Lecture 4: CS395T Numerical Optimization for Graphics and AI — Theory of Constrained Optimization

Qixing Huang
The University of Texas at Austin
huangqx@cs.utexas.edu

Disclaimer

This note is adapted from

- Section 12 of *Numerical Optimization* by Jorge Nocedal and Stephen J. Wright. Springer series in operations research and financial engineering. Springer, New York, NY, 2. ed. edition, (2006)

1 Introduction

The second part of this class is about minimizing functions subject to constraints on the variables. A general formulation for these problems is

$$\begin{aligned} \underset{\boldsymbol{x}}{\text{minimize}} \quad & f(\boldsymbol{x}) \\ \text{subject to} \quad & c_i(\boldsymbol{x}) = 0, \quad i \in \mathcal{E}, \\ & c_i(\boldsymbol{x}) \geq 0, \quad i \in \mathcal{I}, \end{aligned} \tag{1}$$

where f and the functions c_i are all smooth, real-valued functions on a subset of $\mathbb{R}^n$, and $\mathcal{I}$ and $\mathcal{E}$ are two finite sets of indices.

If we define the feasible set Ω to be the set of points $\boldsymbol{x}$ that satisfies the constraints, that is,

$$\Omega = \{\boldsymbol{x} | c_i(\boldsymbol{x}) = 0, i \in \mathcal{E}; c_i(\boldsymbol{x}) \geq 0, i \in \mathcal{I}\}, \tag{2}$$

then we can always rewrite (1) more compactly as

$$\min_{\boldsymbol{x} \in \Omega} f(\boldsymbol{x}) \tag{3}$$

In this lecture, we will go through some optimally conditions. They are generalized from their counterparts in the constrained case. They are summarized below.

2 Local and Global Solutions

We have seen already that global solutions are difficult to find even when there are noconstraints. The situation may be improved when we add constraints, since the feasible setmight exclude many of the local minima and it may be comparatively easy to pick the global minimum from those that remain. However, constraints can also make things much more difficult. As an example, consider the problem

$$\min_{\boldsymbol{x}\in\mathbb{R}^n} \|\boldsymbol{x}\|_2^2, \quad \text{subject to } \|\boldsymbol{x}\|_2^2 \geq 1.$$

Without the constraint, this is a convex quadratic problem with unique minimizer $\boldsymbol{x} = \mathbf{0}$. When the constraint is added, any vector $\boldsymbol{x}$ with $\|\boldsymbol{x}\|_2 = 1$ solves the problem. There are infinitely many such vectors (hence, infinitely many local minima) whenever $n \geq 2$.

A second example shows how addition of a constraint produces a large number of local solutions that do not form a connected set. Consider

$$\min \ (x_2 + 100)^2 + 0.01x_1^2, \quad \text{subject to} x_2 - \cos(x_1) \geq 0/$$

Without the constraint, the problem has the unique solution $(-100, 0)$. With the constraint there are local solutions near the points

$$(x_1, x_2) = (k\pi, -1), \quad \text{for } k = \pm 1, \pm 3, \pm 5, \cdots$$

Definitions of the different types of local solutions are simple extensions of the corresponding definitions for the unconstrained case, except that now we restrict consideration to the feasible points in the neighborhood of $\boldsymbol{x}^\star$. We have the following definition.

Definition 2.1. *A vector $\boldsymbol{x}^\star$ is a local solution to (3) if $\boldsymbol{x}^\star \in \Omega$ and there is a neighborhood $\mathcal{N}$ of $\boldsymbol{x}^\star$ such that $f(\boldsymbol{x}) \geq f(\boldsymbol{x}^\star)$ for $\boldsymbol{x} \in \mathcal{N} \cap \Omega$.*

Definition 2.2. *A vector $\boldsymbol{x}^\star$ is a strict local solution (also called a strong local solution) if $\boldsymbol{x}^\star \in \Omega$ and there is a neighborhood $\mathcal{N}$ of $\boldsymbol{x}^\star$ such that $f(\boldsymbol{x}) \geq f(\boldsymbol{x})^\star)$ for all $\boldsymbol{x} \in \mathcal{N} \cap \Omega$ with $\boldsymbol{x} \neq \boldsymbol{x}^\star$.*

Definition 2.3. *A point $\boldsymbol{x}^\star$ is an isolated local solution if $\boldsymbol{x}^\star \in \Omega$ and there is a neighborhood $\mathcal{N}$ of $\boldsymbol{x}^\star$ such that $\boldsymbol{x}^\star$ is the only local minimize*

3 Examples

To introduce the basic principles behind the characterization of solutions of constrained optimization problems, we work through three simple examples. The ideas discussed here will be made rigorous in the sections that follow. We start by noting one item of terminology that recurs throughout the rest of the class: At a feasible point $\boldsymbol{x}$, the inequality constraint $i \in \mathcal{I}$ is said to be active if $c_i(\boldsymbol{x}) = 0$ and inactive if the strict inequality $c_i(\boldsymbol{x}) > 0$ is satisfied.

3.1 A Single Equality Constraint

Our first example is a two-variable problem with a single equality constraint:

$$\min \ x_1 + x_2 \quad s.t. \ x_1^2 + x_2^2 - 2 = 0 \tag{4}$$

In the language of (1), we have $f(\boldsymbol{x}) = x_1 + x_2, \mathcal{I} = \emptyset, \mathcal{E} = \{1\}$ and $c_1(\boldsymbol{x}) = x_1^2 + x_2^2 - 2$. We can see by inspection that the feasible set for this problem is the circle of radius $\sqrt{2}$centered at the origin-just the boundary of this circle, not its interior. The solution $\boldsymbol{x}^\star$ is obviously $(-1, -1)^T$. From any other point on the circle, it is easy to find a way to move that stays feasible (that is, remains on the circle) while decreasing f. For instance, from the point $\boldsymbol{x} = (\sqrt{2}, 0)^T$ any move in the clockwise direction around the circle has the desired effect. We also see that at the solution $\boldsymbol{x}^\star$, the constraint normal $\nabla c_1(\boldsymbol{x}^\star)$ is parallel to $\nabla f(\boldsymbol{x}^\star)$. That is, there is a scalar $\lambda_1^\star$ such that

$$\nabla f(\boldsymbol{x}^\star) = \lambda_1^\star \nabla c_1^\star(\boldsymbol{x}^\star). \tag{5}$$

(In this particular case, we have $\lambda_1^\star = -\frac{1}{2}$.)

We can derive (5) by examining first-order Taylor series approximations to the objective and constraint functions.To retain feasibility with respect to the function $c_1(\boldsymbol{x}) = 0$, we require that $c_1(\boldsymbol{x}+\boldsymbol{d}) = 0$; that is,

$$0 = c_1(\boldsymbol{x}+\boldsymbol{d}) \approx c_1(\boldsymbol{x}) + \nabla c_1(\boldsymbol{x})^T \boldsymbol{d} = \nabla c_1(\boldsymbol{x})^T \boldsymbol{d}. \tag{6}$$

Hence, the direction $\boldsymbol{d}$ retains feasibility with respect to c_1, to first order, when it satisfies

$$\nabla c_1(\boldsymbol{x})^T \boldsymbol{d} = 0. \tag{7}$$

Similarly, a direction of improvement must produce a decrease in f, so that

$$0 > f(\boldsymbol{x}+\boldsymbol{d}) - f(\boldsymbol{x}) \approx \nabla f(\boldsymbol{x})^T \boldsymbol{d},$$

or, to first order,

$$\nabla f(\boldsymbol{x})^T \boldsymbol{d} < 0. \tag{8}$$

If there exists a direction $\boldsymbol{d}$ that satisfies both (7) and (8), we conclude that improvement on our current point $\boldsymbol{x}$ is possible. It follows that a necessary condition for optimality for the problem (4) is that there exist no direction $\boldsymbol{d}$ satisfying both (7) and (8).

It is easy to check that the only way that such a direction cannot exist is if $\nabla f(\boldsymbol{x})$ and $\nabla c_1(\boldsymbol{x})$ are parallel, that is, if the condition

$$\nabla f(\boldsymbol{x}) = \lambda_1 \nabla c_1(\boldsymbol{x})$$

holds at $\boldsymbol{x}$, for some scalar λ_1. If this condition is not satisfied, the direction defined by

$$\boldsymbol{d} == \big(I - \frac{\nabla c_1(\boldsymbol{x}) \nabla c_1(\boldsymbol{x})^T}{\|\nabla c_1(\boldsymbol{x})\|^2}\big) \nabla f(\boldsymbol{x}). \tag{9}$$

satisfies both conditions (7) and (8). By introducing the Lagrangian function

$$L(\boldsymbol{x}, \lambda) = f(\boldsymbol{x}) - \lambda_1 c_1(\boldsymbol{x}). \tag{10}$$

and noting that $\nabla_{\boldsymbol{x}} f(\boldsymbol{x}, \lambda_1) = \nabla f(\boldsymbol{x}) - \lambda_1 \nabla c_1(\boldsymbol{x})$, we can state the condition (5) equivalently as follows: At the solution $\boldsymbol{x}^\star$,, there is a scalar $\lambda_1^\star$ such that

$$\nabla_{\boldsymbol{x}} L(\boldsymbol{x}^\star, \lambda_1^\star) = 0. \tag{11}$$

This observation suggests that we can search for solutions of the equality-constrained problem (4) by searching for stationary points of the Lagrangian function. The scalar quantity λ_1 in (10) is called a Lagrange multiplier for the constraint $c_1(\boldsymbol{x}) = 0$.

Though the condition (5) (equivalently, (11)) appears to be necessary for an optimal solution of the problem (4), it is clearly not sufficient. For instance, (5) is satisfied at the point $\boldsymbol{x} = (1, 1)^T$ (with $\lambda_1 = \frac{1}{2}$), but this point is obviously not a solution-in fact, it maximizes the function f on the circle. Moreover, in the case of equality-constrained problems, we cannot turn the condition (5) into a sufficient condition simply by placing some restriction on the sign of λ_1. To see this, consider replacing the constraint $x_1^2 + x_2^2 - 2 = 0$ by its negative $2 - x_1^2 - x_2^2 = 0$. The solution of the problem is not affected, but the value of $\lambda_1^\star$ that satisfies the condition (5) changes from $\lambda_1^\star = -\frac{1}{2}$ to $\lambda_1^\star = \frac{1}{2}$.

3.2 A Single Inequality Constraint

This is a slight modification of the first example, in which the equality constraint is replaced by an inequality. Consider

$$\min x_1 + x_2 \quad s.t.\ 2 - x_1^2 - x_2^2 \geq 0, \tag{12}$$

for which the feasible region consists of the circle of problem (4) and its interior. Note that the constraint normal ∇c_1 points toward the interior of the feasible region at each point on the boundary of the circle. By inspection, we see that the solution is still $\boldsymbol{x}^\star = (-1, -1)^T$ and that the condition (5) holds for the value

$\lambda_1^\star = \frac{1}{2}$. However, this inequality-constrained problem differs from the equality-constrained problem (4) in that the sign of the Lagrange multiplier plays a significant role.

As before, we conjecture that a given feasible point $\boldsymbol{x}$ is not optimal if we can find a step $\boldsymbol{d}$ that both retains feasibility and decreases the objective function f to first order. The main difference between problems (4) and (12) comes in the handling of the feasibility condition. As in (8), the direction $\boldsymbol{d}$ improves the objective function, to first order, if $\nabla f(\boldsymbol{x})^T \boldsymbol{d} < 0$. Meanwhile, the direction $\boldsymbol{d}$ retains feasibility if

$$0 \leq c_1(\boldsymbol{x} + \boldsymbol{d}) \approx c_1(\boldsymbol{x}) + \nabla c_1(\boldsymbol{x})^T \boldsymbol{d}.$$

so, to first order, feasibility is retained if

$$c_1(\boldsymbol{x}) + \nabla c_1(\boldsymbol{x})^T \boldsymbol{d} \geq 0. \tag{13}$$

In determining whether a direction $\boldsymbol{d}$ exists that satisfies both (8) and (13), we consider the following two cases:

Case I: Consider first the case in which $\boldsymbol{x}$ lies strictly inside the circle, so that the strict inequality $c_1(\boldsymbol{x}) > 0$ holds. In this case, any vector $\boldsymbol{d}$ satisfies the condition (13), provided only that its length is sufficiently small. In particular, whenever$\nabla f(\boldsymbol{x}^\star) \neq 0$, we can obtain a direction $\boldsymbol{d}$ that satisfies both (8) and (13) by setting

$$d = -c_1(\boldsymbol{x}) \frac{\nabla f(\boldsymbol{x})}{\|\nabla f(\boldsymbol{x})\|}.$$

The only situation in which such a direction fails to exist is when

$$\nabla f(\boldsymbol{x}) = 0. \tag{14}$$

Case II: Consider now the case in which $\boldsymbol{x}$ lies on the boundary of the circle, so that $c_1(\boldsymbol{x}) = 0$. The conditions (8) and (13) therefore become

$$\nabla f(\boldsymbol{x})^T \boldsymbol{d} < 0, \quad \nabla c_1(\boldsymbol{x})^T \boldsymbol{d} \geq 0.$$

The first of these conditions defines an open half-space, while the second defines a closed half-space. It is clear from this figure that the two regions fail to intersect only when $\nabla f(\boldsymbol{x})$ and $\nabla c_1(\boldsymbol{x})$ point in the same direction, that is, when

$$\nabla f(\boldsymbol{x}) = \lambda_1 \nabla c_1(\boldsymbol{x}), \quad \text{for some } \lambda_1 \geq 0. \tag{15}$$

Note that the sign of the multiplier is significant here. If (5) were satisfied with a negative value of λ_1, then $\nabla f(\boldsymbol{x})$ and $\nabla c_1(\boldsymbol{x})$ would point in opposite directions, and we can see that the set of directions that satisfy both (8) and (13) would make up an entire open half-plane. The optimality conditions for both cases I and II can again be summarized neatly with reference to the Lagrangian function. When no first-order feasible descent direction exists at some point $\boldsymbol{x}^\star$, we have that

$$\nabla_{\boldsymbol{x}^\star} L(\boldsymbol{x}^\star, \lambda_1^\star) = 0, \quad \text{for some } \lambda_1^\star \geq 0. \tag{16}$$

where we also require that

$$\lambda_1^\star c_1(\boldsymbol{x}_1^\star) = 0. \tag{17}$$

This condition is known as a complementarity condition; it implies that the Lagrange multiplier λ_1 can be strictly positive only when the corresponding constraint c_1 is active. Conditions of this type play a central role in constrained optimization. In case I, we have that $c_1(\boldsymbol{x}^\star) > 0$, so (17) requires that $\lambda_1^\star = 0$. Hence, (16) reduces to $\nabla f(\boldsymbol{x}^\star) = 0$, as required by (14). In case II, (17) allows $\lambda_1^\star$ to take on a nonnegative value, so (16) becomes equivalent to (17).

Exercise. Please apply the methodology described above to analyze the following constrained optimization problem that has two inequality constraints:

$$\min \; x_1 + x_2 \quad s.t. \quad 2 - x_1^2 - x_2^2 \geq 0, \; x_2 \geq 0. \tag{18}$$

4 First-Order Optimality Conditions

To define the first-order optimality conditions, we first consider the notion of active sets:

Definition 4.1. *The active set* $\mathcal{A}(\boldsymbol{x})$ *at any feasible* $\boldsymbol{x}$ *consists of the equality constraint indices from* $\mathcal{E}$ *together with the indices of the inequality constraints* i *for which* $c_i(\boldsymbol{x}) = 0$*; that is,*

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

We will also need the characterization of a local solution:

Definition 4.2. *A vector* $\boldsymbol{x}^\star$ *is a local solution of the problem (3) if* $\boldsymbol{x}^\star \in \Omega$ *and there is a neighborhood* $\mathcal{N}$ *of* $\boldsymbol{x}^\star$ *such that* $f(\boldsymbol{x}) \geq f(\boldsymbol{x}^\star)$ *for* $\boldsymbol{x} \in \mathcal{N} \cap \Omega$.

To define the optimality conditions, we will also need the so-called LICQ condition.

Definition 4.3. *Given the point* $\boldsymbol{x}$ *and the active set* $\mathcal{A}(\boldsymbol{x})$ *defined in definition 4.1, we say that the linear independence constraint qualification (LICQ) holds if the set of active constraint gradients* $\{\nabla c_i(\boldsymbol{x}), i \in \mathcal{A}(\boldsymbol{x})\}$ *is linearly independent.*

Now we are ready to introduce first-order necessary conditions, which often known as the Karush-Kuhn-Tucker conditions, or KKT conditions for short.

Theorem 4.1. *Consider the Lagrangian given by*

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

Suppose $\boldsymbol{x}^\star$ *is a local solution of (1), that the functions* f *and* c_i *in (1) are continuously differentiable, and that the LICQ holds at* $\boldsymbol{x}^\star$*. Then there is a Lagrangian multipler vector* $\lambda^\star$*, with components* $\lambda_i^\star, i \in \mathcal{E} \cup \mathcal{I}$*, such that the following conditions are satisfied at* $(\boldsymbol{x}^\star, \lambda^\star)$

$$\begin{aligned}
\nabla_{\boldsymbol{x}} L(\boldsymbol{x}^\star, \lambda^\star) &= 0, && (19)\\
c_i(\boldsymbol{x}^\star) &= 0, \quad \text{for all } i \in \mathcal{E}, && (20)\\
c_i(\boldsymbol{x}^\star) &\geq 0, \quad \text{for all } i \in \mathcal{I}, && (21)\\
\lambda_i^\star &\geq 0, \quad \text{for all } i \in \mathcal{I}, && (22)\\
\lambda_i^\star c_i(\boldsymbol{x}^\star) &= 0, \quad \text{for all } i \in \mathcal{I} \cup \mathcal{E}. && (23)
\end{aligned}$$

5 Second-Order Optimality Conditions

To derive second-order optimality conditions, we begin with defining the feasible direction set, which we define as follows.

Definition 5.1. *Given a feasible point* $\boldsymbol{x}$ *and the active constraint set* $\mathcal{A}(\boldsymbol{x})$ *of Definition 4.1, the set of linearized feasible directions* $\mathcal{F}(\boldsymbol{x})$ *is*

$$\mathcal{F}(\boldsymbol{x}) = \left\{ \boldsymbol{d} \middle| \begin{array}{ll} \boldsymbol{d}^T \nabla c_i(\boldsymbol{x}) = 0, & \text{for all } i \in \mathcal{E}, \\ \boldsymbol{d}^T \nabla c_i(\boldsymbol{x}) \geq 0, & \text{for all } i \in \mathcal{A}(\boldsymbol{x}) \cap \mathcal{I} \end{array} \right\} \tag{24}$$

Definition 5.2. *Given* $\mathcal{F}(\boldsymbol{x}^\star)$ *from Definition 5.1 and some Lagrangian multipler vector* $\lambda^\star$ *satisfying the KKT conditions, we define the critical cone* $\mathcal{C}(\boldsymbol{x}^\star, \lambda^\star)$ *as follows:*

$$\mathcal{C}(\boldsymbol{x}^\star, \lambda^\star) = \{\boldsymbol{w} \in \mathcal{F}(\boldsymbol{x}^\star) | \boldsymbol{w}^T \nabla c_i(\boldsymbol{x}^\star) = 0, \text{all } i \in \mathcal{A}(\boldsymbol{x}^\star) \cap \mathcal{I} \text{ with } \lambda_i^\star > 0\} \tag{25}$$

Equivalently,

$$\boldsymbol{w} \in \mathcal{C}(\boldsymbol{x}^\star, \lambda^\star) \leftrightarrow \begin{cases} \boldsymbol{w}^T \nabla c_i(\boldsymbol{x}^\star) = 0, & \text{for all } i \in \mathcal{E}, \\ \boldsymbol{w}^T \nabla c_i(\boldsymbol{x}^\star) = 0, & \text{for all } i \in \mathcal{A}(\boldsymbol{x}^\star) \cap \mathcal{I} \text{ with } \lambda_i^\star > 0, \\ \boldsymbol{w}^T \nabla c_i(\boldsymbol{x}^\star) \geq 0, & \text{for all } i \in \mathcal{A}(\boldsymbol{x}^\star) \cap \mathcal{I} \text{ with } \lambda_i^\star = 0. \end{cases} \tag{26}$$

The critical cone contains those directions $\boldsymbol{w}$ that would tend to "adhere" to the active inequality constraints even when we were to make small changes to the objective (those indices $i \in \mathcal{I}$ for which the Lagrange multiplier component $\lambda_i^\star$ is positive), as well as to the equality constraints. An important property of these directions is:

$$\boldsymbol{w} \in \mathcal{C}(\boldsymbol{x}^\star, \lambda^\star) \to \boldsymbol{w}^T \nabla f(\boldsymbol{x}^\star) = \sum_{i \in \mathcal{E} \cup \mathcal{I}} \lambda_i^\star \boldsymbol{w}^T \nabla c_i(\boldsymbol{x}^\star) = 0.$$

Theorem 5.1. ***(Second-Order Necessary Conditions**.) Suppose $\boldsymbol{x}^\star$ is a local solution of (1) and that LICQ condition is satisfied. Let $\lambda^\star$ be the Lagrangian multiplier vector for which the KKT conditions are satisfied. Then*

$$\boldsymbol{w}^T \nabla^2_{\boldsymbol{xx}} L(\boldsymbol{x}^\star, \lambda^\star) \boldsymbol{w} \geq 0, \quad \text{for all } \boldsymbol{w} \in \mathcal{C}(\boldsymbol{x}^\star, \lambda^\star).$$

The corresponding Second-Order Sufficient Conditions are given below

Theorem 5.2. ***(Second-Order Sufficient Conditions**.) Suppose that for some feasible point $\boldsymbol{x}^\star \in \mathbb{R}^n$ there is a Lagrangian multipler vector $\lambda^\star$ such that the KKT conditions are satisfied. Suppose also that*

$$\boldsymbol{w}^T \nabla^2_{\boldsymbol{xx}} L(\boldsymbol{x}^\star, \lambda^\star) \boldsymbol{w} > 0, \quad \text{for all } \boldsymbol{w} \in \mathcal{C}(\boldsymbol{x}^\star, \lambda^\star) \backslash \{\boldsymbol{0}\}.$$

Then $\boldsymbol{x}^\star$ is a strict local solution for (1).