Lecture 9

Compactness, I

Let's review what we've done in the last week:

1. We've defined the notion of a topology $\mathcal{T}$ on a set X. The definition was a bit abstract; a topology is a collection of subsets of X, which we will call *open*, satisfying some conditions. When a set is given a topology, we call it a topological space.

2. We also saw that a given set may have many different topologies, so we often have to specify which topology we're talking about. (For example, when $X = \mathbb{R}^n$, we could give X the discrete, the trivial, or the standard topology. Our favorite should be the standard topology.)

3. We also saw that any poset admits a topology, called the Alexandroff topology.

4. We defined a function between topological spaces to be continuous if preimages of open subsets are open.

5. We then gave a way to turn subsets of spaces into spaces, by giving subsets the *subspace* topology. So for example, things like simplices, spheres, and disks are all of a sudden topological spaces!

Today, we're going to learn about a very important property that some topological spaces have, called *compactness*.

In a highly sophisticated way, a space being "compact" will be very much like a set being "finite." Compactness is some way that we can talk about a space as being "small enough" to comprehend using very nice tools.

This doesn't mean that non-compact spaces are inaccessible. (Even for sets, we can understand infinite sets like $\mathbb{Z}$ just fine.) But there are more tools available for studying compact spaces.

9.1 Continuous functions compose

Before we go any further, let me state an incredibly important property of continuous functions.

Proposition 9.1.1. Let X, Y, Z be topological spaces. Fix a continuous function $f : X \to Y$ and another continuous function $g : Y \to Z$. Then the composition

$$g \circ f : X \to Z$$

is also continuous.

Try to prove this as an exercise before looking at the proof:

Proof. Let $W \subset Z$ be an open subset. We must prove that $(g \circ f)^{-1}(W)$ is an open subset of X. Observe:

$$\begin{aligned}(g \circ f)^{-1}(W) &= \{x \in X \mid (g \circ f)(x) \in W\} \\ &= \{x \in X \mid g(f(x)) \in W\} \\ &= \{x \in X \mid f(x) \in g^{-1}(W)\} \\ &= f^{-1}(g^{-1}(W)).\end{aligned}$$

But since g is assumed continuous, we know that $g^{-1}(W)$ is open. Hence $f^{-1}(g^{-1}(W))$ is open (because f is also assumed continuous).

This completes the proof. □

9.2 Open covers and subcovers

Do you remember when we proved that $\mathbb{R}^n$ could be written as a union of open balls? There were many, many different ways that we could do this—we could take a union of open balls all centered at the origin, or we could take a union of open balls with all kinds of different centers, too.

Regardless, either collection is an example of a *cover*.

Definition 9.2.1. Let X be a set. A *cover* of X is a collection of subsets $\{U_\alpha\}_{\alpha\in\mathcal{A}}$ such that

$$\bigcup_{\alpha\in\mathcal{A}} U_\alpha = X.$$

Oftentimes, instead of working with the generality of a function $\mathcal{A} \to \mathcal{P}(X)$, we will assume that a cover is specified by a subset of $\mathcal{P}(X)$ (that is, we may assume that $\mathcal{A} \to \mathcal{P}(X)$ is an injection).

In this case, we may denote an open cover by a fancy U, so

$$\mathcal{U} = \{U_\alpha\}_{\alpha\in\mathcal{A}}.$$

The word "cover" is used because the collection "covers" all of X. (Every element of x is inside one of the U_α.)

Definition 9.2.2. Choose a subset $\mathcal{B} \subset \mathcal{A}$ of the indexing set $\mathcal{A}$. If the collection $\{U_\beta\}_{\beta\in\mathcal{B}}$ is a cover of X, we call it a *subcover* of the original cover.

Example 9.2.3. Let $\mathcal{A} = \mathbb{R}^n \times \mathbb{R}_{>0}$, and for all $(x, r) \in \mathcal{A}$, define $U_{(x,r)} = \text{Ball}(x, r)$. We saw long ago that this collection $\{U_{(x,r)}\}_{(x,r)\in\mathcal{A}}$ is a cover of $\mathbb{R}^n$.

Now choose the subset $\mathcal{B} \subset \mathcal{A}$ consisting of those (x, r) for which x is the origin of $\mathbb{R}^n$, and r is a positive integer. Then the collection $\{U_\beta\}_{\beta\in\mathcal{B}}$ is also a cover of $\mathbb{R}^n$. Hence $\{U_\beta\}_{\beta\in\mathcal{B}}$ is a subcover of $\{U_\alpha\}_{\alpha\in\mathcal{A}}$.

Definition 9.2.4. Let X be a topological space. A cover $\{U_\alpha\}_{\alpha\in\mathcal{A}}$ is called an *open cover* if, for every $\alpha \in \mathcal{A}$, the set $U_\alpha \subset X$ is open.

In other words, an open cover is a cover consisting of open subsets.

Example 9.2.5. Both the covers in Example 9.2.3 are open covers of $\mathbb{R}^n$.

The following exercise shows that covers of X induce covers of subspaces; moreover, open covers of X induce open covers of A, simply by intersecting elements of the cover with A:

Exercise 9.2.6. Let $\{U_\alpha\}_{\alpha\in\mathcal{A}}$ be an open cover of a topological space X. Fix a subset $A \subset X$, and endow it with the subspace topology.

Then the collection $\{V_\alpha\}_{\alpha\in\mathcal{A}}$, where

$$V_\alpha := U_\alpha \cap A$$

is an open cover of A.

Proof. We must verify two facts: That $\{V_\alpha\}_{\alpha\in\mathcal{A}}$ is a cover of A, and that each V_α is open (in A).

Each V_α is open by definition of subspace topology.

Clearly, the union $\bigcup_{\alpha\in\mathcal{A}} V_\alpha$ is a subset of A because each V_α is a subset of A.

So we must only show that A is a subset of this union. So fix $a \in A$. Then $a \in X$ because A is a subset of X. In particular, because $\{U_\alpha\}_{\alpha\in\mathcal{A}}$ is a cover of X, there is some $\alpha \in \mathcal{A}$ for which $a \in U_\alpha$. Because $a \in A$ by assumption, we conclude that $a \in U_\alpha \cap A = V_\alpha$. This finishes the proof. □

9.3 The definition of compactness

Here is one of the most important definitions in topology; it is one of the most confusing as well, but it is incredibly powerful.

Definition 9.3.1. Let X be a topological space. We say that X is *compact* if every open cover of X admits a finite subcover.

In other words, X is called compact if the following holds. For every open cover $\{U_\alpha\}_{\alpha\in\mathcal{A}}$, there exists some *finite* subset $\mathcal{B} \subset \mathcal{A}$ so that the collection $\{U_\beta\}_{\beta\in\mathcal{B}}$ is a cover of X.

9.4 Straightforward examples

Whenever you are given a new definition, it is good to ask for examples and non-examples. Here are a few:

Example 9.4.1. Let X be a finite set, and let $\mathcal{T}$ be a topology on the finite set. (So that we may consider X to be a topological space.)

I claim that X is compact.

Here is a proof: Any topology $\mathcal{T} \subset \mathcal{P}(X)$ has only finitely many elements, so X has only finitely many open subsets to begin with. In particular, fix an open cover $\{U_\alpha\}_{\alpha\in\mathcal{A}}$. It may be that $\mathcal{A}$ itself is an infinite set, we can choose some *finite* subset $\mathcal{B} \subset \mathcal{A}$ so that for every $\alpha \in \mathcal{A}$, there is some $\beta \in \mathcal{B}$ for which $U_\alpha = U_\beta$. Then $\{U_\beta\}_{\beta\in\mathcal{B}}$ is a finite subcover.

Example 9.4.2. Let $X = \mathbb{R}^n$ with the standard topology. This is *not* compact.

To prove that something is not compact, we need only find one example of an open cover that does not admit a finite subcover.

So let's choose the following open cover: Let $\mathcal{A} = \mathbb{Z}_{>0}$ (so that $\mathcal{A}$ is the set of positive integers), and define, for all $n \in \mathcal{A}$,

$$U_n := \text{Ball}(0, n)$$

to be the open ball of radius n about the origin of $\mathbb{R}^n$. We have seen that $\{U_n\}_{n \in \mathbb{Z}_{>0}}$ is an open cover of $\mathbb{R}^n$.

Now, if $\mathcal{B} \subset \mathcal{A}$ is any finite subset (so that $\mathcal{B}$ is some finite collection of positive integers), there is a maximal element of $\mathcal{B}$. Call this maximal element N. Then

$$\bigcup_{\beta \in \mathcal{B}} U_\beta = U_N = \text{Ball}(0, N).$$

In particular, the collection $\{U_\beta\}_{\beta \in \mathcal{B}}$ is not a cover of $\mathbb{R}^n$ because any element $x \in \mathbb{R}^n$ having distance larger than N from the origin is not inside the union $\bigcup_{\beta \in \mathcal{B}} U_\beta$.

Thus, the open cover $\{U_n\}_{n \in \mathbb{Z}_{>0}}$ admits no finite subcover. This shows that $\mathbb{R}^n$ is not compact.

Example 9.4.3. Let X be an infinite set, and let $\mathcal{T}_{disc}$ be the discrete topology on X. (This means that any subset of X is declared open.)

I claim that $(X, \mathcal{T}_{disc})$ is not a compact space.

Let $\mathcal{A} = X$, and for all $x \in \mathcal{A}$, declare U_x to be the one-element set $\{x\} \subset X$. By definition of the discrete topology, U_x is open in X. Hence $\{U_x\}_{x \in \mathcal{A}}$ is an open cover of X.

On the other hand, this does not admit a finite subcover. For a finite subset of $\mathcal{A}$ is a finite collection $\{x_1, \ldots, x_n\}$ of some points in X. The union of the sets $\{x_i\}_{i=1,\ldots,n}$ is clearly the set $\{x_1, \ldots, x_n\}$, which is not all of X because X is assumed infinite.

This shows that any infinite set, when equipped with the discrete topology, is not compact.

Here is another example, which we state as a proposition:

Proposition 9.4.4. Let X be a compact space. Suppose Y is a space homeomorphic to X. Then Y is also compact.

This proposition is a verification that "compactness" is a notion that depends only on the topology of space—after all, homeomorphisms are equivalences of topological spaces.

Proof. Let $\mathcal{V}$ be an open cover of Y. We must prove that $\mathcal{V}$ admits a finite subcover.

Fix a continuous function $\phi : X \to Y$. Then for all $V \in \mathcal{V}$, the pre-image $\phi^{-1}(V)$ is an open subset of X. Let $\mathcal{U}$ be the collection of open subsets given by

$$\mathcal{U} := \{U \subset X \,|\, U = \phi^{-1}(V) \text{ for some } V \in \mathcal{V}.$$

In other words, $U \in \mathcal{U}$ if and only if U arises as the preimage of some $V \in \mathcal{V}$. Then $\mathcal{U}$ is an open cover of X, because $x \in X \implies f(x) \in V$ for some $V \in \mathcal{V}$. (This last claim uses that $\mathcal{V}$ is a cover of Y.)

Because X is compact, we may choose a finite subcover of $\mathcal{U}$. So let $U_1, \ldots, U_n$ be the finite collection of open sets of X for which $\bigcup_{i=1,\ldots,n} U_i = X$ and for which $U_i \in \mathcal{U}$ for all i.

Now suppose that ϕ is further a homeomorphism.[1] Then ϕ is a bijection, so if the collection $\{U_1, \ldots, U_n\}$ is a cover of X, then the collection $\{\phi(U_1), \ldots, \phi(U_n)\}$ is a cover of Y. Moreover, again because ϕ is a bijection, the fact that $U_i = \phi^{-1}(V_i)$ means $\phi(U_i) = V_i$. So the collection $\{V_1, \ldots, V_n\}$ is a finite subcover of $\mathcal{V}$.

We have shown that any open cover of Y admits a finite subcover, completing the proof. □

9.5 Some compact posets

To see the simplest examples of compact topological spaces that have infinitely many elements, we turn to posets.

Let $Z = \mathbb{Z}_{>0} \cup \{\infty\}$. In other words, as a set, Z is obtained by adding one new element to the set of all positive integers. We call this new element ∞. You'll see why in a moment.

Note that Z has countably infinite cardinality.

We can give Z a poset structure as follows. If $a, b \in \mathbb{Z}_{>0} \subset Z$, we declare *aleqb* if a is less than or equal to b in the usual sense (of integers being less than or equal to each other). If $a \in \mathbb{Z}_{>0}$ and $b = \infty$, we declare $a \leq b$. We also declare that $\infty \leq \infty$ to ensure reflexivity.

I will leave it to you to see that this relation is transitive and antisymmetric.

[1] This does NOT mean that every function from X to YH is a homeomorphism; but this is a round-about way of saying that we can choose some homeomorphism from X to Y, and we are calling it ϕ for no good reason.

Then let us endow Z with the Alexandroff topology.

Proposition 9.5.1. Z is compact.

Proof. Let $U \subset Z$ be an open subset. By definition of Alexandroff topology, we know that for every $a \in U$ and every $b \in Z$ satisfying $a \leq b$, we know that $b \in U$. Thus, whenever U is non-empty, we know that $\infty \in U$.

Well, if $\mathcal{U}$ is an open cover of Z, there is some element $U \in \mathcal{U}$ containing $1 \in Z$. Then by definition of Alexandrov topology, U contains every element b satisfying $b > 1$. In other words, $U = Z$.

Thus, any open cover $\mathcal{U}$ must satisfy $Z \in \mathcal{U}$, and in particular, any open cover admits a subcover consisting of only one open subset: Z. □

If you study the above proof, you realize that the only property of Z we used is that there is some element (namely, 1) that is less than or equal to any other element. Such an element doesn't always exist in a poset. (For example, the poset $\mathbb{Z}$ doesn't have such an element.) But here is another example of a compact poset:

Proposition 9.5.2. Let A be a set. (Infinite or otherwise.) Let $\mathcal{P}(A)$ be the power set, considered as a poset via the relation $\subset$. Then $\mathcal{P}(A)$, given the Alexandroff topology, is compact.

Proof. Any open cover $\mathcal{U}$ of $\mathcal{P}(A)$ must have some $U \in \mathcal{U}$ for which $\emptyset \in U$. But by definition of Alexandroff topology, any $B \in \mathcal{P}(A)$ satisfying $\emptyset \subset B$ must be contained in U; meaning U contains every element of $\mathcal{P}(A)$.

In other words, any open cover $\mathcal{U}$ must satisfy $\mathcal{P}(A) \in \mathcal{U}$; so $\mathcal{U}$ admits a finite subcover—in fact, a subcover consisting of a single element called $\mathcal{P}(A)$. □

9.6 Closed intervals

The following is the most important example of a compact space. Most other examples of compact subspaces of Euclidean space are in one way or another constructed out of this one:

Theorem 9.6.1. Let $A = [0, 1] \subset \mathbb{R}$ be the closed interval from 0 to 1. Then A, given the subspace topology, is compact.

You may use this theorem freely from now on. It will soon be superseded by the Heine-Borel theorem, but the proof of the Heine-Borel theorem will actually depend on Theorem 9.6.1.

In fact, there is nothing special about $[0, 1]$. *Any* closed and bounded interval is compact as a consequence of the above theorem.

Corollary 9.6.2. Fix two real numbers a, b satisfying $a < b$. Then the closed interval $[a, b]$ (endowed with the subspace topology inherited from $\mathbb{R}$) is compact.

Proof. By Proposition 9.4.4, we are finished if we can exhibit a homeomorphism between $[a, b]$ and $[0, 1]$.

Consider the function

$$f : \mathbb{R} \to \mathbb{R}, \qquad x \mapsto (b - a)x + a.$$

f is continuous, as we know from calculus. (This is relying on the Theorem from previous lectures that continuity in the sense of topology is equivalent to continuity in the sense of calculus.)

We know that, because $[0, 1]$ is given the subspace topology, the inclusion function $i_{[0,1]} : [0, 1] \to \mathbb{R}$ is also continuous. Moreover, by Proposition 8.3.2, the composition of conitnuous functions is still continuous, so we see that the function

$$j = f \circ i_A[0, 1] : [0, 1] \to \mathbb{R}, \qquad x \mapsto (b - a)x + a$$

is continuous. What is the image of j? It is precisely the interval $[a, b]$. (Note that $j(0) = a$ and $j(1) = b$; I'll let you fill in the rest of the details.) Thus, by the universal property of the subspace topology of $[a, b] \subset \mathbb{R}$, the function

$$j' : [0, 1] \to [a, b], \qquad x \mapsto (b - a)x + a$$

is continuous. It is straightforward to verify that j' is both an injection and a surjection, hence a bijection.

In fact, you can write an inverse to j' as follows:

$$g : [a, b] \to [0, 1] \qquad x \mapsto \frac{1}{b - a}(x - a).$$

A similar argument to the demonstration that j' is continuous shows that g is also continuous.

Because j' is a continuous bijection whose inverse is also continuous, j' is a homeomorphism. □

9.7 Proof that $[0, 1]$ is compact

I will write the proof here, but you may take the above theorem for granted. You will prove it in an Analysis class. The most important part of the proof is an understanding of the construction of $\mathbb{R}$.

First, let us recall the most "con*sequential*"[2] property of $\mathbb{R}$:

Proposition 9.7.1. Let $a_1, a_2, \ldots$ be a bounded, increasing sequence of rational numbers. This means that $a_i \leq a_{i+1}$ for all i, and that there is some real number A so that $a_i < A$ for all i.

Then the sequence $a_1, a_2, \ldots$ converges to some $b \in \mathbb{R}$.

I state this as a proposition, but it is a direct consequence of the construction of the real line as a Cauchy complete entity. I will not get much more into this here. Indeed, the fact one can *construct* a set satisfying the above Proposition is the theorem that allows us to construct the real line.

Proof of Theorem 9.6.1. Reduction I: Let us first assume that $\mathcal{U}$ is an open cover of $[0, 1]$ by open balls. More accurately, suppose that for every $U \in \mathcal{U}$, there is an $x \in \mathbb{R}$ and $r > 0 \in \mathbb{R}$ such that $U = \text{Ball}(x, r) \cap [0, 1]$. We will explain at the end why this reduction allows us to conclude the theorem.

Given such an open cover $\{U_\alpha\}_{\alpha \in \mathcal{A}}$ of $[0, 1]$, we now show that there is a finite subcover. For every $\alpha \in \mathcal{A}$, we choose the numbers (x_α, r_α) so that $U_\alpha = \text{Ball}(x_\alpha, r_\alpha)$.

To begin our proof, let $\mathcal{A}_1 \subset \mathcal{A}$ consist of those α for which $0 \in \text{Ball}(x_\alpha, r_\alpha)$, and let

$$a_1 := \sup_{\alpha \in \mathcal{A}_1} (x_\alpha + r_\alpha).$$

If $a_1 > 1$, we are done, for then there is some (x_α, r_α) for which $\text{Ball}(x_\alpha, r_\alpha)$ contains both 0 and 1, hence the entire interval $[0, 1]$. (This exhibits an open subcover with a single element—$[0, 1]$ itself.)

So suppose $a_1 \leq 1$. We let $\mathcal{A}_2 \subset \mathcal{A}$ consist of those α for which $a_1 \in \text{Ball}(x_\alpha, r_\alpha)$. ANd as before, we define

$$a_2 := \sup_{\alpha \in \mathcal{A}_2} (x_\alpha + r_\alpha).$$

[2]This is a bad joke.

Inductively, if it turns out that $a_n < 1$, we may define $\mathcal{A}_{n+1} \subset \mathcal{A}$ to consist of those α for which $a_n \in \text{Ball}(x_\alpha, r_\alpha)$, and set

$$a_{n+1}$$

to equal the lesser of 1 and $\sup_{\alpha \in \mathcal{A}_{n+1}} x_\alpha + r_\alpha$. (So $a_{n+1} \leq 1$.) Note that

$$a_n \leq a_{n+1}. \tag{9.7.0.1}$$

Why? By definition of open cover, there is some element $U = \text{Ball}(x, r)$ in $\mathcal{U}$ that contains a_n, so by the recentering lemma, there is some $\epsilon > 0$ so that $\text{Ball}(a_n, \epsilon)$ is fully contained in U. In particular, $a_{n+1} \geq x + r > a_n + \epsilon$, so $a_{n+1} > a_n + \epsilon$. In particular, (9.7.0.1) follows.

Then we have an increasing sequence $a_1, a_2, \ldots$ which is bounded (as $a_n \leq 1$ for all i). Moreover, by choosing rational numbers r_n so that $a_n < r_n < a_{n+1}$, we obtain a real number b to which $r_1, r_2, \ldots$ converge by Proposition 9.7.1. I promise it is straightforward to check that the sequence $a_1, a_2, \ldots$ also converges to b.

On the other hand, $[0, 1]$ is closed. (It is also straightforward to see that the complement is open.) Hence $b \in [0, 1]$. But by definition of our open cover, there is some (x, r) for which $\text{Ball}(x, r)$ contains b. Further by definition of convergence, this means that for every n large enough, $a_n \in \text{Ball}(x, r)$, contradicting the definition of a_{n+1} unless $a_{n+1} = 1$. In other words, for n large enough, all the a_n equals 1. So there is only a finite collection of open numbers $0 = a_0, a_1, a_2, \ldots, a_n = 1$ picked out by the above process.

Now choose, for each i, a pair (x_i, r_i) so that $\text{Ball}(x_i, r_i)$ contains a_i and $\text{Ball}(x_i, r_i) \in \mathcal{U}$. By definition of the a_i, we can choose these so that $\text{Ball}(x_i, r_i)$ intersects $\text{Ball}(x_{i+1}, r_{i+1})$ for each i. In particular, these form a finite subcover.

Now let us explain why Reduction One is enough to complete the proof. Given an arbitrary open cover $\mathcal{V} = \{V_\beta\}_{\beta \in \mathcal{B}}$, we may consider a much bigger open cover $\mathcal{U} = \{U_\alpha\}_{\alpha \in \mathcal{A}}$ consisting of only open balls, satisfying the property that for every open ball $U_\alpha \in \mathcal{U}$, there exists some β such that $U_\alpha \subset V_\beta$. (For example, $\mathcal{A}$ could consist of triplets (x, r, β) for which $\text{Ball}(x, r) \subset V_\beta$.) Then Reduction One allows us to choose a finite collection $\alpha_1, \ldots, \alpha_n \in \mathcal{A}$ for which $\{U_{\alpha_1}, \ldots, U_{\alpha_n}\}$ forms a subcover. Then, for each i, choosing a β_i such that $V_{\beta_i} \supset U_{\alpha_i}$, we see that $\{V_{\beta_1}, \ldots, V_{\beta_n}\}$ forms a subcover of $\mathcal{V}$. □