of positive scalar $\{t_k\}_{k \in \mathbb{N}}$ such that:

$$\lim_{k \to \infty} t_k = 0 \qquad , \qquad \lim_{k \to \infty} \frac{z_k - x}{t_k} = d.$$

The set of all tangent vectors to $\mathcal{F}$ at x is called the tangent cone and is denoted by $T_{\mathcal{F}}(x)$.

Intuitively, the tangent cone to $\mathcal{F}$ at a point x are equal to the set direction where one can depart from x such that we still remain in $\mathcal{F}$. Consider the following examples

Example 1.2. Consider two examples:

1. $\mathcal{F} = \overline{\mathcal{B}(0,1)}$ is the closed unit ball: the tangent cone to $\mathcal{F}$ at x is $\mathbb{R}^d$ if $x \in \mathcal{B}(0,1)$ and $\{d \mid d^\top x \leq 0\}$ otherwise.

2. $\mathcal{F} = \{(x,x) \mid x \in \mathbb{R}^n\} \cup \{(x,-x) \mid x \in \mathbb{R}^n\}$: the tangent cone to $\mathcal{F}$ at $(0,0)$ is: $\{t(1,1) \mid t \in \mathbb{R}\} \cup \{t(1,-1) \mid t \in \mathbb{R}\}$.

The tangent cone allows us to obtain the following necessary for an optimal solution.

Theorem 1.3 (Necessary conditions for an optimal solution)**.** *Consider a constrained optimization problem of the form* (CO)*. If x is a (local) optimal solution, then:*

$$\nabla f(x)^\top d \geq 0, \forall d \in T_{\mathcal{F}}(x).$$

Proof. For the sake of contradiction, assume there exists $d \in T_{\mathcal{F}}(x)$ such that $\nabla f(x)^\top d < 0$. By the definition of tangent cone, there exists a sequence of $d_k \in \mathbb{R}^n$ and $t_k > 0, k \in \mathbb{N}$ such that:

1. The sequence d_k converges to d.

2. The sequence t_k converges to 0.

3. $x + t_k d_k \in \mathcal{F}, \forall k \in \mathbb{N}$.

Consider the Taylor expansion of f at x, we have:

$$f(x + t_k d_k) = f(x) + t_k \nabla f(x)^\top d_k + R(t_k d_k),$$

where $R(t_k d_k)$ is a correction term satisfying: $R(t_k d_k) = o(t_k)$. Since $d_k \to d$, $\nabla f(x)^\top d_k \to \nabla f(x)^\top d$. Thus, for k sufficiently large, one has $\nabla f(x)^\top d_k < 0$. Again, choose k sufficiently large one more time so that $R(t_k d_k)$ is dominated by $t_k \nabla f(x)^\top d_k$. This causes a contradiction. $\square$

Effectively, Theorem 1.3 proof is very similar to the classical results about local minimum implying critical point. Nevertheless, in the constrained case, one needs to restrict the set of feasible direction to the tangent cone. Theorem 1.3 is, however, not satisfying because given a problem of the form (CO), Theorem 1.3 does not tell us how to verify this condition. In general, the set $\mathcal{F}$ and the tangent cone $T_{\mathcal{F}}(x)$ are not simple to compute.

In the following, we consider an alternative version of $T_{\mathcal{F}}(x)$, which is tractable and can be manipulated more efficiently.

Definition 1.4 (Active set)**.** Consider a feasible set $\mathcal{F}$ defined by the constraints of (CO). The active set $\mathcal{A}(x)$ at a feasible point x is define as:

$$\mathcal{A}(x) := \mathcal{E} \cup \{i \in \mathcal{I} \mid c_i(x) = 0\}.$$

Definition 1.5 (Linearized feasible directions)**.** Given a feasible set x and the active constrained set $\mathcal{A}(x)$ (cf. Definition 1.4), the set of linearized feasible direction $\mathcal{F}(x)$ is given by:

$$\mathcal{F}(x) := \{d \mid d^\top \nabla c_i(x) = 0, \forall i \in \mathcal{E}\} \cap \{d \mid d^\top \nabla c_i(x) \geq 0, \forall i \in \mathcal{A}(x) \cap \mathcal{I}\}.$$

Literally, Definition 1.5 defines linearized feasible directions as d that make $c_i, i \in \mathcal{E}$ constant and $c_i, i \in \mathcal{A}(x) \cap \mathcal{I}$ remain positive (using the first-order Taylor expansion). We have the following relation between the tangent cone and the set of linear feasible directions:

Lemma 1.6 (Relation between $T_{\mathcal{F}}(x)$ and $\mathcal{F}(x)$)**.** *Consider $\mathcal{F}$ defined by the constraints of* (CO) *and a feasible solution $x \in \mathcal{F}$, then:*

$$T_{\mathcal{F}}(x) \subseteq \mathcal{F}(x).$$

Proof. Let $d \in T_{\mathcal{F}}(x)$. By definition of the tangent cone, there exists a sequence $d_k \to d$ and $t_k \to 0, t_k > 0, \forall k \in \mathbb{N}$ such that $x + t_k d_k \in \mathcal{F}$. We consider two elements:

1. For $i \in \mathcal{E}$, we have:

$$0 = \lim_{k\to\infty} \frac{c_i(x + t_k d_k) - c_i(x)}{t_k} = \lim_{k\to\infty} \frac{t_k \nabla c_i(x)^\top d_k + o(t_k d_k)}{t_k} = \nabla f(x)^\top d.$$

2. For $i \in \mathcal{A}(x) \cap \mathcal{I}$, we have:

$$0 \geq \lim_{k\to\infty} \frac{c_i(x + t_k d_k) - c_i(x)}{t_k} = \nabla f(x)^\top d.$$

□

Example 1.7. We consider the two previous examples:

1. $\mathcal{F} = \overline{\mathcal{B}(0,1)}$ is defined by the constraint $\|x\|^2 \leq 1$. We have:

$$\mathcal{F}(x) = \begin{cases} \mathbb{R}^2, & \text{if } x \in \mathcal{B}(0,1) \\ -x^\top d \geq 0, & \text{otherwise} \end{cases} = T_{\mathcal{F}}(x).$$

2. $\mathcal{F} = \{(x,x) \mid x \in \mathbb{R}^n\} \cup \{(x,-x) \mid x \in \mathbb{R}^n\}$ is defined by the constraint: $x^2 - y^2 = 0$. We have:

$$\mathcal{F}(0) = \{d \mid d^\top 0 = 0\} = \mathbb{R}^2 \supsetneq T_{\mathcal{F}}(0).$$

In general, we do not have $T_{\mathcal{F}}(x) = \mathcal{F}(x)$. The condition for which equality happens is called constraint qualifications. We consider several constrain qualifications:

1. Linearity Constraint Qualification (LCQ): If $c_i, i \in \mathcal{E} \cup \mathcal{I}$, then $T_{\mathcal{F}}(x) = \mathcal{F}(x)$. The converse of Lemma 1.6 is quite simple to prove (one can take the sequence $z_k = x + td_k$ and t_k is an arbitrary sequence of positive scalars converging to 0).

2. Linear Independence Constraint Qualification (LICQ): If $\nabla c_i(x), i \in \mathcal{A}(x)$ is linearly independent, then $T_{\mathcal{F}}(x) = \mathcal{F}(x)$. Note that LICQ is verified by the first example, and not by the second one.

3. Slatter's condition: if $c_i, i \in \mathcal{E}$ are linear and $c_i, i \in \mathcal{I}$ are convex and there exists a point x such that $c_i(x) = 0, \forall i \in \mathcal{E}$ and $c_i(x) > 0, \forall i \in \mathcal{I}$, then $T_{\mathcal{F}}(x) = \mathcal{F}(x), \forall x$. Slatter condition is verified by the first example, and not by the second one.

Assuming these constraint qualifications, we can state the Karush-Kuhn-Tucker (a.k.a KKT) conditions as:

Theorem 1.8 (KKT conditions)**.** *Consider a constrained optimization problem of the form* (CO)*. Assume that a constraint qualification is satisfied at a feasible point x, then if x is a (local) solution of* (CO)*, then there exists $\lambda \in \mathbb{R}^{|\mathcal{I}|}$ and $\nu \in \mathbb{R}^{|\mathcal{E}|}$ such that:*

$$\begin{aligned} \nabla f(x) - \sum_{i\in\mathcal{E}} \nu_i \nabla c_i(x) - \sum_{i\in\mathcal{I}} \lambda_i \nabla c_i(x) = 0 && \text{(Stationarity)} \\ c_i(x) = 0, \forall i \in \mathcal{E} && \text{(Primal feasibility)} \\ c_i(x) \geq 0, \forall i \in \mathcal{E} && \text{(Primal feasibility)} \\ \lambda \geq 0, && \text{(Dual feasibility)} \\ \sum_{i\in\mathcal{I}} c_i(x)\lambda_i = 0. && \text{(Complementary slackness)} \end{aligned}$$

Proof. In this proof, we admit without proving the following result (you will see this in the tutorial session):

Lemma 1.9 (Farkas lemma)**.** *Let* $\mathbf{B} \in \mathbb{R}^{n\times m}$ *and* $\mathbf{C} \in \mathbb{R}^{n\times p}$*. Define the cone:*

$$K := \{\mathbf{B}y + \mathbf{C}w \mid y \in \mathbb{R}^m_+, w \in \mathbb{R}^p\}.$$

Only one of the two alternatives is true:

1. $g \in K$.
2. *There exists* $d \in \mathbb{R}^n$ *such that* $g^\top d < 0$, $\mathbf{B}^\top d \geq 0, \mathbf{C}^\top d = 0$.

If $x \in \mathcal{F}$ is a local solution, due to Theorem 1.3, we have:

$$g^\top d \geq 0, \forall d \in T_{\mathcal{F}}(x).$$

However, since $T_{\mathcal{F}}(x) = \mathcal{F}(x)$, it implies that:

$$g^\top d \geq 0, \forall d \in K := \{\mathbf{B}\lambda + \mathbf{C}\nu \mid \lambda \geq 0\},$$

where $\mathbf{B} = (\nabla c_i(x))_{i\in\mathcal{I}\cap\mathcal{A}(x)}, \mathbf{C} = (\nabla c_i(x))_{i\in\mathcal{E}}$. By the Farkas lemma, $g \in K$. Therefore, there exists $\lambda \geq 0$ and $\nu \in \mathbb{R}^{|\mathcal{E}|}$ such that:

$$\nabla f(x) = \sum_{i\in\mathcal{I}\cap\mathcal{A}(x)} \lambda_i \nabla c_i(x) + \sum_{i\in\mathcal{E}} \nu_i \nabla c_i(x).$$

By choosing $\lambda_i = 0, i \in \mathcal{I} \setminus \mathcal{A}(x)$, we have:

$$\nabla f(x) - \sum_{i\in\mathcal{I}} \lambda_i \nabla c_i(x) - \sum_{i\in\mathcal{E}} \nu_i \nabla c_i(x) = 0.$$

Since x is feasible, the two equations of primal feasibility are satisfied. By our choice of $\lambda \geq 0$, the equation of dual feasibility is also satisfied.

Lastly, for $i \in \mathcal{I}\cap\mathcal{A}(x)$, $\lambda_i c_i(x) = 0$. For the remaining indices, we choose $\lambda_i = 0$ and hence, the complementary slackness is also satisfied. □

2 Duality theory

In Theorem 1.8, we saw the appearance of a function of the form:

$$\mathcal{L}(x, \lambda, \nu) = f(x) - \sum_{i\in\mathcal{I}} \lambda_i c_i(x) - \sum_{i\in\mathcal{E}} \nu_i c_i(x).$$

Indeed, the stationarity equation implies that $\nabla_x \mathcal{L}(x, \lambda, \nu) = 0$. This function is known as the Lagrangian function of (CO). The variables λ, ν are called Lagrangian multipliers. The dual objective function is defined as:

$$q(\lambda, \nu) = \inf_x \mathcal{L}(x, \lambda, \nu).$$

The domain of the dual function q is given by: $\mathcal{D} := \{(\lambda, \nu) \mid q(\lambda, \nu) > -\infty\}$. We have the following property of dual functions.

Lemma 2.1 (Concavity of dual functions)**.** *The function* $q(\lambda, \nu)$ *is concave. Thus, its domain is convex.*

Proof. For any two pairs of $(\lambda_i, \nu_i), i = 1, 2$ and $t \in [0, 1]$, we get:

$$\begin{aligned} tq(\lambda_1, \nu_1) + (1-t)q(\lambda_2, \nu_2) &= \inf_x \; t\mathcal{L}(x, \lambda_1, \nu_1) + \inf_x (1-t)\mathcal{L}(x, \lambda_2, \nu_2) \\ &\leq \inf_x t\mathcal{L}(x, \lambda_1, \nu_1) + (1-t)\mathcal{L}(x, \lambda_2, \nu_2) \\ &= \inf_x \; \mathcal{L}(x, t\lambda_1 + (1-t)\lambda_2, t\nu_1 + (1-t)\nu_2) \\ &= q(t\lambda_1 + (1-t)\lambda_2, t\nu_1 + (1-t)\nu_2). \end{aligned}$$

Thus, the function is concave. Its domain is, thus, convex. □

The dual function is always concave, regardless the property of $f(x), c_i(x)$. Dual function can be used to lower-bound the optimal value of a constrained problem. The dual optimization problem (corresponding to the primal one, i.e., (CO)) is:

$$\underset{\lambda, \nu}{\text{Maximize}} \quad q(\lambda, \nu) \qquad \text{such that} \qquad \lambda \geq 0. \tag{D}$$

This is the main consequence of the following result:

Theorem 2.2 (Weak duality)**.** *Let $p^\star$ and $d^\star$ be the optimal value of* (CO) *and* (D) *respectively. We have: $p^\star \geq d^\star$.*

Proof. For any $x \in \mathcal{F}$ and $\lambda \geq 0$, we have:

$$f(x) \geq \mathcal{L}(x, \lambda, \nu) = f(x) - \underbrace{\sum_{i \in \mathcal{I}} \lambda_i c_i(x)}_{\geq 0} - \underbrace{\sum_{i \in \mathcal{E}} \nu_i c_i(x)}_{=0} \geq \inf_x \mathcal{L}(x, \lambda, \nu) = q(\lambda, \nu).$$

Therefore, the minimum value of $f(x), x \in \mathcal{F}$ is bigger than the maximum of $q(\lambda, \nu), \lambda \geq 0$. □

Therefore, maximizing the concave function $q(\lambda, \nu)$ (equivalent to minimizing the convex function $-q$) can give us a lowerbound for the optimization problem (regardless of the property of the original (CO)).

A more interesting case is the strong duality theory, which is usually studied in the context of convex optimization. In that case, we have $p^\star = d^\star$, which allows one to reduce a general optimization problem to a convex one.