

17. Tychonoff's theorem, and more on compactness

1 Motivation

In this section we will prove Tychonoff's theorem. Before we can do that, we will discuss a different characterization of compactness in terms of filters and ultrafilters. We will have to do some work to create the tools we need, but they make the proof itself very straightforward.

After that we will discuss some other notions related to compactness.

Tychonoff's theorem, which says that compactness is (arbitrarily) productive, is an absolutely extraordinary result. The intuition that we have developed thus far for compact topological spaces is that they feel small. A space can be very large in terms of its number of points, but if it is compact it acts like a finite space in most ways; absent any other intuition, finite spaces are the only obvious ones for which all covers have finite subcovers. We also know that taking large products can break many of the "smallness" properties we have seen. Separable spaces feel small in some sense, but we now know that large products of separable spaces may not be separable. Second countable spaces feel small in a different sense, and again we know that large (uncountable) products of second countable spaces are almost never second countable.

In fact, even *finiteness* itself is not countably productive. $\{0, 1\}$ is a finite topological space, and $\{0, 1\}^{\mathbb{N}}$ is *uncountable*, let alone finite. Compactness, however, is arbitrarily productive. Any product of compact spaces, no matter how large, is again compact. In this way, compact spaces can feel even smaller than finite spaces; they feel more like single points. The more one studies compact spaces in higher mathematics, the more one gets the feeling that compact sets are more like single points than anything else, as this intuition underlies many definitions one may see later. We will see just a bit of this intuition on the Big List.

2 Filters and compactness

Before we prove Tychonoff's theorem, we will establish a framework for characterizing compactness in terms of filters, ultrafilters, and sets with the finite intersection property. We will redefine any of these concepts that we have seen earlier.

Definition 2.1. *Let X be a set. A collection \mathcal{A} of subsets of X has the finite intersection property or FIP if for every finite collection $\{A_1, \dots, A_n\} \subseteq \mathcal{A}$, their common intersection $\bigcap_{k=1}^n A_k$ is not empty.*

Note that a collection having the FIP does not imply that the intersection of the whole collection is not empty. For example, the collection of all subsets of the natural numbers of the form $\{n, n+1, n+2, \dots\}$ has the FIP, but the intersection of all subsets of this form is empty.

This is a boring-looking definition, but it is intimately connected to compactness.

Proposition 2.2. *A topological space (X, \mathcal{T}) is compact if and only if for every collection \mathcal{A} of closed subsets of X with the finite intersection property, $\bigcap \mathcal{A} \neq \emptyset$.*

Proof. (\Rightarrow). Suppose (X, \mathcal{T}) is compact, and let \mathcal{A} be a collection of closed subsets of X with the finite intersection property. Consider the collection of open sets $\mathcal{U} = \{X \setminus C : C \in \mathcal{A}\}$.

First notice that no finite subcollection of \mathcal{U} covers X . Indeed, if $X \setminus C_1, \dots, X \setminus C_n$ is a finite subcollection of \mathcal{U} , then by DeMorgan's laws

$$\bigcup_{k=1}^n (X \setminus C_k) = X \setminus \left(\bigcap_{k=1}^n C_k \right) \neq X,$$

since the intersection in the middle term is nonempty by assumption. This implies that \mathcal{U} is not a cover of X , since if it was a cover it would have a finite subcover by compactness. In other words:

$$X \neq \bigcup_{U \in \mathcal{U}} U = \bigcup_{C \in \mathcal{A}} (X \setminus C) = X \setminus \left(\bigcap_{C \in \mathcal{A}} C \right).$$

This means the big intersection in the last term is nonempty, which is what we wanted to show.

(\Leftarrow). We prove this direction by contrapositive. So suppose \mathcal{U} is a cover of X that has no finite subcover. Then $\bigcup \mathcal{U} = X$, and so we have:

$$\bigcap_{U \in \mathcal{U}} (X \setminus U) = \emptyset,$$

but by assumption any finite subset $\mathcal{F} \subseteq \mathcal{U}$ is not a cover, meaning $\bigcup_{U \in \mathcal{F}} U \neq X$. Then we have:

$$\bigcap_{U \in \mathcal{F}} (X \setminus U) \neq \emptyset.$$

This means that $\mathcal{A} = \{X \setminus U : U \in \mathcal{U}\}$ is a collection of closed sets with the finite intersection property whose intersection is empty, as required. \square

Before we go on, we recall the definition of a filter and ultrafilter.

Definition 2.3. *Let X be a set. A nonempty collection $\mathcal{F} \subseteq \mathcal{P}(X)$ is called a filter on X if the following three properties are satisfied.*

1. $\emptyset \notin \mathcal{F}$.
2. \mathcal{F} is closed upwards: if $A \in \mathcal{F}$ and $A \subseteq B$, then $B \in \mathcal{F}$.
3. \mathcal{F} is closed under finite intersections: if $A, B \in \mathcal{F}$, then $A \cap B \in \mathcal{F}$.

A filter \mathcal{F} on a set X is called an ultrafilter if it is not properly contained in any other filter on X .

Also recall the following important characterization of an ultrafilter:

Proposition 2.4. *Let $\mathcal{U} \subseteq \mathcal{P}(X)$ be a filter. Then \mathcal{U} is an ultrafilter if and only if for every $A \subseteq X$, either $A \in \mathcal{U}$ or $X \setminus A \in \mathcal{U}$.*

And finally here is another fact we will use:

Proposition 2.5. *Let X be a set and let \mathcal{U} be an ultrafilter on X . Suppose X is written as a union of finitely many sets $X = X_1 \cup \cdots \cup X_n$. Then there is a k such that $X_k \in \mathcal{U}$.*

Proof. Suppose for the sake of contradiction that $X_k \notin \mathcal{U}$ for all $k = 1, \dots, n$. Then by the previous proposition, $X \setminus X_k \in \mathcal{U}$ for all $k = 1, \dots, n$. These are now finitely many sets in \mathcal{U} , so their intersection must be in \mathcal{U} . But their intersection is empty:

$$\bigcap_{k=1}^n (X \setminus X_k) = X \setminus \left(\bigcup_{k=1}^n X_k \right) = X \setminus X = \emptyset,$$

contradicting the fact that \mathcal{U} is a filter. □

With these definitions and basic tools established, we can give a characterization of compactness in terms of ultrafilters.

Definition 2.6. *Let (X, \mathcal{T}) be a topological space, let $\mathcal{F} \subseteq \mathcal{P}(X)$ be a filter, and suppose $x \in X$. Then \mathcal{F} is said to converge to x , denoted $\mathcal{F} \rightarrow x$, if every open subset U containing x is an element of \mathcal{F} .*

A great deal of elaboration on this definition can be found in my supplementary notes on nets and filters. This is a very powerful definition which is much better than the definition of sequence convergence. For now though, the following proposition is our payoff.

Proposition 2.7 (Ultrafilter characterization of compactness). *A topological space (X, \mathcal{T}) is compact if and only if every ultrafilter on X converges.*

Proof. (\Rightarrow). Suppose (X, \mathcal{T}) is compact, and suppose for the sake of contradiction that \mathcal{U} is an ultrafilter on X that does not converge to any point. This means that for every $x \in X$, there is an open set U_x containing x such that $U_x \notin \mathcal{U}$. Then $\{U_x : x \in X\}$ is an open cover of X , and so since X is compact it has a finite subcover $\{U_{x_1}, U_{x_2}, \dots, U_{x_n}\}$. In other words

$$X = U_{x_1} \cup \cdots \cup U_{x_n}.$$

But then by Proposition 2.5, $U_{x_k} \in \mathcal{U}$ for some $k = 1, \dots, n$, a contradiction.

(\Leftarrow). Suppose every ultrafilter on X converges, and let \mathcal{A} be a collection of closed subsets of X with the FIP. We want to show that $\bigcap \mathcal{A} \neq \emptyset$.

\mathcal{A} generates a filter on X , by first adding all finite intersections of elements of \mathcal{A} (which, note, does not add \emptyset since \mathcal{A} has the FIP), and then adding all supersets of the resulting collection. Call this filter \mathcal{F} . Then by Zorn's Lemma \mathcal{F} is contained in an ultrafilter \mathcal{U} on X .

By assumption, $\mathcal{U} \rightarrow x$ for some $x \in X$. We show that $x \in \bigcap \mathcal{A}$. Fix any $A \in \mathcal{A}$, and let U be any open set containing x . Then $U \in \mathcal{U}$ since $\mathcal{U} \rightarrow x$. Now we have that A and U are both elements of \mathcal{U} , and therefore $A \cap U \neq \emptyset$. Since we can do this for any such U , we have shown that $x \in \overline{A} = A$. Repeating this argument for each $A \in \mathcal{A}$ shows that $x \in A$ for all $A \in \mathcal{A}$, or in other words $x \in \bigcap \mathcal{A}$. \square

3 Tychonoff's theorem

We need just one more tool.

Proposition 3.1. *Let X and Y be sets, let \mathcal{F} be a filter on X , and let $f : X \rightarrow Y$ be a function. Then the collection*

$$f_*(\mathcal{F}) := \{ B \subseteq Y : f^{-1}(B) \in \mathcal{F} \}$$

is a filter on Y . If \mathcal{F} is an ultrafilter on X , then $f_(\mathcal{F})$ is an ultrafilter on Y .*

Proof. This proof is entirely routine symbol-pushing, but we provide it here in case the reader is skeptical.

We first check that $f_*(\mathcal{F})$ is a filter. $Y \in f_*(\mathcal{F})$ since $f^{-1}(Y) = X \in \mathcal{F}$, and so $f_*(\mathcal{F})$ is not empty.

1. $f^{-1}(\emptyset) = \emptyset \notin \mathcal{F}$, so $\emptyset \notin f_*(\mathcal{F})$.
2. $f_*(\mathcal{F})$ is closed upwards: Let $A \in f_*(\mathcal{F})$, and let $B \supseteq A$ be a superset of A . Then $f^{-1}(B) \supseteq f^{-1}(A)$. $f^{-1}(A) \in \mathcal{F}$ since $A \in f_*(\mathcal{F})$, and so $f^{-1}(B) \in \mathcal{F}$ as well since \mathcal{F} is closed upwards. But then $B \in f_*(\mathcal{F})$ by definition.
3. $f_*(\mathcal{F})$ is closed under finite intersections. Let $A, B \in f_*(\mathcal{F})$. Then by definition $f^{-1}(A)$ and $f^{-1}(B)$ are elements of \mathcal{F} . But $f^{-1}(A \cap B) = f^{-1}(A) \cap f^{-1}(B)$, the latter of which is in \mathcal{F} since \mathcal{F} is closed under finite intersections.

Now assume \mathcal{F} is an ultrafilter, and we show that $f_*(\mathcal{F})$ is an ultrafilter using the characterization in Proposition 2.4. Let $A \subseteq Y$ be a set, and assume $A \notin f_*(\mathcal{F})$. Then we want to show that $Y \setminus A \in f_*(\mathcal{F})$. But

$$f^{-1}(Y \setminus A) = X \setminus f^{-1}(A),$$

and the latter set is in \mathcal{F} since $f^{-1}(A)$ is not in \mathcal{F} by assumption, and \mathcal{F} is an ultrafilter. \square

As you would hope, continuous functions respect filter convergence in this sense.

Proposition 3.2. *Let (X, \mathcal{T}) and (Y, \mathcal{U}) be a topological space, and let $f : X \rightarrow Y$ be a continuous function. Suppose \mathcal{F} is a filter on X converging to a point x . Then $f_*(\mathcal{F}) \rightarrow f(x)$.*

Proof. Suppose $\mathcal{F} \rightarrow x$, and let U be any open subset of Y containing $f(x)$. We want to show that $U \in f_*(\mathcal{F})$, or in other words that $f^{-1}(U) \in \mathcal{F}$. But this is immediate, since $f^{-1}(U)$ is an open set containing x , and therefore is in \mathcal{F} since $\mathcal{F} \rightarrow x$. \square

This property actually *characterizes* continuity, as is proved in the supplementary notes on nets and filters. That is, a function $f : X \rightarrow Y$ is continuous *if and only if* it respects the convergence of all filters.

Anyway, this allows us to state and prove the powerful tool we will use to prove Tychonoff's theorem. This result is the analogue, in “filter world”, of the results that say a sequence converges in a product if and only if all of the component sequences converge, or that a function to a product is continuous if and only if all the component functions are continuous.

Lemma 3.3. *Let I be a nonempty indexing set, and let $\mathcal{X} = \{(X_\alpha, \mathcal{T}_\alpha) : \alpha \in I\}$ be a collection of nonempty topological spaces. Let $X = \prod_{\alpha \in I} X_\alpha$ be their product, equipped with the product topology. Let \mathcal{F} be a filter on X , and let $x \in X$. Then $\mathcal{F} \rightarrow x$ if and only if $(\pi_\alpha)_*(\mathcal{F}) \rightarrow \pi_\alpha(x)$ for all $\alpha \in I$.*

Proof. (\Rightarrow). Suppose $\mathcal{F} \rightarrow x$, and fix some $\alpha \in I$. We want to show that $(\pi_\alpha)_*(\mathcal{F}) \rightarrow \pi_\alpha(x)$. But this follows immediately from the previous proposition and the fact that $\pi_\alpha : X \rightarrow X_\alpha$ is continuous.

(\Leftarrow). Now suppose that $(\pi_\alpha)_*(\mathcal{F}) \rightarrow \pi_\alpha(x)$ for all $\alpha \in I$. We want to show that $\mathcal{F} \rightarrow x$, or in other words that for every open set U containing x in the product topology on X , that $U \in \mathcal{F}$. Note that it suffices to prove this for basic open sets since \mathcal{F} is closed upwards.

So, let U be a basic open set containing x . By definition of the product topology, U is of the form

$$U = \pi_{\alpha_1}^{-1}(V_1) \cap \cdots \cap \pi_{\alpha_n}^{-1}(V_n),$$

where for each $k = 1, \dots, n$, V_k is an open subset of X_{α_k} . Now for each $k = 1, \dots, n$, since $x \in U$ we also have $\pi_{\alpha_k}(x) \in V_k$, and since by assumption $(\pi_{\alpha_k})_*(\mathcal{F}) \rightarrow \pi_{\alpha_k}(x)$, this means $V_k \in (\pi_{\alpha_k})_*(\mathcal{F})$. By definition of $(\pi_{\alpha_k})_*$, this means $\pi_{\alpha_k}^{-1}(V_k) \in \mathcal{F}$. All of this is true for every $k = 1, \dots, n$, and so U is also in \mathcal{F} , being a finite intersection of elements of \mathcal{F} . \square

Phew! We are finally ready. Actually, the previous framework of definitions and results did all the heavy lifting.

Theorem 3.4 (Tychonoff's theorem). *Compactness is productive.*

Proof. Let I be a nonempty indexing set, and let $\mathcal{X} = \{(X_\alpha, \mathcal{T}_\alpha) : \alpha \in I\}$ be a collection of nonempty, compact topological spaces. We want to prove that $X = \prod_{\alpha \in I} X_\alpha$ is compact with

its product topology. We show this by using the result of Proposition 2.7, which is to say by showing that every ultrafilter on X converges. So fix an ultrafilter \mathcal{U} on X .

Fix $\alpha \in I$. By Proposition 3.1, $(\pi_\alpha)_*(\mathcal{U})$ is an ultrafilter on X_α . Since X_α is compact, $(\pi_\alpha)_*(\mathcal{U}) \rightarrow x_\alpha$ for some $x_\alpha \in X_\alpha$ by Proposition 2.7. Let $x \in X$ be the point such that $\pi_\alpha(x) = x_\alpha$ for all $\alpha \in I$. Then $(\pi_\alpha)_*(\mathcal{U}) \rightarrow \pi_\alpha(x)$ for all $\alpha \in I$, and therefore by Proposition 3.3, $\mathcal{U} \rightarrow x$. \square

4 Some definitions related to compactness

Just for some extra flavour, we are going to spend a bit of time talking about some topological properties that are very similar to compactness. This whole section should be considered supplementary, but some students will find it interesting. This is just the tip of the iceberg.

First, a definition usually first given in second year calculus (though likely without this name).

Definition 4.1. *A topological space (X, \mathcal{T}) is called sequentially compact if every sequence in X has a convergent subsequence.*

As the name implies, this is an attempt at characterizing of compactness using only sequences. We should immediately be skeptical that sequences are capable of characterizing compactness, however. As before, we expect first countability to patch the hole (if there is one); it seems like a first countable topological space should be compact if and only if it is sequentially compact. Surprisingly, this is not true.

Example 4.2. ω_1 with its order topology is first countable and sequentially compact, but not compact. We know all of these things already, essentially. We have not specifically proved that every sequence in ω_1 has a convergent subsequence, but it is an easy exercise to prove this to yourself given what we know.

This is not good. The problem is essentially that open covers can be very large, and a local countability property cannot help sequences capture all the information they need to capture in order to characterize compactness. If the open covers are restricted to be smaller, things get better.

Definition 4.3. *A topological space (X, \mathcal{T}) is called countably compact if every **countable** open cover of X has a finite subcover.*

Obviously compactness implies countable compactness, but not the other way around. Countable compactness can be characterized in terms of accumulation points as well. We will avoid proving this equivalence for now, as it is somewhat tedious.

Proposition 4.4. *A topological space (X, \mathcal{T}) is countably compact if and only if every infinite subset of X has an ω -accumulation point. That is, for every infinite $S \subseteq X$, there is a point $x \in X$ such that every open U containing x contains infinitely many points of S .*

Being an ω -accumulation point of a set S is like a stronger version of being an element of \bar{S} .

Another approach to the first countability issue is to define something a little more general than sequential compactness.

Definition 4.5. A topological space (X, \mathcal{T}) is called limit point compact if every infinite subset of X has a limit point. That is, for every infinite $S \subseteq X$, there is a point $x \in X$ such that for every open U containing x , $(U \cap S) \setminus \{x\} \neq \emptyset$.

These definitions interact in some unusual ways, and we will not get into most of it here. We will just note a few things, starting with how first countability fits into this picture.

Proposition 4.6. Every sequentially compact space is countably compact.

Proof. Suppose (X, \mathcal{T}) is sequentially compact, and let $S \subseteq X$ be an infinite set. We show that S has an ω -accumulation point, which implies that (X, \mathcal{T}) is countably compact by Proposition 4.4.

Let $\{s_n : n \in \mathbb{N}\}$ be any enumeration of any countably infinite subset of S . Then we may regard S as a sequence, which by definition of sequential compactness must have a convergent subsequence. The point to which this subsequence converges is an ω -accumulation point of S . □

Proposition 4.7. A first countable topological space is countably compact if and only if it is sequentially compact.

Proof. Every sequentially compact space is countably compact by the previous proposition, so for the other direction let (X, \mathcal{T}) be a first countable, countably compact topological space, and let $\{x_n\}_{n \in \mathbb{N}}$ be a sequence in X . We want to show that X has a convergent subsequence.

If $\{x_n : n \in \mathbb{N}\}$ is finite, then by the pigeonhole principle there must be an infinite $A \subseteq \mathbb{N}$ and an $x \in X$ such that $x_n = x$ for all $n \in A$. That is, there must be a constant subsequence. This subsequence obviously converges, and we are done.

So assume that the sequence has infinitely many different values. Then $\{x_n : n \in \mathbb{N}\}$ is infinite, and so by Proposition 4.4 it has an ω -accumulation point, which we call x . Let $\mathcal{B}_x = \{B_n : n \in \mathbb{N}\}$ be a nested, countable local basis at x . Then for each k we can pick an element x_{n_k} from our sequence in the set B_k . Then $\{x_{n_k}\}_{k \in \mathbb{N}}$ is easily seen to be a subsequence converging to x . □

Example 4.8. There is no implication in general between compactness and sequential compactness. We saw earlier that ω_1 is sequentially compact but not compact. On the other hand, $X = [0, 1]^{[0,1]}$ is compact (by Tychonoff's theorem!) but not sequentially compact. We will not give the full details of the latter fact here, but the following sequence in X is an example of a sequence with no convergent subsequence. For each $n \in \mathbb{N}$, let $f_n \in X$ be:

$$f_n : [0, 1] \rightarrow [0, 1] \quad \text{defined by} \quad f_n(x) = \text{the } n^{\text{th}} \text{ digit in the binary decimal expansion of } x.$$

(For points with more than one binary expansion, pick the finite one.) What a weird sequence! The details of this can be found in *Counterexamples in Topology*.

Second countability is a property that says any collection of open sets is essentially a countable collection of open sets. In light of this, the following proposition should not be surprising.

Proposition 4.9. *A second countable topological space is compact if and only if it is countably compact.*

Proof. Exercise. It is easy to see that compact implies countably compact. For the other direction, start with an open cover \mathcal{U} , and use second countability to define an associated countable cover to which you can apply countable compactness. \square

Proposition 4.10. *If (X, \mathcal{T}) is first countable, T_1 , and limit point compact, then it is sequentially compact.*

Proof. Exercise. This is a quite straightforward use of first countability to define a subsequence by induction. \square

This is a lot to take in. There are *many* more propositions along these lines we could prove. Instead, we next examine this in the context of metric spaces, which are very well-behaved.

5 Compactness in metric spaces

In a metric space, all of these bizarre definitions that are close to the definition of compactness calm down and coincide. We start off with a natural definition, then state the theorem that justifies this.

Definition 5.1. *A metric space (X, d) is called totally bounded if for every $\epsilon > 0$ there exist finitely many points x_1, \dots, x_n such that $X = \bigcup_{k=1}^n B_\epsilon(x_k)$.*

That is, the space can be covered by finitely many ϵ -balls, for any $\epsilon > 0$. This is stronger than being simply being bounded. This property somehow separates out metrics that are “artificially” designed specifically to be bounded. For example, \mathbb{R} with the metric $\bar{d}(x, y) = \min\{1, |x - y|\}$ is bounded but not totally bounded (prove this to yourself).

Theorem 5.2. *Let (X, d) be a metric space, thought of as a topological space with the metric topology generated by d . Then the following are equivalent:*

1. X is compact.
2. X is sequentially compact.
3. X is countably compact.

4. X is limit point compact.

5. X is complete and totally bounded (this should feel similar the Heine-Borel theorem).

Proof. Exercise. A section of the Big List is dedicated to guiding the reader through this, in parts.

(None of this material will be tested; it is just here for flavour.)

□