Math 484: Nonlinear Programming[1] Mikhail Lavrov

Chapter 6, Lecture 2: Guarantees on the Penalty Method

April 5, 2019 *University of Illinois at Urbana-Champaign*

1 Notation and terminology

In the previous lecture, we converted the problem

$$(P) \qquad \begin{cases} \underset{\mathbf{x}\in\mathbb{R}^n}{\text{minimize}} & f(\mathbf{x}) \\ \text{subject to} & \mathbf{g}(\mathbf{x}) \le \mathbf{0} \end{cases}$$

into an unconstrained optimization problem: minimizing the modified objective function

$$F_k(\mathbf{x}) = f(\mathbf{x}) + k \cdot \left[(g_1^+(\mathbf{x}))^2 + (g_2^+(\mathbf{x}))^2 + \cdots + (g_m^+(\mathbf{x}))^2\right]$$

over all $\mathbf{x} \in \mathbb{R}^n$, where k is some large number.

We'll need some notation and terminology to talk about this precisely. I have already introduced one new bit of notation: we put a subscript F_k indicating the dependence on k, because we will need to consider instances of this function for different values of k.

- We call k the *penalty factor*. We always assume that $k > 0$, but you should imagine k to be much larger than that.
- The sum $(g_1^+(\mathbf{x}))^2 + (g_2^+(\mathbf{x}))^2 + \cdots + (g_m^+(\mathbf{x}))^2$ is called the *badness* of $\mathbf{x}$: it measures by how much $\mathbf{x}$ violates the constraints.
- Together, $k \cdot \left[(g_1^+(\mathbf{x}))^2 + (g_2^+(\mathbf{x}))^2 + \cdots + (g_m^+(\mathbf{x}))^2\right]$ is called the *penalty* or *penalty term*.
- If a global minimizer of F_k turns out to exist, we will write $\mathbf{x}^*(k)$ to denote that global minimizer.

2 Guarantees assuming convergence

We work with problems satisfying the following three assumptions, which are all reasonably general and nearly impossible to do without.

- The problem P is feasible: there exists at least one point $\mathbf{x}^{(0)} \in \mathbb{R}^n$ which satisfies $g_i(\mathbf{x}^{(0)}) \le 0$ for all i. If this does not hold, then it's impossible to say anything about how F_k behaves, because there's no actual solution for us to approach.
- For all sufficiently large k, the function F_k has a global minimizer $\mathbf{x}^*(k)$. Again, if this does not hold, we wouldn't have anything to talk about, because the penalty method won't work.

[1] This document comes from the Math 484 course webpage: https://faculty.math.illinois.edu/~mlavrov/courses/484-spring-2019.html

Technically, it would be enough here to assume that $\mathbf{x}^*(k)$ exists for infinitely many values of k, and only focus on the values of k for which this is true. But this is usually pointless generality: it's hard to come up with a natural optimization problem for which this is true and yet $\mathbf{x}^*(k)$ doesn't exist for all sufficiently large k.

- The functions $f, g_1, g_2, \dots, g_m$ are all continuous. This is the weakest condition we could ask for, and we need to ask for it once we start talking about limits, because otherwise we don't know if the limits exist.

When the penalty method works, with the Courant–Beltrami penalty function, it works not by giving us the right answer for some sufficiently large k, but by giving us the right answer in the limit. That is, as $k \to \infty$, $\mathbf{x}^*(k) \to \mathbf{x}^*$ for some $\mathbf{x}^* \in \mathbb{R}^n$.

We're not guaranteed convergence. However, if we do have convergence, we know that the limit $\mathbf{x}^*$ is the optimal solution!

Theorem 2.1. *As indicated above, suppose that P is feasible (there exists some $\mathbf{x}^{(0)} \in \mathbb{R}^n$ satisfying $\mathbf{g}(\mathbf{x}^{(0)}) \le \mathbf{0}$), $f, g_1, g_2, \dots, g_m$ are all continuous, and for all sufficiently large k, at least one global minimizer $\mathbf{x}^*(k)$ of F_k exists.*

Additionally, suppose that as $k \to \infty$, $\mathbf{x}^*(k) \to \mathbf{x}^*$ for some $\mathbf{x}^* \in \mathbb{R}^n$. (If there are multiple global minimizers of F_k, we assume that $\mathbf{x}^*(k)$ is one of them, chosen to have this convergence.)

Then $\mathbf{x}^*$ is an optimal solution of P.

Proof. Continuity of $f, g_1, g_2, \dots, g_m$ lets us know that lots of limits have the expected behavior. As $k \to \infty$ and $\mathbf{x}^*(k) \to \mathbf{x}^*$, we know that $f(\mathbf{x}^*(k)) \to f(\mathbf{x}^*)$ and that $g_i(\mathbf{x}^*(k)) \to g_i(\mathbf{x}^*)$ for all i. The function $t \mapsto \max\{0, t\}^2$ is continuous; therefore $g_i^+(\mathbf{x}^*(k))^2 \to g_i^+(\mathbf{x}^*)^2$ for all i, as well. By the sum of limits, the badness converges:

$$\lim_{k\to\infty} \sum_{i=1}^m (g_i^+(\mathbf{x}^*(k)))^2 = \sum_{i=1}^m (g_i^+(\mathbf{x}^*))^2.$$

What about the limit of F_k? Here, we run into trouble. There are two possibilities:

- If the badness $\sum_{i=1}^m (g_i^+(\mathbf{x}^*))^2$ is nonzero (that is, $\mathbf{x}^*$ is not feasible for P) then $F_k(\mathbf{x}^*(k)) \to \infty$ as $k \to \infty$: we have

$$\lim_{k\to\infty} F_k(\mathbf{x}^*(k)) = f(\mathbf{x}^*) + \lim_{k\to\infty} \left(k \cdot \sum_{i=1}^m (g_i^+(\mathbf{x}^*(k)))^2 \right)$$

  and the limit of the penalty term is the product of a penalty factor (which goes to infinity) and a badness (which approaches a nonzero constant).

- If the badness $\sum_{i=1}^m (g_i^+(\mathbf{x}^*))^2$ is zero (that is, $\mathbf{x}^*$ is feasible for P) then the limit of the penalty term has the undefined form $\infty \cdot 0$. In this case, we can't predict the behavior of the limit of F_k, at least not using this approach!

In general, either of these cases could occur. However, we have made an additional assumption changing the situation: that there is a point $\mathbf{x}^{(0)}$ which is feasible for P.

This point is not necessarily great for optimization; $f(\mathbf{x}^{(0)})$ could be very large. But when applying the penalty method, $F_k(\mathbf{x}^{(0)}) = f(\mathbf{x}^{(0)})$ for all k (the badness of $\mathbf{x}^{(0)}$ is zero) which is a constant independent of k.

This prevents things from getting too bad as $k \to \infty$. It is impossible that the optimal value of F_k diverges as $k \to \infty$, because the optimal value should always be at least as good as $f(\mathbf{x}^{(0)})$. So among the two possibilities above, we know we're not in the first case (where F_k diverges) and must be in the second case (where we couldn't determine the limit of F_k directly).

This proves that the badness of $\mathbf{x}^*$ is 0: that $\mathbf{x}^*$ is a feasible point of P.

To prove that it's optimal, let $\mathbf{y}$ be any other feasible point of P. Then for each k, we have

$$f(\mathbf{y}) = F_k(\mathbf{y}) \geq F_k(\mathbf{x}^*(k)) \geq f(\mathbf{x}^*(k)).$$

To unpack this: $f(\mathbf{y}) = F_k(\mathbf{y})$ because $\mathbf{y}$ is feasible, so it has no penalty. $F_k(\mathbf{y}) \geq F_k(\mathbf{x}^*(k))$ because $\mathbf{x}^*(k)$ is a global minimizer of F_k. Finally, $F_k(\mathbf{x}^*(k)) \geq f(\mathbf{x}^*(k))$ because $F_k(\mathbf{x}^*(k))$ consists of $f(\mathbf{x}^*(k))$ plus a nonnegative penalty term.

Since $f(\mathbf{y}) \geq f(\mathbf{x}^*(k))$, the same inequality holds in the limit. This means that $f(\mathbf{y}) \geq f(\mathbf{x}^*)$, which is exactly what we wanted to know, to show that $\mathbf{x}^*$ is optimal. □

In this proof, the assumption that $\mathbf{x}^*(k)$ exists for all large k and converges to $\mathbf{x}^*$ as $k \to \infty$ can be relaxed: it's enough to have $\mathbf{x}^*(k)$ exist and converge for some unbounded sequence of penalty factors $k_1 < k_2 < k_3 < \cdots$. This will be relevant for us in the next lecture. The proof is the same, but with slightly more complicated notation.

It's also worth mentioning that the same proof holds, not just for the Courant–Beltrami penalty function, but for many others. All we used was that the penalty term is continuous (assuming the continuity of $g_1, \dots, g_m$), vanishes for feasible points, and is positive for infeasible points.

3 Example

Consider the simple optimization problem

$$\begin{aligned} \underset{x,y\in\mathbb{R}}{\text{minimize}} \quad & x^2 + y^2 \\ \text{subject to} \quad & 5 - x - 2y \leq 0. \end{aligned}$$

We begin by writing down

$$F_k(x, y) = x^2 + y^2 + k \cdot \max\{0, 5 - x - 2y\}^2.$$

Note that this is a convex function (and in general, F_k is a convex function when we start with a convex program). Both 0 and $5 - x - 2y$ are convex; the maximum of convex functions is convex; $t \mapsto t^2$ is convex and increasing on nonnegative inputs, so it preserves convexity; therefore F_k is a sum of convex functions.

This is good, because it means that if we solve for the critical points of F_k, we will find the global minimizers. (This is always important to check!)

When $5 - x - 2y \le 0$, we have $F_k(x, y) = x^2 + y^2$; setting $\nabla F_k(x, y)$ to $(0, 0)$ gives us $x = y = 0$, which doesn't satisfy the condition for the case we're looking at. So there are no critical points that appear from this case.

When $5 - x - 2y > 0$, we have $F_k(x, y) = x^2 + y^2 + k(5 - x - 2y)^2$, and

$$\nabla F_k(x, y) = \begin{bmatrix} 2x - 2k(5 - x - 2y) \\ 2y - 4k(5 - x - 2y) \end{bmatrix}.$$

If we set it equal to 0, we get $x = k(5 - x - 2y)$ and $y = 2k(5 - x - 2y)$, so in particular $y = 2x$. Therefore $x = k(5 - x - 4x)$, which we can solve to get $x = \frac{5k}{5k+1}$ and so $y = \frac{10k}{5k+1}$. With this, $x + 2y = \frac{25k}{5k+1} < 5$, so $5 - x - 2y > 0$ and we did find a true critical point.

Because F_k is convex, this critical point is a global minimizer.

Now we take the limit as $k \to \infty$, and find that $(\frac{5k}{5k+1}, \frac{10k}{5k+1}) \to (\frac{5}{5}, \frac{10}{5}) = (1, 2)$. Because the limit exists, Theorem 2.1 tells us that it's an optimal solution to the original problem.

4 Upcoming: guaranteeing convergence

Theorem 2.1 is usually enough to justify applying the penalty method in retrospect. When our hypotheses hold, if we apply the penalty method to P, find some global minimizers $\mathbf{x}^*(k)$, and realize that they converge to something as $k \to \infty$ (or even that a subsequence of them converges), then we know that what they converge to is an optimal solution.

Next time, we will show another result:

Theorem 4.1. *Suppose that P is feasible, $g_1, g_2, \dots, g_m$ are continuous, and f is coercive.*

Then there is some unbounded sequence $k_1 < k_2 < k_3 < \cdots$ with global minimizers $\mathbf{x}^*(k_1), \mathbf{x}^*(k_2), \dots$ that converge to $\mathbf{x}^*$ (which is an optimal solution of P by Theorem 2.1).

This is a different kind of guarantee. Theorem 2.1 guarantees correctness: when it applies, if the penalty method gives an answer, it's the correct answer. Theorem 4.1 promises that for a certain class of problems, the penalty method *will* produce an answer.