

(1)

Compact Sets

Suppose that you are sitting by a restaurant table facing the street. Suddenly you feel a tickle in your nose, which causes you to sneeze on the table. Vainly you try to cover the evidence of your social ineptitude by the napkins from the box on your table, when an angry tree-hugging environmental activist runs in your direction. The following conversation takes place:



(2)

Activist: Hey man, you are, like, wasting trees or something?

You: I beg your pardon?

Activist: You are using too much paper to cover your snaright. Just imagine how many trees would still be growing if you used napkins $\frac{1}{4}$ the size of these ones? (Hands you scissors). Here! Reduce your carbon print.

Reluctantly, you do your little part to save the rain forest and cut each napkin in the box into 4 identical napkins of smaller size. You then proceed to cover the stains on the table with the smaller napkins, but the activist is unrelenting.

Activist: Look how much paper is wasted! This is not nearly good enough. Cut each napkin further into 4 smaller ones.

At this point you should realize that the size of your paper napkin is never small enough for the tree-hugging hippie. Whenever you succeed in covering the stain with a napkin of diagonal length $\epsilon > 0$, the hippie demands that you cover

(3)

the stain with finitely many napkins of diagonal length $\epsilon/2$. It would appear you are doomed to cut napkins for the rest of your life, right? Perhaps not! What if there exists a diagonal length small enough that you would have to use infinitely many napkins to cover the stain? Surely no hippie would want your napkins to be brought down to that size. After all, infinitely many napkins would mean infinitely many trees cut down. To divine your fate, we must decide whether the stain on your table is totally bounded or not.

Def: A set $A \subset (M, d)$ is said to be totally bounded if, given any $\epsilon > 0$, there exist finitely many points $x_1, \dots, x_n \in M$ such that $A \subset \bigcup_{i=1}^n B_\epsilon(x_i)$. That is, each $x \in A$ is within ϵ of some x_i .

The set $\{x_1, \dots, x_n\}$ is called an ϵ -net for A .

In the definition of a totally bounded set A , we could easily insist that each ϵ -ball be centered at a point of A .

Indeed, given $\epsilon > 0$, choose $x_1, \dots, x_n \in M$ so that $A \subset \bigcup_{i=1}^n B_{\epsilon/2}(x_i)$.

We may certainly assume that $A \cap B_{\epsilon/2}(x_i) \neq \emptyset$ for each i ,

and so we may choose a point $y_i \in B_{\epsilon/2}(x_i) \cap A$ for each i . By

the triangle inequality, we then have $A \subset \bigcup_{i=1}^n B_\epsilon(y_i)$ (Why?).

(4)

That is, A can be covered by finitely many ϵ -balls, each centered at a point in A . More to the point, a set A is totally bounded if and only if A can be covered by finitely many arbitrary sets of diameter at most ϵ , for any $\epsilon > 0$.

Lemma: A is totally bounded iff, given $\epsilon > 0$, there are finitely many sets $A_1, \dots, A_n \subset A$, with $\text{diam}(A_i) < \epsilon$ for all i , such that $A \subset \bigcup_{i=1}^n A_i$

Proof: First suppose that A is totally bounded. Given $\epsilon > 0$, we may choose $x_1, \dots, x_n \in M$ such that $A \subset \bigcup_{i=1}^n B_\epsilon(x_i)$. As above, A is then covered by the sets $A_i = A \cap B_\epsilon(x_i) \subset A$ and $\text{diam}(A_i) \leq \text{diam}(B_\epsilon(x_i)) \leq 2\epsilon$ for each i .

Conversely, given $\epsilon > 0$, suppose that there are finitely many sets $A_1, \dots, A_n \subset A$, with $\text{diam}(A_i) < \epsilon$ for all i , such that $A \subset \bigcup_{i=1}^n A_i$. Given $x_i \in A_i$, we then have $A_i \subset B_{2\epsilon}(x_i)$

for each i and, hence, $A \subset \bigcup_{i=1}^n A_i \subset \bigcup_{i=1}^n B_{2\epsilon}(x_i)$. Since ϵ is arbitrary in either case, we are done.

Notice that the condition in the lemma above demands that A_1, \dots, A_n be subsets of A . This is no real constraint since,

(5)

after all, if A is covered by $B_1, \dots, B_n \subset M$, then A is also covered by the sets $A_i = A \cap B_i \subset A$ and $\text{diam}(A_i) \leq \text{diam}(B_i)$.

Ex. (a) By the triangle inequality, a totally bounded set is necessarily bounded. (Why?) Notice also that any subset of a totally bounded set is again totally bounded.

(b) A finite set is always totally bounded. In a discrete space, a set is totally bounded iff it is finite. (Why?)

(c) In \mathbb{R} , we do not get anything new: A subset of \mathbb{R} is totally bounded iff it is bounded. Thus, total boundedness is apparently not a topological property; it depends intimately on the metric at hand.

(d) In general, not every bounded set is totally bounded.

The discrete metric gives us a clue as to how we might construct such a set. Recall the sequence $e^{(n)} = (0, \dots, 0, 1, 0, \dots)$ in l_1 , where the single nonzero entry is in the n^{th} place. Then, $\{e^{(n)}; n \geq 1\}$ is a bounded set in l_1 , since $\|e^{(n)}\|_1 = 1$ for all n , but not totally bounded. Why? Because $\|e^{(m)} - e^{(n)}\|_1 = 2$ for $m \neq n$; thus, $\{e^{(n)}; n \geq 1\}$ cannot be covered by finitely many balls of radius < 1 . In fact, the set $\{e^{(n)}; n \geq 1\}$ is discrete in its relative metric.

(6)

We next give a sequential criterion for total boundedness. The key observation is isolated in:

Lemma: Let $\{x_n\}$ be a sequence in (M, d) , and let $A = \{x_n : n \geq 1\}$ be its range.

(i) If $\{x_n\}$ is Cauchy, then A is totally bounded.

(ii) If A is totally bounded, then $\{x_n\}$ has a Cauchy subsequence.

Proof: (i) Let $\epsilon > 0$. Then, since $\{x_n\}$ is Cauchy, there is some index $N \geq 1$ such that $\text{diam}\{x_n : n \geq N\} < \epsilon$. Thus:

$$A = \underbrace{\{x_1\} \cup \dots \cup \{x_{N-1}\} \cup \{x_n : n \geq N\}}_{N \text{ sets of diameter } < \epsilon}$$

(ii) If A is a finite set, we are done. (Why?) So, suppose that A is an infinite totally bounded set. Then A can be covered by finitely many sets of diameter < 1 . One of these sets, at least, must contain infinitely many points of A . Call this set A_1 . But then A_1 is also totally bounded, and so it can be covered by finitely many sets of diameter $< \frac{1}{2}$. One of these, call it A_2 , contains infinitely many points of A_1 . Continuing this process, we find a decreasing sequence of sets $A \supset A_1 \supset A_2 \supset \dots$,

(7)

where each A_k contains infinitely many x_n and where $\text{diam}(A_k) < \frac{1}{k}$. In particular, we may choose a subsequence $\{x_{n_k}\}$ with $x_{n_k} \in A_k$ for all k (How?). That $\{x_{n_k}\}$ is Cauchy is now clear since $\text{diam}\{x_{n_j} : j \geq k\} \leq \text{diam}(A_k) < \frac{1}{k}$.

Ex. (a) The sequence $x_n = (-1)^n$ in \mathbb{R} shows that a Cauchy subsequence is the best that we can hope for in the above lemma. (ii)

(b) Note that the sequence $\{e^{in}\}$ in \mathbb{C} has no Cauchy subsequence.

We are finally ready for our sequential characterization of total boundedness:

Thm: A set A is totally bounded iff every sequence in A has a Cauchy subsequence.

Proof: The forward implication is clear from the above lemma. To prove the backward implication, suppose that A cannot be covered by finitely many ϵ -balls. Thus, by induction, we can find a sequence $\{x_n\}$ in A such that $d(x_n, x_m) \geq \epsilon$ whenever $m \neq n$. (How?) But then, $\{x_n\}$ has no Cauchy subsequence.

(8)

All of this should remind you of the Bolzano-Weierstrass Theorem - and for good reasons:

Corollary: (The Bolzano-Weierstrass Theorem) Every bounded infinite subset of \mathbb{R} has a limit point in \mathbb{R} .

Proof: Let A be a bounded infinite subset of \mathbb{R} . Then, in particular, there is a sequence $\{x_n\}$ of distinct points in A . Since A is totally bounded, there is a Cauchy subsequence $\{x_{n_k}\}$ of $\{x_n\}$. But Cauchy sequences in \mathbb{R} converge, and so $\{x_{n_k}\}$ converges to some $x \in \mathbb{R}$. Thus, x is a limit point of A .

Before we are ready to talk about compact sets, we need a few results about completeness.

Def: A metric space M is said to be complete if every Cauchy sequence in M converges to a point in M .

Ex. (a) \mathbb{R} is complete. This is a consequence of the least upper bound axiom; in fact, as we will see, the completeness of \mathbb{R} is equivalent to the least upper bound axiom.

(b) \mathbb{R}^n is complete (because \mathbb{R} is)

(c) Any discrete space is complete (trivially).

(d) $(0,1)$ is not complete. (Why?) Hence completeness is not preserved by continuous functions.

(9)

(e) C_0 , l_1 , l_2 , l_p , and l_∞ are all complete. The proofs are all very similar; we sketch the proof for l_2 below.

(f) $C[a, b]$ is complete. The proof is not terribly difficult, but it will best serve our purposes to postpone it.

The proof that l_2 is complete is based on a few simple principles that will generalize to all sorts of different settings. This generality will become all the more apparent if we introduce a slight change in our notation. Since a sequence is just another name for a function on \mathbb{N} , let's agree to write an element $f \in l_2$ as $f = \{f(k)\}_{k=1}^{\infty}$, in which case $\|f\|_2 = \left(\sum_{k=1}^{\infty} |f(k)|^2 \right)^{1/2}$.

For example, the notorious vectors $e^{(n)}$ will now be written as e_n , where $e_n(k) = \delta_{n,k}$ (This is Kronecker's delta, defined by $\delta_{n,k} = 1$ if $n=k$ and $\delta_{n,k} = 0$ otherwise.)

Let $\{f_n\}$ be a sequence in l_2 , where now we write $f_n = \{f_n(k)\}_{k=1}^{\infty}$, and suppose that $\{f_n\}$ is Cauchy in l_2 . That is, suppose that for each $\epsilon > 0$ there exists an n_0 such that $\|f_m - f_n\|_2 < \epsilon$ whenever $m, n \geq n_0$. Of course, we want to show that $\{f_n\}$ converges, in the metric of l_2 , to some $f \in l_2$. We will break the proof into three steps:

(10)

Step 1, $f(k) = \lim_{n \rightarrow \infty} p_n(k)$ exists in \mathbb{R} for each k .

To see why, note that $|p_n(k) - p_m(k)| \leq \|p_n - p_m\|_2$ for any k , and hence $\{p_n(k)\}_{n=1}^{\infty}$ is Cauchy in \mathbb{R} for each k . Thus, f is the obvious candidate for the limit of $\{p_n\}$, but we still have to show that the convergence takes place in the metric space h_2 ; that is, we need to show that $f \in h_2$ and that $\|p_n - f\|_2 \rightarrow 0$ (as $n \rightarrow \infty$).

Step 2. $f \in h_2$; that is, $\|f\|_2 < \infty$

We know that $\{p_n\}$ is bounded in h_2 (why?); say $\|p_n\|_2 \leq B$ for all n . Thus, for any fixed $N < \infty$ we have:

$$\sum_{k=1}^N |f(k)|^2 = \lim_{n \rightarrow \infty} \sum_{k=1}^N |p_n(k)|^2 \leq B^2.$$

Since this holds for any N , we get that $\|f\|_2 \leq B$

Step 3. Now we repeat Step 2 (more or less) to show that $p_n \rightarrow f$ in h_2 .

Given $\epsilon > 0$, choose n_0 so that $\|p_n - p_m\|_2 < \epsilon$ whenever $m, n \geq n_0$.

Then, for any N and any $n \geq n_0$,

$$\sum_{k=1}^N |f(k) - p_n(k)|^2 = \lim_{m \rightarrow \infty} \sum_{k=1}^N |p_m(k) - p_n(k)|^2 \leq \epsilon^2$$

Since this holds for any N , we have $\|f - p_n\|_2 \leq \epsilon$ for all $n \geq n_0$

(11)

That is, $f_n \mapsto f$ in k_2 .

Ex. (a) Just having a candidate for a limit is not enough. Consider the sequence $\{f_n\}$ in k_2 defined by $f_n = (1, \dots, 1, 0, \dots)$ where the first n entries are 1 and the rest are 0. The "obvious" limit is $f = (1, 1, \dots)$ (all 1), but $\|f - f_n\|_\infty = 1$ for all n . What's wrong?

(b) Worse still, sometimes the "obvious" limit is not even in the space. Consider the same sequence as in (a) and note that each f_n is actually an element of c_0 . This time, the natural candidate f is not in c_0 . Again, what's wrong?

As you can see, there can be a lot of details to check in a proof of completeness, and it would be hardly to have at least a few easy cases available. For example, when is a subset of a complete space complete? The answer is given as:

Thm: Let (M, d) be a complete metric space and let A be a subset of M . Then (A, d) is complete if and only if A is closed in M .

Proof: First suppose that (A, d) is complete, and let $\{x_n\}$ be a sequence in A that converges to some point $x \in M$. Then

(12)

$\{x_n\}$ is Cauchy in (A, d) and so converges to some point of A . That is, we must have $x \in A$ and, hence, A is closed.

Next, suppose that $\{x_n\}$ is a Cauchy sequence in (A, d) .

Then $\{x_n\}$ is also Cauchy in (M, d) (Why?). Hence, $\{x_n\}$ converges to some point $x \in M$. But A is closed and so, in fact, $x \in A$. Thus, (A, d) is complete.

Ex. (a) $[0, 1]$, $[0, \infty)$, \mathbb{N} , and \mathbb{A} are all complete.

(b) It follows from the theorem above that if a metric space is both complete and totally bounded, then every sequence in M has a convergent subsequence. In particular, any closed, bounded subset of \mathbb{R} is both complete and totally bounded. Thus, for example, every sequence in $[a, b]$ has a convergent subsequence. As you can easily imagine, the interval $[a, b]$ is a great place to do analysis! We will pursue the consequences of this felicitous combination of properties when we explore compact sets.

Our next result underlines the fact that complete spaces have a lot in common with \mathbb{R} .

(13)

Thm: For any metric space (M, d) , the following statements are equivalent:

(i) (M, d) is complete.

(ii) (The Nested Set Theorem) Let $F_1 \supset F_2 \supset \dots$ be a decreasing sequence of non-empty closed sets in M with $\text{diam}(F_n) \rightarrow 0$. Then, $\bigcap_{n=1}^{\infty} F_n \neq \emptyset$ (in fact, it contains exactly one point).

(iii) (The Bolzano-Weierstrass Theorem) Every infinite, totally bounded subset of M has a limit point in M .

Proof: (i) \Rightarrow (ii) Given $\{F_n\}$ as in (ii), choose $x_n \in F_n$ for each n . Then, since the F_n decrease, $\{x_k : k \geq n\} \subset F_n$ for each n , and hence $\text{diam} \{x_k : k \geq n\} \rightarrow 0$ as $n \rightarrow \infty$. That is, $\{x_n\}$ is Cauchy. Since M is complete, we have $x_n \rightarrow x$ for some $x \in M$. But the F_n are closed, and so we must have $x \in F_n$ for all n . Thus, $\bigcap_{n=1}^{\infty} F_n \neq \emptyset$.

(ii) \Rightarrow (iii) Let A be an infinite, totally bounded subset of M . Recall that we have shown that A contains a Cauchy sequence $\{x_n\}$ comprised of distinct points ($x_n \neq x_m$ for $n \neq m$). Now, setting $A_n = \{x_k : k \geq n\}$, we get $A \supset A_1 \supset A_2 \supset \dots$, each A_n is nonempty (even infinite), and $\text{diam}(A_n) \rightarrow 0$. That is, (ii) almost applies. But, clearly, $\overline{A_n} \supset \overline{A_{n+1}} \neq \emptyset$ for each n , and

(14)

$\text{diam}(\bar{A}_n) = \text{diam}(A_n) \rightarrow 0$ as $n \rightarrow \infty$. Thus there exists an $x \in \bigcap_{n=1}^{\infty} \bar{A}_n \neq \emptyset$. Now $x_n \in A_n$ implies that $d(x_n, x) \leq \text{diam}(\bar{A}_n) \rightarrow 0$. That is, $x_n \rightarrow x$ and so x is a limit point of A .

(iii) \Rightarrow (i) Let $\{x_n\}$ be Cauchy in (M, d) . We just need to show that $\{x_n\}$ has a convergent subsequence. Now, the set $A = \{x_n : n \geq 1\}$ is totally bounded (why?). If A happens to be finite, we are done. (why?) Otherwise, (iii) tells us that A has a limit point $x \in M$. It follows that some subsequence of $\{x_n\}$ converges to x . (why?)

We are finally ready to study compact metric spaces.

Def: A metric space (M, d) is said to be compact if it is both complete and totally bounded.

Ex. (a) A subset K of \mathbb{R} is compact iff K is closed and bounded. This fact is usually referred to as the Heine-Borel theorem. Hence, a closed bounded interval $[a, b]$ is compact. Also, the Cantor set Δ is compact. The interval $(0, 1)$, on the other hand, is not compact.

(b) A subset K of \mathbb{R}^n is compact iff K is closed and bounded. (Why?)

(c) It is important that we not confuse the first two examples with the general case. Recall that the set $\{e_n : n \geq 1\}$ is closed and bounded in l_∞ but not totally bounded - hence not compact. Taking this a step further, notice that the closed ball $\{x : \|x\|_\infty \leq 1\}$ in l_∞ is not compact, whereas any closed ball in \mathbb{R}^n is compact.

(d) A subset of a discrete space is compact iff it is finite (Why?)

We now present a sequential characterization of compactness:

Thm: (M, d) is compact iff and only iff every sequence in M has a subsequence that converges to a point in M .

Proof:

$$\left\{ \begin{array}{l} \text{totally bounded} \\ + \\ \text{complete} \end{array} \right\} \iff \left\{ \begin{array}{l} \text{every sequence in } M \text{ has} \\ \text{a Cauchy subsequence} \\ + \\ \text{Cauchy sequences converge} \end{array} \right\}$$

Corollary: Let A be a subset of a metric space M , iff A is compact, then A is closed in M . If M is compact and A is closed, then A is compact.

Proof: Suppose that A is compact, and let $\{x_n\}$ be a

(16)

sequence in A that converges to a point $x \in M$. Then, from the above theorem, $\{x_n\}$ has a subsequence that converges in A , and hence we must have $x \in A$. Thus, A is closed.

Next, suppose that M is compact and that A is closed in M . Given an arbitrary sequence $\{x_n\}$ in A . Theorem above supplies a subsequence of $\{x_n\}$ that converges to a point $x \in M$. But since A is closed, we must have $x \in A$. Thus, A is compact.

Compactness is a topological property. That is, the image of a compact set under a continuous function must be compact.

Thm: Let $f: (M, d) \rightarrow (N, \rho)$ be continuous, & K is compact in M , then $f(K)$ is compact in N .

Proof: Let $\{y_n\}$ be a sequence in $f(K)$. Then $y_n = f(x_n)$ for some sequence $\{x_n\}$ in K . But, since K is compact, $\{x_n\}$ has a convergent subsequence, say $x_{n_k} \rightarrow x \in K$. Then, since f is continuous, $y_{n_k} = f(x_{n_k}) \rightarrow f(x) \in f(K)$. Thus $f(K)$ is compact.

The theorem above gives us a wealth of useful information.

In particular, it tells us that real-valued continuous functions on compact spaces are quite well behaved:

Corollary: Let (M, d) be compact. If $f: M \rightarrow \mathbb{R}$ is continuous, then f is bounded. Moreover, f attains its maximum and minimum values.

Proof: $f(M)$ is compact in \mathbb{R} ; hence it is closed and bounded. Moreover, $\sup f(M)$ and $\inf f(M)$ are actually elements of $f(M)$. (Why?) That is, there exist $x, y \in M$ such that $f(x) \leq f(t) \leq f(y)$ for all $t \in M$. (In this case we would write $f(x) = \min_{t \in M} f(t)$ and $f(y) = \max_{t \in M} f(t)$).

Corollary: If $f: [a, b] \rightarrow \mathbb{R}$ is continuous, then the range of f is a compact interval $[c, d]$ for some $c, d \in \mathbb{R}$.

Corollary: If M is a compact metric space, then $\|f\|_{\infty} = \max_{t \in M} |f(t)|$ defines a norm on $C(M)$, the vector space of continuous real-valued functions on M .

It appears that compactness is the analogue of finite sets. To better appreciate this, we will need a slightly more esoteric characterization of compactness. A bit of preliminary detail-checking will ease the transition.

Lemma: In a metric space M , the following are equivalent:

(a) \mathcal{L} is any collection of open sets in M with

$\bigcup \{G : G \in \mathcal{L}\} \supset M$, then there are finitely many sets $G_1, \dots, G_n \in \mathcal{L}$ with $\bigcup_{i=1}^n G_i \supset M$.

(b) \mathcal{F} is any collection of closed sets in M such that

$\bigcap_{i=1}^n F_i \neq \emptyset$ for all choices of finitely many sets $F_1, \dots, F_n \in \mathcal{F}$,

then $\bigcap \{F : F \in \mathcal{F}\} \neq \emptyset$

You should attempt to prove the above lemma yourself.

As you might guess, De Morgan's laws do all of the work.

The first condition is usually paraphrased by saying, in less than perfect English, "every open cover has a finite subcover."

The second condition is abbreviated by saying "every collection of closed sets with 'finite intersection property' has nonempty intersection."

These may at first seem to be unwieldy statements to work with, but each is worth the trouble. Here's why we care: Condition (a) implies

that M is totally bounded because, for any $\epsilon > 0$, the collection $\mathcal{L} = \{B_\epsilon(x) : x \in M\}$ is an open cover for M .

Condition (b) implies that M is complete because it easily implies the nested set theorem. Put the two together and we've got our new characterization of compactness.

Thm: M is compact if and only if it satisfies either (hence both) conditions (a) or (b) in the lemma.

Proof: As noted above, condition (a) and (b) imply that M is totally bounded and complete, hence compact. So we need to show that compactness will imply, say (a). To this end, suppose that M is compact, and suppose that \mathcal{L} is an open cover for M that admits no finite subcover. We will work toward a contradiction.

Now M is totally bounded, so M can be covered by finitely many closed sets of diameter at most 1 . It follows that at least one of these, call it A_1 , cannot be covered by finitely many sets from \mathcal{L} . Certainly $A_1 \neq \emptyset$ (since the empty set is easy to cover!) Note that A_1 must be infinite. Next, A_1 is totally bounded, so A_1 can be covered by finitely many closed sets of diameter at most $\frac{1}{2}$. At least one of these, call it A_2 , cannot be covered by finitely many sets from \mathcal{L} .

(20)

Continuing, we get a decreasing sequence $A_1 \supset A_2 \supset \dots \supset A_n \supset \dots$, where A_n is closed, nonempty (infinite, actually), has diam $A_n \leq \frac{1}{n}$, and cannot be covered by finitely many sets from \mathcal{L} .

Now here's the fly in the ointment! Let $x \in \bigcap_{n=1}^{\infty} A_n (\neq \emptyset$ because M is complete). Then $x \in G \in \mathcal{L}$ for some G (since \mathcal{L} is an open cover) and so, since G is open, $x \in B_\epsilon(x) \subset G$ for some $\epsilon > 0$. But for any n with $1/n < \epsilon$ we would then have $x \in A_n \subset B_\epsilon(x) \subset G$. That is, A_n is covered by a single set from \mathcal{L} . This is the contradiction we were looking for.

Corollary: M is compact if and only if every decreasing sequence of nonempty closed sets has nonempty intersection; that is, if and only if, whenever $F_1 \supset F_2 \supset \dots$ is a sequence of nonempty closed sets in M , we have $\bigcap_{n=1}^{\infty} F_n \neq \emptyset$.

Proof: The forward implication is clear from the above theorem.

So suppose that every nested sequence of nonempty closed sets in M has nonempty intersection, and let $\{x_n\}$ be a sequence in M . Then there is some point x in the nonempty set $\bigcap_{n=1}^{\infty} \overline{\{x_k : k \geq n\}}$. (Why?) It follows that some subsequence of $\{x_n\}$ must converge to x .