

Chapter 1, Lecture 6: Closed and bounded sets

February 1, 2019

University of Illinois at Urbana-Champaign

1 Existence of minimizers in \mathbb{R}

We are now capable of identifying critical points as local minimizers (or maximizers) pretty well. We don't yet know how to tell when a point is a global minimizer. But global minimizers should all be local ones, so we can just look at all of those...

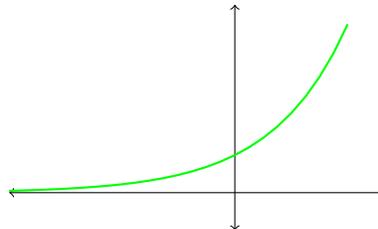
... Unless, of course, a function doesn't *have* a global minimizer.

We are going to consider two ways to guarantee that it does. We can make some assumptions about the function. But we can also (and this is easier) make assumptions about the domain D of a function $f : D \rightarrow \mathbb{R}$ that guarantee the existence of minimizers for any (continuous) f .

We'll look at examples where $D \subseteq \mathbb{R}$, first. There are essentially two reasons why a (continuous) function $f : D \rightarrow \mathbb{R}$ can fail to have a global minimizer or maximizer.

1.1 The domain D is unbounded

The first reason is D might go off to positive or negative infinity. The function $f(x) = e^x$, with $D = \mathbb{R}$, is a good illustration:



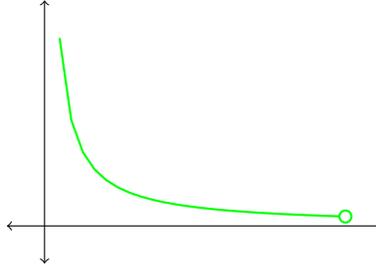
There is no global maximizer because $f(x) \rightarrow \infty$ as $x \rightarrow \infty$: we can make the function arbitrarily large by making x sufficiently large, and there is no limit to how large x can get.

There is no global minimizer for a similar reason. Sure, we can't make $f(x)$ arbitrarily small. However, we can make $f(x)$ arbitrarily close to 0 by taking a sufficiently large negative x . We can never reach this minimum value of 0, so there is no global minimizer: no matter how close to 0 the value of $f(x)$ gets, it can always get closer.

1.2 The domain D is missing a boundary point

The second reason is that D might be missing the point at which a minimizer or maximizer "ought to" occur. A good example of this is the function $f(x) = \frac{1}{x}$ when $D = (0, 1)$:

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



There is no global maximizer because $f(x) \rightarrow \infty$ as $x \rightarrow 0$. We can make $f(x)$ arbitrarily large by making x sufficiently close to 0. The function $f(x)$ wouldn't be defined at 0, but that's not a problem, since 0 isn't in the domain D anyway.

There is no global minimizer because the global minimizer "should" be at the point $(1, 1)$, but 1 is not in the domain D . We can make $f(x)$ arbitrarily close to 1 by making x sufficiently close to 1, but we can never reach 1 (and no matter how close we get, it's always possible to get closer).

2 Closed and bounded sets

It is not obvious, but:

- These are the *only* two ways that a function can fail to have a global minimizer or maximizer.
- Generalizing the two bad properties, the same thing is true for $D \subseteq \mathbb{R}^n$.

A subset $D \subseteq \mathbb{R}^n$ is called *unbounded* if it "goes off to infinity", and *bounded*² otherwise. Formally, we say that $D \subseteq \mathbb{R}^n$ is bounded if there is some (sufficiently large) radius $R > 0$ such that

$$D \subseteq B(\mathbf{0}, R) = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\| < R\}.$$

If we have a function $f : D \rightarrow \mathbb{R}$ and D is unbounded, it's possible for it to fail to have a global minimizer or global maximizer for the same reason that e^x doesn't have them.

To generalize the second bad property, we need to define open sets and closed sets. There is some flexibility in the definitions here. In your textbook, open and closed sets are defined as follows:

- A set $D \subseteq \mathbb{R}^n$ is *open* if all its points are interior points: for any $\mathbf{x} \in D$, there is some sufficiently small $r > 0$ such that the ball $B(\mathbf{x}, r)$ is contained in D .
- A set $D \subseteq \mathbb{R}^n$ is *closed* if it contains all its limit points: whenever we have a sequence $\mathbf{x}^{(0)}, \mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots \in D$ such that $\mathbf{x}^{(n)} \rightarrow \mathbf{x}^*$ as $n \rightarrow \infty$, \mathbf{x}^* is also contained in D .

These are very different-looking definitions. However, we can show that they are complementary: A set $D \subseteq \mathbb{R}^n$ is open if and only if its complement $\mathbb{R}^n \setminus D$ is closed. In fact, it's common to take this as the definition of a closed set: to define $D \subseteq \mathbb{R}^n$ as closed whenever $\mathbb{R}^n \setminus D$ is open.

Don't be misled by terminology: it's possible (and common) for a set to be neither open nor closed. For example, if $D = [0, 1) = \{x \in \mathbb{R} : 0 \leq x < 1\}$, then

²Despite the slightly unfortunate terminology, the words "bounded" and "unbounded" have nothing to do with the boundary of a set; we have separate terminology for boundaries.

- D is not open: if we take $x = 0$, then $x \in D$, but there is no $r > 0$ such that the ball $B(x, r) = (-r, r)$ is a subset of D .
- D is not closed: if we take the sequence $\frac{1}{2}, \frac{2}{3}, \frac{3}{4}, \dots, 1 - \frac{1}{n}, \dots$, then every element of the sequence is in D , but the limit (which is 1) is not.

It's also possible, but unusual, for a set to be both open and closed. The empty set \emptyset is open and closed, because both properties are satisfied trivially: there are no cases to check. All of \mathbb{R}^n is also both open and closed. These are the only two examples.

Closed sets and open sets have a lot to do with strict and non-strict inequalities. For example, the open ball

$$B(\mathbf{x}, r) = \{\mathbf{y} \in \mathbb{R}^n : \|\mathbf{x} - \mathbf{y}\| < r\}$$

is an open set, while the closed ball

$$\overline{B}(\mathbf{x}, r) = \{\mathbf{y} \in \mathbb{R}^n : \|\mathbf{x} - \mathbf{y}\| \leq r\}$$

is a closed set. In general, a set of the form

$$\{\mathbf{x} \in \mathbb{R}^n : f_1(x) < b_1, f_2(x) < b_2, \dots, f_m(x) < b_m\}$$

will be open, provided that f_1, \dots, f_m are continuous functions. A set of the form

$$\{\mathbf{x} \in \mathbb{R}^n : f_1(x) \leq b_1, f_2(x) \leq b_2, \dots, f_m(x) \leq b_m\}$$

will be closed if f_1, \dots, f_m are continuous. Not all open and closed sets have a nice, natural description of this form. But many of the ones we actually do, and it's a good way to get intuition for which sets are open and which are closed.

For now, the reason to introduce open and closed sets is for the following theorem (which we will not prove yet):

Theorem 2.1 (Extreme value theorem). *Let $D \subseteq \mathbb{R}^n$, $D \neq \emptyset$, be a closed and bounded set, and let $f : D \rightarrow \mathbb{R}$ be a continuous function. Then f has a global minimizer (and a global maximizer) on D .*

3 Coercive functions

Using this theorem, we can define a fairly large class of functions which are guaranteed to have local minimizers—even on all of \mathbb{R}^n .

First of all: define a *sublevel set* of a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ (at a level $c \in \mathbb{R}$) to be the set

$$L_c^-(f) := \{\mathbf{x} \in \mathbb{R}^n : f(\mathbf{x}) \leq c\}.$$

That is, $L_c^-(f)$ is the set of points where f is at most c . (If you imagine pouring water over the graph of f such that the surface of the water is at height c , then $L_c^-(f)$ is the set of points that end up underwater.)

A *coercive function* is a continuous function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ such that for all $c \in \mathbb{R}$, the sublevel set $L_c^-(f)$ is bounded.³ That is, for any value of c , there is only a bounded region where f can be that small; outside that bounded region, f is large.

Theorem 3.1. *Every coercive function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ has a global minimizer.*

Proof. Let $\mathbf{x}^{(0)} \in \mathbb{R}^n$ be any arbitrary point. (You can take $\mathbf{x}^{(0)} = \mathbf{0}$ if you like.) Let $c = f(\mathbf{x}^{(0)})$, and let D be the sublevel set $L_c^-(f)$.

By definition of coercive function, the set D is bounded. Since it's defined by a single \leq inequality for a continuous function, D is also closed. Finally, $D \neq \emptyset$: we know $\mathbf{x}^{(0)} \in D$. (That's all we need $\mathbf{x}^{(0)}$ for; you can forget about it now.)

Therefore, by the extreme value theorem, the function f has a global minimizer \mathbf{x}^* on D . We're going to show that \mathbf{x}^* is actually a global minimizer on \mathbb{R}^n .

The point \mathbf{x}^* must satisfy $f(\mathbf{x}^*) \leq c$, because \mathbf{x}^* is contained in the sublevel set $L_c^-(f)$, and that is the condition defining the sublevel set. So for any $\mathbf{x} \notin L_c^-(f)$, we have

$$f(\mathbf{x}) > c \geq f(\mathbf{x}^*).$$

This means that \mathbf{x}^* is not only better than every point in D : it is also better than every point outside D . Therefore it is a global minimizer not just over D , but over all of \mathbb{R}^n . \square

This theorem will be very useful when we develop more of an idea of which functions are coercive. If we know that a function f must have a global minimizer, we can avoid using the Hessian matrix of f to test its critical points. Instead, just evaluate f at all of its critical points. The global minimizer must be one of them: the one that has the least value of f .

We'll see more examples of coercive functions later in this course. For now, here are some very simple cases:

- A function $f : \mathbb{R} \rightarrow \mathbb{R}$ is coercive if and only if

$$\lim_{x \rightarrow \infty} f(x) = \lim_{x \rightarrow -\infty} f(x) = +\infty.$$

The sublevel definition of a coercive function is in this case a restatement of what it means for the limit of a function to be infinite.

- If we can write $f : \mathbb{R}^n \rightarrow \mathbb{R}$ as

$$f(x_1, x_2, \dots, x_n) = f_1(x_1) + f_2(x_2) + \dots + f_n(x_n)$$

and f_1, f_2, \dots, f_n are all coercive, then f is coercive. (This is *not* a necessary condition, just a sufficient one.)

Intuitively, if any $|x_i|$ is very large, then $f_i(x_i)$ will be very large, and the other f_j will not be able to compensate. So the only way to get a small value of f is for $|x_i|$ to be small for all i , which means that f 's sublevel sets are bounded.

³The textbook gives a very differently-phrased definition that's equivalent. I'm hoping that this definition will be less confusing.