MATH-329 Nonlinear optimization
Exercise session 8: Convex optimization

Instructor: Nicolas Boumal
TAs: Andreea Musat, Andrew McRae

Document compiled on November 4, 2024

1. Sufficient optimality conditions. Mind the following: in Theorem 9.2, it is important to check *all three* of the following boxes:

(a) We are *minimizing* (and not maximizing).

(b) The cost function f is convex.

(c) The search space S is convex.

To really appreciate this fact, do the following:

1. Give examples of optimization problems which check two of the above but not all three of the above, and for which there exists a non-optimal stationary point.
2. If given a maximization problem, explain how you can get an equivalent minimization problem.
3. If the cost function is not convex, explain how you can get an equivalent problem with a convex cost function (you can even make the cost function linear).
4. If the constraint set is not convex, explain how you can get an equivalent problem with a convex search space (you can even make the problem unconstrained)—you will need to make the cost function a bit weird for this though (hint: "indicator function" with values in $\{0, \infty\}$).

For each of the above, explain how we should understand Theorem 9.2 against the modified problem, specifically to verify that, sadly, there is no free lunch.

Answer.

1.
   - Let $f_a(x) = x^2$ that we want to maximize on $S = [-1, 1]$. The function f and the set S are convex. However the point $x = 0$ satisfies the first-order stationarity condition $\nabla f(x) = 0$ and is not a global (nor local) maximum.
   - Let $f_b(x) = x^4 + x^3 - x^2$ on $S = \mathbb{R}$ be a function that we want to minimize. It is non-convex and has a local maximum at $x = 0$ that satisfies the first-order stationarity condition ($\nabla f(x) = 0$).
   - Let $f_c(x) = x^2$ that we want to minimize on $S = [0, 1] \cup [2, 3]$. The point $x = 2$ is first-order stationary but is not a global minimum.

2. If we have the maximization problem
$$\max_{x \in S} g(x)$$
we can simply let $f(x) = -g(x)$ and minimize f on S, that is, solve
$$\min_{x \in S} f(x).$$
In order to apply Theorem 9.2 we need the new function f to be convex.

3. Let $g\colon S' \to \mathbb{R}$ be a potentially non-convex cost function that we wish to minimize on S'. We define $f(t, x) = t$ and $S = \{(t, x) \in \mathbb{R} \times S' \mid t \geq g(x)\}$. Then minimizing g on S' is equivalent to solve
$$\min_{(t,x) \in S} f(t, x).$$
In order to apply Theorem 9.2 we need the new set S to be convex.

4. Let $g\colon S \to \mathbb{R}$ be a function that we want to minimize on a potentially non-convex set $S \subseteq \mathcal{E}$. We define
$$f(x) = \begin{cases} g(x) & \text{if } x \in S, \\ +\infty & \text{if } x \notin S. \end{cases}$$
Then minimizing g on S is equivalent to minimizing f on $\mathcal{E}$. The function f is in general not convex so we cannot apply Theorem 9.2.

$\blacksquare$

2. Discontinuous projection. Show with a drawing that Proj_S may be discontinuous if S is non-empty and closed but fails to be convex. This reveals why the PGD iteration map $x \mapsto \mathrm{Proj}_S(x - \alpha \nabla f(x))$ could be discontinuous if S is not convex. It would be much harder to analyze the algorithm if we allowed that to happen.

Answer. Issues arise when the projection is not unique. In those cases even if we arbitrarily select one of the projections to ensure that Proj_S is single-valued we realize that it is impossible to make that choice such that Proj_S is continuous. Many examples are possible. Here is one: let S be a circle: $S = \{x \in \mathbb{R}^2 : \|x\| = 1\}$. Then, it is easy to show that (work out the details)
$$\mathrm{Proj}_S(z) = \begin{cases} \frac{1}{\|z\|} z & \text{if } z \neq 0, \\ S & \text{if } z = 0. \end{cases}$$
We see that the projection of the origin to S is not unique. We could assign some arbitrary choice to define $\mathrm{Proj}_S(0)$, but whatever we choose the resulting Proj_S is discontinuous. $\blacksquare$

3. Trust-region subproblem. Let $S = \{x \in \mathbb{R}^n : \|x\| \leq 1\}$ be the unit norm ball. Give an expression for Proj_S. Given a symmetric matrix A of size n and $b \in \mathbb{R}^n$, let $f(x) = \frac{1}{2}x^\top A x + b^\top x$. What is the Lipschitz constant of ∇f? Can you easily compute an upper bound for it? Implement projected gradient descent (see lecture notes) for this problem, with a proper choice of step-size. Do you expect that this method would converge to a global minimizer? Note: this is a fairly terrible algorithm for the trust-region subproblem but it has the merit of being simple.

Answer. It's easy to prove the following formula (work out the details):

$$\mathrm{Proj}_S(z) = \frac{1}{\max(1, \|z\|)} z.$$

The Hessian of f is $\nabla^2 f(x) = A$, so that ∇f is L-Lipschitz continuous with $L = \|A\|$ (the operator norm). It is a well-known fact that the Frobenius norm of A is an upper-bound for A: that's cheap to compute. Here is Matlab code for PGD:

```
% Generate a random symmetric matrix A.
n = 10;
A = randn(n);
A = A+A';

% Force A to be positive definite to see the effect.
% A = A + n*eye(n);

% This is the projector to the unit-norm ball.
Proj = @(z) z / max(1, norm(z));

% Use norm(A, 'F') as an upper bound on norm(A, 2) which is the
% Lipschitz constant of the gradient of .5*x'*A*x.
alpha = 1/norm(A, 'F');

x = randn(n, 1); % random initialization
for k = 1 : 1000 % simplistic stopping criterion
x = Proj(x - alpha*A*x);
end

x'*x
x'*A*x
min(eig(A))
```

The behavior of the algorithm depends on A. If A is positive semidefinite, then f is convex and the algorithm will converge to a global minimum. If A is not positive semidefinite then the problem is not convex: there may be local minima. See https://epubs.siam.org/doi/pdf/10.1137/0804009 if you are interested in what happens in this case. ■

4. Image and inverse image of affine function. Let $f : \mathbb{R}^n \to \mathbb{R}^m$ be an affine function, that is, there exists $A \in \mathbb{R}^{m\times n}, b \in \mathbb{R}^m$ such that $f(x) = Ax + b$ for all $x \in \mathbb{R}^n$.

1. Let $S \subseteq \mathbb{R}^n$ be a convex set. Show that the image of S under f,

$$f(S) = \{f(x) \mid x \in S\},$$

is convex.

2. Let $S \subseteq \mathbb{R}^m$ be a convex set. Show that the inverse image of S under f,

$$f^{-1}(S) = \{x \in \mathbb{R}^n \mid f(x) \in S\},$$

is convex.

3. Let $g : \mathbb{R}^m \to \mathbb{R}^n$ be a convex function. Show that $g \circ f$ is convex.

Answer.

1. Let $S \subseteq \mathbb{R}^n$ be a convex set. Let $u, v \in f(S)$ and $t \in [0,1]$. There exist $x, y \in \mathbb{R}^n$ such that $u = Ax + b$ and $v = Ay + b$. We have $(1-t)u + tv = A((1-t)x + ty) + b = f((1-t)x + ty)$. Since $(1-t)x + ty \in S$ (convexity of S) we find that $(1-t)u + tv \in f(S)$ and $f(S)$ is convex.

2. Let $S \subseteq \mathbb{R}^m$ be a convex set. Let $x, y \in f^{-1}(S)$ and $t \in [0,1]$. The images of x and y, $Ax + b$ and $Ay + b$, are both in S. As S is convex this implies that

$$(1-t)(Ax + b) + t(Ay + b) = A((1-t)x + ty) + b$$

is in S. We conclude that $(1-t)x + ty \in f^{-1}(S)$ and $f^{-1}(S)$ is convex.

3. Let $x, y \in \mathbb{R}^n$ and $t \in [0,1]$. Then

$$\begin{aligned} g(f((1-t)x + ty)) &= g((1-t)f(x) + tf(y)) \\ &\leq (1-t)g(f(x)) + tg(f(y)). \end{aligned}$$

■

Supplementary exercises

1. Convex combination. Let $C \subseteq \mathbb{R}^n$ be a convex set, $x_1, \dots, x_k \in C$ and $\theta_1, \dots, \theta_k \geq 0$ be non-negative coefficients such that $\theta_1 + \cdots + \theta_k = 1$. Show that the convex combination $\theta_1 x_1 + \cdots + \theta_k x_k$ is in C.

Answer. We proceed by induction. The result holds for $k = 2$ by definition of convexity (and also for $k = 1$). Let $k \geq 1$ such that the result holds. Let $x_1, \dots, x_{k+1} \in C$ and $\theta_1, \dots, \theta_{k+1} \geq 0$ such that $\theta_1 + \cdots + \theta_{k+1} = 1$. We let $\bar{\theta} = \theta_1 + \cdots + \theta_k$. Then we have

$$\theta_1 x_1 + \cdots + \theta_{k+1} x_{k+1} = \bar{\theta}\left(\frac{\theta_1}{\bar{\theta}} x_1 + \cdots + \frac{\theta_k}{\bar{\theta}} x_k\right) + \theta_{k+1} x_{k+1}.$$

We observe that

$$\frac{\theta_1}{\bar{\theta}} + \cdots + \frac{\theta_k}{\bar{\theta}} = 1$$

which implies that

$$\frac{\theta_1}{\bar{\theta}} x_1 + \cdots + \frac{\theta_k}{\bar{\theta}} x_k \in C$$

by induction hypothesis. Moreover $\bar{\theta} + \theta_{k+1} = 1$ so $\theta_1 x_1 + \cdots + \theta_{k+1} x_{k+1} \in C$. ■

2. Intersection with a line. Show that a set is convex if and only if its intersection with any line is convex.

Answer. Let $S \subseteq \mathbb{R}^n$ be a convex set. The intersection of two convex sets is convex. Lines are convex so the intersection of S with a line is convex.

Conversely suppose the intersection of $S \subseteq \mathbb{R}^n$ with any line is convex. Let $x, y \in S$ be two distinct points. The intersection of S with the line going through x and y is convex. Therefore any convex combination of x and y is in S. $\blacksquare$

3. Sublevel sets. Let $f\colon \mathbb{R}^n \to \mathbb{R}$ be a convex function. Show that for all $\alpha \in \mathbb{R}$ the sublevel set $\{x \in \mathbb{R}^n \mid f(x) \leq \alpha\}$ is convex.

Answer. Let S be the α-sublevel set. Let $x, y \in S$ and $t \in [0, 1]$. Then

$$\begin{aligned} f((1-t)x + ty) &\leq (1-t)f(x) + tf(y) \\ &\leq \alpha. \end{aligned}$$

So $(1-t)x + ty \in S$ and S is convex. $\blacksquare$