

The Karush-Kuhn-Tucker (KKT) conditions

In this section, we will give a set of sufficient (and at most times necessary) conditions for a \mathbf{x}^* to be the solution of a given convex optimization problem. These are called the Karush-Kuhn-Tucker (KKT) conditions, and they play a fundamental role in both the theory and practice of convex optimization. We have derived these conditions (and have shown that they are both necessary and sufficient) in some special cases in the previous notes

We will start here by considering a general convex program with **inequality** constraints only. This is just to make the exposition easier — after we have this established, we will show how to include equality constraints (which must always be affine in convex programming). A great source for the material in this section is [Lau13, Chap. 10].

Everywhere in this section, the functionals $f_0(\mathbf{x}), f_1(\mathbf{x}), \dots, f_M(\mathbf{x}), f_m : \mathbb{R}^N \rightarrow \mathbb{R}$, are convex and differentiable.

KKT (inequality only)

The KKT conditions for the convex program

$$\begin{aligned} \underset{\mathbf{x}}{\text{minimize}} \quad & f_0(\mathbf{x}) \quad \text{subject to} \quad f_1(\mathbf{x}) \leq 0 & (1) \\ & & f_2(\mathbf{x}) \leq 0 \\ & & \vdots \\ & & f_M(\mathbf{x}) \leq 0 \end{aligned}$$

in $\mathbf{x} \in \mathbb{R}^N$ and $\boldsymbol{\lambda} \in \mathbb{R}^M$ are

$$f_m(\mathbf{x}) \leq 0, \quad m = 1, \dots, M, \quad (\text{K1})$$

$$\boldsymbol{\lambda} \geq \mathbf{0}, \quad (\text{K2})$$

$$\lambda_m f_m(\mathbf{x}) = 0, \quad m = 1, \dots, M, \quad (\text{K3})$$

$$\nabla f_0(\mathbf{x}) + \sum_{m=1}^M \lambda_m \nabla f_m(\mathbf{x}) = \mathbf{0}, \quad (\text{K4})$$

We start by establishing that these are sufficient conditions for a minimizer.

If the KKT conditions hold for \mathbf{x}^* and some $\boldsymbol{\lambda}^* \in \mathbb{R}^M$, then \mathbf{x}^* is a solution to the program (1).

Below, we denote the feasible set as

$$\mathcal{C} = \{\mathbf{x} \in \mathbb{R}^N : f_m(\mathbf{x}) \leq 0, m = 1, \dots, M\}.$$

It should be clear that the convexity of the f_m implies the convexity¹

¹The f_m are convex functions, so their sublevel sets are convex sets, and \mathcal{C}

of \mathcal{C} . The sufficiency proof simply relies on the convexity of \mathcal{C} , the convexity of f_0 , and the concept of a descent/ascent direction (see the previous notes).

Suppose \mathbf{x}^* , $\boldsymbol{\lambda}^*$ obey the KKT conditions. The first thing to note is that if

$$\lambda_1 = \lambda_2 = \dots = \lambda_M = 0,$$

then (K4) implies that

$$\nabla f(\mathbf{x}^*) = \mathbf{0},$$

and hence \mathbf{x}^* is a global min, as by the convexity of f_0 ,

$$f_0(\mathbf{x}) \geq f_0(\mathbf{x}^*) + \langle \mathbf{x} - \mathbf{x}^*, \nabla f_0(\mathbf{x}^*) \rangle = f_0(\mathbf{x}^*),$$

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

Now suppose that $R > 0$ entries of $\boldsymbol{\lambda}^*$ are positive — without loss of generality, we will take these to be the first R ,

$$\lambda_1^* > 0, \quad \lambda_2^* > 0, \quad \dots, \quad \lambda_R^* > 0, \quad \lambda_{R+1}^* = 0, \quad \dots, \quad \lambda_M^* = 0.$$

We can rewrite (K4) as

$$\nabla f_0(\mathbf{x}^*) + \lambda_1^* \nabla f_1(\mathbf{x}^*) + \dots + \lambda_R^* \nabla f_R(\mathbf{x}^*) = \mathbf{0}, \quad (2)$$

and note that by (K3),

$$f_1(\mathbf{x}^*) = 0, \dots, f_R(\mathbf{x}^*) = 0.$$

Consider any $\mathbf{x} \in \mathcal{C}$, $\mathbf{x} \neq \mathbf{x}^*$. As \mathcal{C} is convex, every point in between \mathbf{x}^* and \mathbf{x} must also be in \mathcal{C} , meaning

$$f_m(\mathbf{x}^* + \theta(\mathbf{x} - \mathbf{x}^*)) \leq 0 = f_m(\mathbf{x}^*), \quad m = 1, \dots, R,$$

is an intersection of sublevel sets.

for all $0 \leq \theta \leq 1$. This means that $\mathbf{x} - \mathbf{x}^*$ cannot be an ascent direction, and so

$$\langle \mathbf{x} - \mathbf{x}^*, \nabla f_m(\mathbf{x}^*) \rangle \leq 0, \quad m = 1, \dots, R.$$

It is clear, then, that

$$\langle \mathbf{x} - \mathbf{x}^*, \nabla f_0(\mathbf{x}^*) \rangle \geq 0,$$

as otherwise there is no way (2) can hold with positive λ_m . Along with the convexity of f_0 , this means that

$$f_0(\mathbf{x}) \geq f_0(\mathbf{x}^*) + \langle \mathbf{x} - \mathbf{x}^*, \nabla f(\mathbf{x}^*) \rangle \geq f(\mathbf{x}^*).$$

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

Necessity

To establish the necessity of the KKT conditions, we need one piece of mathematical technology that we have not been exposed to yet. The *Farkas lemma* is a fundamental result in convex analysis; we will prove it in the Technical Details section.

Farkas Lemma:

Let \mathbf{A} be an $M \times N$ matrix and $\mathbf{b} \in \mathbb{R}^M$. The exactly one of the following two things is true:

1. there exists $\mathbf{x} \geq \mathbf{0}$ such that $\mathbf{A}\mathbf{x} = \mathbf{b}$;
2. there exists $\boldsymbol{\lambda} \in \mathbb{R}^M$ such that

$$\mathbf{A}^T \boldsymbol{\lambda} \leq \mathbf{0}, \quad \text{and} \quad \langle \mathbf{b}, \boldsymbol{\lambda} \rangle > 0.$$

With this in place, we can give two different situations under which KKT is necessary. These are by no means the only situations for which this is true, but these two cover a high percentage of the cases encountered in practice.

Suppose \mathbf{x}^* is a solution to a convex program with affine inequality constraints:

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \quad f_0(\mathbf{x}) \quad \text{subject to} \quad \mathbf{A}\mathbf{x} \leq \mathbf{b}.$$

Then there exists a $\boldsymbol{\lambda}^*$ such that \mathbf{x}^* , $\boldsymbol{\lambda}^*$ obey the KKT conditions.

In this case, the constraint functions have the form

$$f_m(\mathbf{x}) = \langle \mathbf{x}, \mathbf{a}_m \rangle - b_m, \quad \text{and so} \quad \nabla f_m(\mathbf{x}) = \mathbf{a}_m,$$

where \mathbf{a}_m^T is the m th row of \mathbf{A} . Since \mathbf{x}^* is feasible, K1 must hold. If none of the constraints are “active”, meaning $f_m(\mathbf{x}^*) < 0$ for $m = 1, \dots, M$ and \mathbf{x}^* lies in the interior of \mathcal{C} , then it must be that $\nabla f(\mathbf{x}^*) = \mathbf{0}$, and K2–K4 hold with $\boldsymbol{\lambda} = \mathbf{0}$.

Suppose that there are R active constraints at \mathbf{x}^* ; without loss of generality, we will take these to be the first R :

$$\begin{aligned} f_1(\mathbf{x}^*) = 0, \quad f_2(\mathbf{x}^*) = 0, \quad \dots, \quad f_R(\mathbf{x}^*) = 0, \\ f_{R+1}(\mathbf{x}^*) < 0, \quad \dots, \quad f_M(\mathbf{x}^*) < 0. \end{aligned}$$

We start by taking $\lambda_{R+1} = \lambda_{R+2} = \dots = \lambda_M = 0$, which means K3 will hold. Suppose that there were no $\boldsymbol{\lambda} \geq \mathbf{0}$ such that

$$\nabla f_0(\mathbf{x}^*) + \lambda_1 \nabla f_1(\mathbf{x}^*) + \dots + \lambda_R \nabla f_R(\mathbf{x}^*) = \mathbf{0}. \quad (3)$$

With $\mathbf{A}' : R \times N$ consisting of the first R rows of \mathbf{A} , and $\mathbf{b}' \in \mathbb{R}^R$ as the first R entries in \mathbf{b} , this means that there is no $\boldsymbol{\lambda}' \in \mathbb{R}^R$ such that

$$\mathbf{A}'^T \boldsymbol{\lambda}' = -\nabla f_0(\mathbf{x}^*), \quad \boldsymbol{\lambda}' \geq \mathbf{0}.$$

By the Farkas lemma, this means that there is a $\mathbf{d} \in \mathbb{R}^N$ such that

$$\mathbf{A}' \mathbf{d} \leq \mathbf{0}, \quad \langle \mathbf{d}, -\nabla f_0(\mathbf{x}^*) \rangle > 0,$$

which means, since $\nabla f_m(\mathbf{x}) = \mathbf{a}_m$,

$$\begin{aligned} \langle \mathbf{d}, \nabla f_0(\mathbf{x}^*) \rangle &< 0 \\ \langle \mathbf{d}, \nabla f_1(\mathbf{x}^*) \rangle &\leq 0 \\ &\vdots \\ \langle \mathbf{d}, \nabla f_R(\mathbf{x}^*) \rangle &\leq 0. \end{aligned}$$

This means that \mathbf{d} is a descent direction for f_0 , and is not an ascent direction for f_1, \dots, f_R . Because the constraint functionals are affine, if $\langle \mathbf{d}, \nabla f_m(\mathbf{x}^*) \rangle = 0$ above, then $f_m(\mathbf{x}^* + t\mathbf{d}) = f_m(\mathbf{x}^*)$ — this means that moving in the direction \mathbf{d} will not increase f_1, \dots, f_m . Since the last $M - R$ constraints are not active, we can move at least a small amount in any direction so that they stay that way. This means that there exists a $t > 0$ such that

$$f_0(\mathbf{x}^* + t\mathbf{d}) < f_0(\mathbf{x}^*),$$

but also maintains feasibility:

$$f_m(\mathbf{x}^* + t\mathbf{d}) \leq 0, \quad m = 1, \dots, M.$$

This directly contradicts the assertion that \mathbf{x}^* is optimal, and so $\lambda_1, \dots, \lambda_R \geq 0$ must exist such that (3) holds.

For general convex inequality constraints, there are various other scenarios under which the KKT conditions are necessary; these are

called **constraint qualifications**. We have already seen that polygonal (affine) constraints qualify. Another set of constraint qualifications are *Slater's condition*:

Slater's condition: There exists at least one strictly feasible point; a \mathbf{x} such that none of the constraints are active:

$$f_1(\mathbf{x}) < 0, f_2(\mathbf{x}) < 0, \dots, f_M(\mathbf{x}) < 0.$$

Suppose that Slater's condition holds for f_1, \dots, f_M , and let \mathbf{x}^* be a solution to

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} f_0(\mathbf{x}) \quad \text{subject to} \quad f_m \leq 0, \quad m = 1, \dots, M.$$

Then there exists a $\boldsymbol{\lambda}^*$ such that $\mathbf{x}^*, \boldsymbol{\lambda}^*$ obey the KKT conditions.

This is proved in much the same way as in the affine inequality case. Suppose that \mathbf{x}^* is a solution, and that

$$\begin{aligned} f_1(\mathbf{x}^*) = 0, f_2(\mathbf{x}^*) = 0, \dots, f_R(\mathbf{x}^*) = 0, \\ f_{R+1}(\mathbf{x}^*) < 0, \dots, f_M(\mathbf{x}^*) < 0. \end{aligned}$$

We take $\lambda_{R+1} = \dots = \lambda_M = 0$, and show that if there is not $\lambda_1, \dots, \lambda_R \geq 0$ such that

$$\nabla f_0(\mathbf{x}^*) + \sum_{m=1}^R \lambda_m \nabla f_m(\mathbf{x}^*) = \mathbf{0}, \quad (4)$$

then there is another feasible point with a smaller value of f_0 .

By the Farkas lemma, if there does not exist a $\lambda_1, \dots, \lambda_R \geq 0$ such that (4) holds, then there must be a $\mathbf{u} \in \mathbb{R}^N$ such that

$$\begin{aligned} \langle \mathbf{u}, \nabla f_0(\mathbf{x}^*) \rangle &< 0 \\ \langle \mathbf{u}, \nabla f_1(\mathbf{x}^*) \rangle &\leq 0 \\ &\vdots \\ \langle \mathbf{u}, \nabla f_R(\mathbf{x}^*) \rangle &\leq 0. \end{aligned}$$

Now let \mathbf{z} be a strictly feasible point, $f_m(\mathbf{z}) < 0$ for all m . We know that

$$0 > f_m(\mathbf{z}) \geq f_m(\mathbf{x}^*) + \langle \mathbf{z} - \mathbf{x}^*, \nabla f_m(\mathbf{x}^*) \rangle \quad \Rightarrow \quad \langle \mathbf{z} - \mathbf{x}^*, \nabla f_m(\mathbf{x}^*) \rangle < 0,$$

for $m = 1, \dots, R$, since then $f_m(\mathbf{x}^*) = 0$. So \mathbf{u} is a descent direction for f_0 , and $\mathbf{z} - \mathbf{x}^*$ is a descent direction for all all of the constraint functionals f_m , $m = 1, \dots, R$ that are active.

We consider a convex combination of these two vectors

$$\mathbf{d}_\theta = (1 - \theta)\mathbf{u} + \theta(\mathbf{z} - \mathbf{x}^*).$$

We know that $\langle \mathbf{d}_\theta, \nabla f_m(\mathbf{x}^*) \rangle < 0$ for all $0 < \theta \leq 1$, $m = 1, \dots, R$. We also know that there is a θ small enough so that \mathbf{d}_θ is a descent direction for f_0 ; there exists $0 < \epsilon_0 < 1$ such that

$$\langle \mathbf{d}_{\epsilon_0}, \nabla f_0(\mathbf{x}^*) \rangle < 0.$$

Finally, we also know that we can move a small enough amount in any direction and keep constraints f_{R+1}, \dots, f_M inactive. Thus there is a $t > 0$ such that

$$f_0(\mathbf{x}^* + t\mathbf{d}_{\epsilon_0}) < f_0(\mathbf{x}^*), \quad f_m(\mathbf{x}^* + t\mathbf{d}_{\epsilon_0}) \leq 0, \quad m = 1, \dots, M,$$

which directly contradicts the assertion that \mathbf{x}^* is optimal.

It should be clear from the two arguments above that Slater's condition can be refined — we only need a point which obeys $f_m(\mathbf{z}) < 0$ for the f_m which are not affine. We now state this formally:

Suppose that $f_1, \dots, f_{M'}$ are affine functionals, and $f_{M'+1}, \dots, f_M$ are convex functional which are not affine. Suppose that Slater's condition holds for $f_{M'+1}, \dots, f_M$, and let \mathbf{x}^* be a solution to

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \quad f_0(\mathbf{x}) \quad \text{subject to} \quad f_m(\mathbf{x}) \leq 0, \quad m = 1, \dots, M.$$

Then there exists a $\boldsymbol{\lambda}^*$ such that $\mathbf{x}^*, \boldsymbol{\lambda}^*$ obey the KKT conditions.

The above statement lets us extend the KKT conditions to optimization problems with linear equality constraints, which we now state.

KKT (with equality constraints)

The KKT conditions for the optimization program

$$\begin{aligned} \min_{\mathbf{x}} f_0(\mathbf{x}) \quad \text{subject to} \quad & f_m(\mathbf{x}) \leq 0, \quad m = 1, \dots, M \quad (5) \\ & h_p(\mathbf{x}) = 0, \quad p = 1, \dots, P \end{aligned}$$

in $\mathbf{x} \in \mathbb{R}^N$, $\boldsymbol{\lambda} \in \mathbb{R}^M$, and $\boldsymbol{\nu} \in \mathbb{R}^P$ are

$$f_m(\mathbf{x}) \leq 0, \quad m = 1, \dots, M, \quad (\text{K1})$$

$$h_p(\mathbf{x}) = 0, \quad p = 1, \dots, P$$

$$\boldsymbol{\lambda} \geq \mathbf{0}, \quad (\text{K2})$$

$$\lambda_m f_m(\mathbf{x}) = 0, \quad m = 1, \dots, M, \quad (\text{K3})$$

$$\nabla f_0(\mathbf{x}) + \sum_{m=1}^M \lambda_m \nabla f_m(\mathbf{x}) + \sum_{p=1}^P \nu_p \nabla h_p(\mathbf{x}) = \mathbf{0}, \quad (\text{K4})$$

We call the $\boldsymbol{\lambda}$ and $\boldsymbol{\nu}$ above **Lagrange multipliers**. Notice that $\boldsymbol{\lambda}$ is constrained to be positive, while $\boldsymbol{\nu}$ can be arbitrary. Also, if the h_p are affine, which they have to be for the program above to be convex, then we can write the equality constraints

$$h_p(\mathbf{x}) = 0, \quad p = 1, \dots, P \quad \text{as} \quad \mathbf{A}\mathbf{x} = \mathbf{b},$$

for some $\mathbf{A} : P \times N$ and $\mathbf{b} \in \mathbb{R}^P$. Also, we can rewrite (K4) as

$$\nabla f_0(\mathbf{x}) + \sum_{m=1}^M \lambda_m \nabla f_m(\mathbf{x}) + \mathbf{A}^T \boldsymbol{\nu} = \mathbf{0}.$$

If the f_m are convex and the h_p affine, then the KKT conditions are sufficient for \mathbf{x}^* to be the solution to the convex program (5).

If Slater's condition holds for the non-affine f_m , then they are also necessary. Almost nothing changes in the proofs above — we could simply separate an equality constraint of the form $\langle \mathbf{x}, \mathbf{a} \rangle = b$ into $\langle \mathbf{x}, \mathbf{a} \rangle - b \leq 0$ and $\langle \mathbf{x}, -\mathbf{a} \rangle + b \leq 0$. Then we can recombine the result, taking $nu = \lambda_1 - \lambda_2$, where λ_1 is the Lagrange multiplier for $\langle \mathbf{x}, \mathbf{a} \rangle - b$ and λ_2 is the same for $\langle \mathbf{x}, -\mathbf{a} \rangle + b$.

Technical Details: Proof of the Farkas Lemma

We prove the Farkas Lemma: if \mathbf{A} is an $M \times N$ matrix and $\mathbf{b} \in \mathbb{R}^M$ is a given vector, then exactly one of the following two things is true:

1. there exists $\mathbf{x} \geq \mathbf{0}$ such that $\mathbf{Ax} = \mathbf{b}$;
2. there exists $\mathbf{v} \in \mathbb{R}^M$ such that

$$\mathbf{A}^T \mathbf{v} \leq \mathbf{0}, \quad \text{and} \quad \langle \mathbf{b}, \mathbf{v} \rangle > 0.$$

It is clear that if the first condition holds, the second cannot, as $\langle \mathbf{b}, \mathbf{v} \rangle = \langle \mathbf{x}, \mathbf{A}^T \mathbf{v} \rangle$ for any \mathbf{x} such that $\mathbf{Ax} = \mathbf{b}$, and $\langle \mathbf{x}, \mathbf{A}^T \mathbf{v} \rangle \leq 0$ for any $\mathbf{x} \geq \mathbf{0}$ and $\mathbf{A}^T \mathbf{v} \leq \mathbf{0}$.

It is more difficult to argue that if the first condition does not hold, the second must. This ends up being a direct result of the separating hyperplane theorem. Let $\mathcal{C}(\mathbf{A})$ be the (convex) cone generated by the columns $\mathbf{a}_1, \dots, \mathbf{a}_N$ of \mathbf{A} :

$$\mathcal{C}(\mathbf{A}) = \left\{ \mathbf{v} \in \mathbb{R}^M : \mathbf{v} = \sum_{n=1}^N \theta_n \mathbf{a}_n, \quad \theta_n \geq 0, \quad n = 1, \dots, N \right\}.$$

Then 1 above is clearly equivalent to $\mathbf{b} \in \mathcal{C}(\mathbf{A})$. Since $\mathcal{C}(\mathbf{A})$ is closed and convex, and \mathbf{b} is a single point, we know that if $\mathbf{b} \notin \mathcal{C}(\mathbf{A})$, then $\mathcal{C}(\mathbf{A})$ and \mathbf{b} are strongly separated by a hyperplane. That is, if $\mathbf{b} \notin \mathcal{C}(\mathbf{A})$ implies that there exists a $\mathbf{v} \in \mathbb{R}^M$ such that

$$\mathbf{v}^T \mathbf{b} > \mathbf{v}^T \boldsymbol{\lambda} \quad \text{for all} \quad \boldsymbol{\lambda} \in \mathcal{C}(\mathbf{A}),$$

which is the same as saying

$$\mathbf{v}^T \mathbf{b} > \sup_{\boldsymbol{\lambda} \in \mathcal{C}(\mathbf{A})} \mathbf{v}^T \boldsymbol{\lambda} = \sup_{\mathbf{x} \geq \mathbf{0}} \mathbf{v}^T \mathbf{Ax}.$$

We know that $\mathbf{0} \in \mathcal{C}(\mathbf{A})$, so we must have $\mathbf{v}^\top \mathbf{b} > 0$. The above equation also gives a finite upper bound (namely whatever the actual value of $\mathbf{v}^\top \mathbf{b}$ is) on the function $\mathbf{v}^\top \mathbf{A} \mathbf{x}$ for all $\mathbf{x} \geq \mathbf{0}$. But this means that $\mathbf{A}^\top \mathbf{v} \leq \mathbf{0}$, as otherwise we would have the following contradiction. If there were some index n such that $(\mathbf{A}^\top \mathbf{v})[n] = \epsilon > 0$, then with $\mathbf{e}_n \geq \mathbf{0}$ as the unit vector

$$\mathbf{e}_n[k] = \begin{cases} 1, & k = n, \\ 0, & k \neq n \end{cases},$$

we have

$$\sup_{\mathbf{x} \geq \mathbf{0}} \mathbf{v}^\top \mathbf{A} \mathbf{x} \geq \sup_{\alpha \geq 0} \mathbf{v}^\top \mathbf{A} (\alpha \mathbf{e}_n) = \sup_{\alpha \geq 0} \alpha \epsilon = \infty,$$

which contradicts the existence of this upper bound.

References

- [Lau13] N. Lauritzen. *Undergraduate Convexity*. World Scientific, 2013.