

Optimality conditions for constrained optimization

When we are solving an unconstrained optimization problem, the goal is clear: we want to find a point where the gradient vanishes. All of the algorithms we looked at over the last few lectures were in service of this condition. Once we add constraints, the optimality conditions are more complicated, and involve relationships between the gradient of the functional we are minimizing along with the gradients of the constraints — these are the so-called Karush-Kuhn-Tucker (KKT) conditions.

We will build up to the KKT conditions slowly. We will first derive a general (and very easy to prove) *geometric* necessary and sufficient condition for \mathbf{x}^* to be a minimizer of a constrained optimization program. We will then show how this simple result immediately yields the KKT conditions for certain kinds of constraints. In the next set of notes, we will derive the KKT conditions, show that they are always sufficient, and discuss conditions under which they are also necessary.

We start by considering the general constrained problem

$$\underset{\mathbf{x} \in \mathcal{C}}{\text{minimize}} \ f(\mathbf{x})$$

where \mathcal{C} is a closed, convex set, and f is again a convex function.

We have the following fundamental result:

Let f be a differentiable convex function, and \mathcal{C} be a closed convex set. Then \mathbf{x}^* is a minimizer of


$$\underset{\mathbf{x} \in \mathcal{C}}{\text{minimize}} f(\mathbf{x})$$

if and only if $\mathbf{x}^* \in \mathcal{C}$ and

$$\langle \mathbf{y} - \mathbf{x}^*, \nabla f(\mathbf{x}^*) \rangle \geq 0$$

for all $\mathbf{y} \in \mathcal{C}$.

This result is geometrically intuitive; it is saying that every vector from \mathbf{x}^* to another point \mathbf{y} in \mathcal{C} must make an **obtuse** angle with $-\nabla f(\mathbf{x}^*)$. That is, there cannot be any descent directions from \mathbf{x}^* that lead to another point in \mathcal{C} . Here is a picture:



To prove this, we first argue that $\langle \mathbf{y} - \mathbf{x}^*, \nabla f(\mathbf{x}^*) \rangle \geq 0$ for all $\mathbf{y} \in \mathcal{C}$ implies that \mathbf{x}^* is optimal. Since f is convex, for any $\mathbf{y} \in \mathcal{C}$

$$f(\mathbf{y}) \geq f(\mathbf{x}^*) + \langle \mathbf{y} - \mathbf{x}^*, \nabla f(\mathbf{x}^*) \rangle,$$

and so

$$f(\mathbf{y}) - f(\mathbf{x}^*) \geq \langle \mathbf{y} - \mathbf{x}^*, \nabla f(\mathbf{x}^*) \rangle \geq 0,$$

Since this holds for every $\mathbf{y} \in \mathcal{C}$, \mathbf{x}^* is a minimizer.

Now suppose that \mathbf{x}^* is a minimizer. If there were a $\mathbf{y} \in \mathcal{C}$ such that $\langle \mathbf{y} - \mathbf{x}^*, \nabla f(\mathbf{x}^*) \rangle < 0$, then $\mathbf{d} = \mathbf{y} - \mathbf{x}^*$ would be a descent direction, and there would exist a $0 < t < 1$ such that

$$f(\mathbf{x}^* + t(\mathbf{y} - \mathbf{x}^*)) < f(\mathbf{x}^*).$$

Since \mathcal{C} is convex and $\mathbf{x}^*, \mathbf{y} \in \mathcal{C}$, we know $\mathbf{x}^* + t(\mathbf{y} - \mathbf{x}^*) \in \mathcal{C}$. But this contradicts the assertion that \mathbf{x}^* is a minimizer, and so no such \mathbf{y} can exist.

Examples

The abstract geometrical result in the previous section will eventually lead us to the Karush-Kuhn-Tucker (KKT) conditions. But we will build up to this by looking at what it tells us in several important (and prevalent) cases.

We assume throughout this section that f is convex, differentiable, and defined on all of \mathbb{R}^N .

Linear constraints

Consider a convex optimization problem with linear¹ constraints,

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \quad f(\mathbf{x}) \quad \text{subject to} \quad \mathbf{A}\mathbf{x} = \mathbf{b},$$

where \mathbf{A} is $M \times N$ and $\mathbf{b} \in \mathbb{R}^M$. At a solution \mathbf{x}^* , we have

$$\langle \mathbf{y} - \mathbf{x}^*, \nabla f(\mathbf{x}^*) \rangle \geq 0,$$

for all \mathbf{y} such that $\mathbf{A}\mathbf{y} = \mathbf{b}$. Since $\mathbf{A}\mathbf{x}^* = \mathbf{b}$ as well, this is equivalent to

$$\langle \mathbf{h}, \nabla f(\mathbf{x}^*) \rangle \geq 0, \quad \text{for all } \mathbf{h} \in \text{Null}(\mathbf{A}).$$

Since $\mathbf{h} \in \text{Null}(\mathbf{A}) \Leftrightarrow -\mathbf{h} \in \text{Null}(\mathbf{A})$, we must have

$$\langle \mathbf{h}, \nabla f(\mathbf{x}^*) \rangle = 0, \quad \text{for all } \mathbf{h} \in \text{Null}(\mathbf{A}),$$

i.e. the gradient is **orthogonal** to the null space of \mathbf{A} . This means that it is in the row space,

$$\nabla f(\mathbf{x}^*) \in \text{Range}(\mathbf{A}^T),$$

and so there is a $\boldsymbol{\nu} \in \mathbb{R}^M$ such that

$$\nabla f(\mathbf{x}^*) + \mathbf{A}^T \boldsymbol{\nu} = \mathbf{0}.$$

¹We really should be saying *affine constraints*, but “linear constraints” is typical nomenclature for this type of problem.


Summary:

\mathbf{x}^* is a solution to

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \ f(\mathbf{x}) \quad \text{subject to} \quad \mathbf{Ax} = \mathbf{b},$$

if and only if

1. $\mathbf{Ax}^* = \mathbf{b}$, and
2. there exists $\boldsymbol{\nu}^* \in \mathbb{R}^M$ such that $\nabla f(\mathbf{x}^*) + \mathbf{A}^T \boldsymbol{\nu}^* = \mathbf{0}$.



Non-negativity constraints

Now consider the convex program

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \quad f(\mathbf{x}) \quad \text{subject to} \quad \mathbf{x} \geq \mathbf{0}.$$

At a solution \mathbf{x}^* , we will have

$$\langle \mathbf{y} - \mathbf{x}^*, \nabla f(\mathbf{x}^*) \rangle \geq 0, \quad \text{for all } \mathbf{y} \in \mathbb{R}_+^N. \quad (1)$$

Since both $\mathbf{0} \in \mathbb{R}_+^N$ and $2\mathbf{x}^* \in \mathbb{R}_+^N$, this means

$$\langle \mathbf{x}^*, \nabla f(\mathbf{x}^*) \rangle = 0, \quad (2)$$

and so

$$\langle \mathbf{y}, \nabla f(\mathbf{x}^*) \rangle \geq 0, \quad \text{for all } \mathbf{y} \in \mathbb{R}_+^N,$$

meaning that the gradient has only non-negative values as well,

$$\nabla f(\mathbf{x}^*) \geq \mathbf{0}. \quad (3)$$

The conditions (2) and (3) are sufficient as well, as together they imply (1). Condition (3) is the same as saying there exists a $\boldsymbol{\lambda}^* \geq \mathbf{0}$ such that

$$\nabla f(\mathbf{x}^*) - \boldsymbol{\lambda}^* = \mathbf{0}.$$

We can also see that (2) and (3), along with the fact that $\mathbf{x}^* \in \mathbb{R}_+^N$, mean that $\nabla f(\mathbf{x}^*)$ and \mathbf{x}^* can only be non-zero at different indices:

$$\begin{aligned} [\nabla f(\mathbf{x}^*)]_n > 0 &\Rightarrow x_n = 0, \\ x_n > 0 &\Rightarrow [\nabla f(\mathbf{x}^*)]_n = 0. \end{aligned}$$


Summary:

\mathbf{x}^* is a solution to

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \quad f(\mathbf{x}) \quad \text{subject to} \quad \mathbf{x} \geq \mathbf{0},$$

if and only if

1. $\mathbf{x}^* \geq \mathbf{0}$,
and there exists a $\boldsymbol{\lambda}^* \in \mathbb{R}^N$ such that
2. $\boldsymbol{\lambda}^* \geq \mathbf{0}$, and
3. $\lambda_n x_n = 0$ for all $n = 1, \dots, N$, and
4. $\nabla f(\mathbf{x}^*) - \boldsymbol{\lambda}^* = \mathbf{0}$.



A single convex inequality constraint

Now consider the convex program

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \quad f(\mathbf{x}) \quad \text{subject to} \quad g(\mathbf{x}) \leq 0,$$

where g is also a differentiable convex function. We will argue that in this case, the optimality conditions for \mathbf{x}^* ,

$$g(\mathbf{x}^*) \leq 0, \quad \text{and} \quad \langle \mathbf{y} - \mathbf{x}^*, \nabla f(\mathbf{x}^*) \rangle \geq 0, \quad \text{for all } \mathbf{y} \text{ with } g(\mathbf{y}) \leq 0,$$

are equivalent to one of these two conditions holding,

1. $g(\mathbf{x}^*) < 0$ and $\nabla f(\mathbf{x}^*) = \mathbf{0}$, or
2. $g(\mathbf{x}^*) = 0$ and the gradients of g and f are negatively aligned:

$$\nabla g(\mathbf{x}^*) = -\lambda \nabla f(\mathbf{x}^*), \quad \text{for some } \lambda > 0.$$

Establishing this relies on the following geometric fact²:

Let \mathbf{u}, \mathbf{v} be vectors in \mathbb{R}^N . If no \mathbf{d} exists such that

$$\langle \mathbf{d}, \mathbf{u} \rangle < 0, \quad \text{and} \quad \langle \mathbf{d}, \mathbf{v} \rangle < 0 \quad \text{simultaneously,} \quad (4)$$


then \mathbf{u} and \mathbf{v} are negatively aligned,

$$\mathbf{u} = -\lambda \mathbf{v}, \quad \text{for some } \lambda > 0. \quad (5)$$

The converse also holds, as if (5) is true, there is no way (4) can be true.

²This is a special case of the famous Gordan Theorem.

The argument for this is simple. The sets $\{\mathbf{x} : \langle \mathbf{x}, \mathbf{u} \rangle < 0\}$ and $\{\mathbf{x} : \langle \mathbf{x}, \mathbf{v} \rangle < 0\}$ are open half spaces, and these half spaces are disjoint if and only if (5) holds.



Suppose now that there is a \mathbf{x}^* such that $g(\mathbf{x}^*) = 0$ and a $\lambda > 0$ so that $\nabla g(\mathbf{x}^*) = -\lambda \nabla f(\mathbf{x}^*)$. Let \mathbf{x} be any other feasible point; $g(\mathbf{x}) \leq 0$. Then, by the convexity of g ,

$$g(\mathbf{x}^* + \theta(\mathbf{x} - \mathbf{x}^*)) \leq 0, \quad \text{for all } 0 \leq \theta \leq 1.$$

Since the above is true for *all* θ in this range, we know that $\mathbf{x} - \mathbf{x}^*$ cannot be an ascent direction for g from \mathbf{x}^* . Thus

$$\langle \mathbf{x} - \mathbf{x}^*, \nabla g(\mathbf{x}^*) \rangle \leq 0.$$

Since $\nabla g(\mathbf{x}^*) = -\lambda \nabla f(\mathbf{x}^*)$, we now know

$$\langle \mathbf{x} - \mathbf{x}^*, \nabla f(\mathbf{x}^*) \rangle \geq 0.$$

Then by the convexity of f ,

$$\begin{aligned} f(\mathbf{x}) &\geq f(\mathbf{x}^*) + \langle \mathbf{x} - \mathbf{x}^*, \nabla f(\mathbf{x}^*) \rangle \\ &\geq f(\mathbf{x}^*), \end{aligned}$$

and so \mathbf{x}^* is a minimizer.

Now suppose that \mathbf{x}^* , with $g(\mathbf{x}^*) \leq 0$, is a minimizer. We know that $\nabla g(\mathbf{x}^*)$ and $\nabla f(\mathbf{x}^*)$ must be negatively aligned, as otherwise our geometric fact dictates that there is a \mathbf{d} that is a descent direction for both g and f , meaning there is a $0 < t < 1$ such that

$$\begin{aligned} f(\mathbf{x}^* + t\mathbf{d}) &< f(\mathbf{x}^*), \text{ and} \\ g(\mathbf{x}^* + t\mathbf{d}) &< g(\mathbf{x}^*) \leq 0. \end{aligned}$$

This would mean that there is a feasible point at which f is smaller than it is at \mathbf{x}^* , directly contradicting the assertion that \mathbf{x}^* is a minimizer. Thus no such \mathbf{d} can exist.

We can collect all of this into the following summary:

\mathbf{x}^* is a solution to

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \quad f(\mathbf{x}) \quad \text{subject to} \quad g(\mathbf{x}) \leq 0,$$

if and only if

1. $g(\mathbf{x}^*) \leq 0$,
and there exists a $\lambda^* \in \mathbb{R}$ such that
2. $\lambda^* \geq 0$, and
3. $\lambda^* g(\mathbf{x}^*) = 0$, and
4. $\nabla f(\mathbf{x}^*) + \lambda^* \nabla g(\mathbf{x}^*) = \mathbf{0}$.

