

Analysis, Part I: Metric Spaces
Jay Havaldar

The general idea is to begin by defining the notion of convergence for arbitrary spaces.

Definition: A **metric space** is a set X equipped with a function $d : X \times X \rightarrow [0, \infty]$ such that: - $d(x, y) = 0$ iff $x = y$ (Positive definiteness) - $d(x, y) = d(y, x)$ (Symmetry) - $d(x, z) \leq d(x, y) + d(y, z)$. (Triangle Inequality)

We now define convergence:

Definition: A sequence x_n converges to x with respect to a metric d if $\lim_{n \rightarrow \infty} d(x_n, x)$ exists and is equal to 0.

It is not hard to check that a sequence converges in the Euclidean, taxi-cab, or sup norm metric iff each of its components converges.

Lemma

Limits are unique in a metric space.

Since a sequence gets arbitrarily close to its limit, we can prove that if there are two limits, they are also arbitrarily close -- and indeed so they must be equal.

0.1 Point-Set Topology

Definition: A **neighborhood** of x_0 in X is a set $B(x_0, r) = \{y \in X \mid d(x_0, y) < r, \text{ for some } r > 0\}$. This is sometimes called an open ball.

For later on, we define the useful notion of a limit point.

Definition: A limit point x of E is a point such that every neighborhood of x intersects E in some point other than x itself.

Definition: $x_0 \in \text{int}(E)$, the interior of E , if some neighborhood of x_0 is contained in E .

Definition: $x_0 \in \text{ext}(E)$ if there is some neighborhood of x_0 which does not intersect E .

Definition: A boundary point of E is neither in its interior or exterior. The set of all boundary points of E is denoted ∂E .

It is clear that around a boundary point, every neighborhood intersects both E and the complement of E , which is what we would expect.

Definition: The closure \bar{E} of E is the set of points $x \in X$ so that every neighborhood of x has nonempty intersection with E ; indeed this definition immediately suggests the following proposition.

Proposition

The following are equivalent for $x_0 \in X$ in a metric space X . - $x_0 \in \bar{E}$ - $x_0 \in \text{int}(E) \cup \partial E$ - There exists a sequence $x_n \in X$ which converges to x_0 .

Now, we can define the fundamental concepts of topology: open and closed sets.

Definition: E is closed if it contains its boundary.

Definition: E is open if it does not intersect its boundary. Equivalently, if $E = \text{int}(E)$.

Proposition

- E is closed iff it contains all its limit points.
- A singleton set is automatically closed.
- Finite unions and infinite (or finite) intersections of closed sets are closed.
- Finite intersections and infinite (or finite) unions of open sets are open.
- $\text{int}E$ is the largest open set which is contained in E .
- \bar{E} is the smallest closed set which contains E .

All these facts are proven in any topology book, or in my topology notes¹.

0.2 Relative Topology

Note that open sets are determined by both the metric itself and by the choice of ambient space. For example, an open interval of the x -axis is open when considered as a part of \mathbb{R} , but considered in the plane it is not open.

Definition: Given a metric space X and $Y \subseteq X$, we call E relatively open with respect to Y if it is open with respect to the metric space induced on Y by the metric of X .

Proposition

- E is relatively open with respect to Y iff $E = V \cap Y$ for some open set $V \subset X$.
- E is relatively closed with respect to Y iff $E = K \cap Y$ for some closed set $K \subset X$.

This is exactly what we do when constructing the subspace topology, as discussed in the aforementioned notes.

0.3 Cauchy Sequences

Lemma

If a sequence converges, every subsequence converges as well.

If we pick a large enough N , then for $n \geq N$, every element of any subsequence will get arbitrarily close to a limit.

¹<https://jhavaldar.github.io/notes/2017/08/04/top1.html>

Proposition

Suppose we have a sequence

x_n following are equivalent: - For every $\epsilon > 0, N \in \mathbb{Z}$, there exists an $n \geq N$ so that $d(x_n, L) < \epsilon$. - There exists a subsequence of x_n which converges to L .

The backwards direction is obvious. For the forward direction, we start by constructing a convergent subsequence. For each $N > 0$, and each natural number k , there exists an integer $M > k$ so that:

$$d(a_M, L) < \frac{1}{N}$$

So $M(k, N)$ exists for each choice of k, N . We define the sequence:

$$m_0 = 0, m_1 = M(m_0, 1), m_2 = M(m_1, 2) \dots$$

This is a subsequence of

x_n by definition. And for each $\epsilon < \frac{1}{N}$, we have constructed x_{m_N} which lies within distance ϵ of L . Thus, we have found a subsequence which converges to L .

We now move on to an important class of sequences.

Definition:

x_n is called a Cauchy sequence if for each $\epsilon > 0$, there exists an integer N so that if $j, k \geq N$, then $d(x_j, x_k) < \epsilon$.

Lemma

A convergent sequence is Cauchy.

For each epsilon, we pick N_ϵ so that for all $n \geq N_\epsilon$, $d(x_n, L) < \epsilon$. But then by the triangle inequality, for $j, k \geq N_{\epsilon/2}$, we have:

$$d(x_j, x_k) \leq d(x_j, L) + d(x_k, L) \leq \epsilon$$

Proposition

Every subsequence of a Cauchy sequence is Cauchy.

This proof is very similar to the proof that every subsequence of a convergent sequence converges.

We remark here that not every Cauchy sequence converges, so the converse of the above lemma is not in general true. For example, take the successive truncated decimal approximations of π ; this sequence is Cauchy, but it does not converge in \mathbb{Q} . We naturally ask, in what kind of a space is a Cauchy sequence always convergent?

Lemma

Let x_n be a Cauchy sequence in a metric space X . Suppose some subsequence converges to $L \in X$. Then x_n converges to L .

We noted that if a sequence converges, every subsequence converges. Here, we have something like a converse. If a Cauchy sequence has a convergent subsequence, then the sequence converges as well.

We now explore the conditions under which Cauchy sequences and convergent sequences are equivalent.

Definition: A metric space X is complete if every Cauchy sequence converges in X

Proposition

Let X be a metric space, and let Y be a subspace with the induced metric. Then, if Y is complete, then Y is closed in X . Conversely, if X is complete, and Y is closed, then Y is complete.

Closed subsets of complete spaces are complete; complete subsets of arbitrary spaces are closed. In a sense, then, complete metric spaces are intrinsically closed, regardless of the ambient context; and if we take closed subsets of complete spaces, this property is preserved.

To prove the first proposition, suppose Y is complete. Take any convergent sequence in X . Since convergent sequences are Cauchy, its limit must be in Y . Therefore, Y contains all its limit points and it is closed.

To prove the second proposition, suppose that Y is closed, and X is complete. Take a Cauchy sequence in Y . Since X is complete, the sequence converges in X . Therefore, the limit of this Cauchy sequence lies in \bar{Y} ; and since Y is closed, indeed the limit lies in Y . Therefore, the Cauchy sequence in Y converges in Y .

Note that proving both statements tells us that in a complete metric space, closed sets are complete and vice versa. In general, however, completeness is strictly stronger than the closed condition.

0.4 Compact Metric Spaces

Definition: A metric space is compact iff every sequence has a convergent subsequence.

Compactness, unlike some of the other properties mentioned, does not depend on the ambient context. We'll see why in a bit. First, we explain some of the conditions associated with compactness.

Definition: Let $Y \subset X$. Y is bounded if there exists an open ball around some point $x \in X$ which entirely contains Y .

Proposition

If X is compact, then it is complete and bounded.

Let

x_n be a Cauchy sequence in a compact space. Then, there exists a convergent subsequence. By an earlier lemma, this means that x_n converges as well. Therefore, X is complete.

To prove the boundedness condition, we use the triangle equality. Pick an arbitrary $x \in X$. If X is not bounded, then we can construct a sequence x_n so that $d(x_n, x) > n$. Fix another integer N . Then, for $m > N$, we have:

$$d(x, x_m) \leq d(x, x_N) + d(x_m, x_N) \quad d(x, x_m) - d(x, x_N) \leq d(x_m, x_N)$$

Since we fixed N , we can pick m large enough that the quantity on the left hand side blows up. It is not hard to see that this sequence has no convergent subsequence, since no subsequence is Cauchy; we fail to meet the condition that $d(x_m, x_N)$ goes to 0.

Corollary

If X is a metric space, and Y is compact in X , then Y is closed and bounded.

Completeness implies closedness, and so we are done.

Theorem

In Euclidean space, a set E is compact iff it is closed and bounded.

So the converse of the above theorem holds true not in general, but in Euclidean spaces. This is the famous Heine-Borel theorem from real analysis. The proof of the converse goes something like this: in a closed and bounded space, every sequence has a monotone subsequence, and such subsequences must converge.

We now have an alternative definition of compactness, which may be considered the "true" definition since it differs in some kinds of topological spaces from the one we constructed earlier. You shouldn't feel bad if this definition is puzzling; it took a lot of work to find the right condition to discuss these notions.

Theorem

Let $Y \subseteq X$ be a compact subspace of a metric space. Let $V_\alpha, \alpha \in A$ be an open cover for Y , i.e. a collection of (possibly finite) open sets whose union contains Y . Then there is a finite subset $F \subseteq A$ so that:

$$Y \subseteq \bigcup_{\alpha \in F} V_\alpha$$

More succinctly: every open cover of Y has a finite subcover.

Proof We assume that there exists no finite subcover. Then, we pick $y \in Y$ arbitrarily. We know that some neighborhood of y is contained in V_α for some α . We define:

$$r(y) = \sup\{r \mid B(y, r) \subset V_\alpha \text{ for some } \alpha\}$$

So $r(y)$ is the size of the largest ball around y which is contained in a single element of the open cover. We define:

$$r_0 = \inf\{r(y) \mid y \in Y\}$$

So r_0 is the smallest radius for which a ball of size r_0 around **every** $y \in Y$ is contained in a single element of the open cover.

There are two cases here: $r_0 = 0$, or $r_0 > 0$.

Case 1:

Let $r_0 = 0$. For each positive integer n , we can pick (by the axiom of choice) some element y_n so that $r(y_n) < \frac{1}{n}$. The limit of the sequence $r(y_n)$ must be zero.

Since Y is compact, we have a convergent subsequence y_{n_k} , which converges to some $L \in Y$. If we pick small $\epsilon > 0$ we can have $B(L, \epsilon) \subset V_\alpha$ (since V_α is open).

Furthermore, there is some integer N so that for all $n_k > N$, $y_{n_k} \in B(L, \frac{\epsilon}{2})$.

Putting these facts together, we pick an arbitrary $y \in B(y_{n_k}, \frac{\epsilon}{2})$; by the triangle inequality, we have for $n_k \geq N$:

$$d(y, L) \leq d(y, y_{n_k}) + d(y_{n_k}, L) \leq \epsilon$$

So indeed this means that $B(y_{n_k}, \frac{\epsilon}{2}) \subseteq B(L, \epsilon) \subseteq V_\alpha$. But this means that $r(y_{n_k}) \geq \frac{\epsilon}{2}$ for each $n_k \geq N$. This contradicts our earlier assumption that the limit of the sequence $r(y_n)$ must be zero.

Case 2:

In this case $r_0 > 0$. So in particular, $r(y) > r_0/2$ for each $y \in Y$. Equivalently, for each $y \in Y$, there is some α so that $B(y, r_0/2) \subseteq V_\alpha$.

We pick $y_1 \in Y$ and construct an open ball around y_1 of radius $r_0/2$. This ball cannot cover Y , since that would mean we have a finite subcover V_α .

Since the above set does not contain Y , we can pick y_2 from outside the set (i.e. farther than distance $r_0/2$ away from y_1) and add to our construction an open ball around y_2 of radius $r_0/2$. Our construction still cannot cover Y , since that would mean we have a finite subcover.

We continue so on to get a sequence y_n so that $d(y_k, y_j) \geq r_0/2$ for $k > j$. This sequence is by construction not Cauchy, and no subsequence is Cauchy. So there is no convergent subsequence; indeed, this means that Y is not compact.

Corollary

Let X be a metric space. Take K_1, K_2, \dots to be nonempty compact nesting sets, so that: $K_1 \supset K_2 \supset K_3 \supset \dots$. Then, $\bigcap_{n=1}^{\infty} K_n$ is non-empty.

Let $V_n = K_1 - K_n$ and assume that $\bigcap_{n=1}^{\infty} K_n$ is empty. Then, that means that for each $x \in K_1$, there exists some j so that $x \notin K_j \implies x \in V_j$. Therefore, V_n form an open cover for K_1 .

Since K_1 is compact, V_n have a finite subcover. Of the finite subcover, we pick the one with the smallest index V_m (which necessarily contains the union of all the elements of the subcover). Then, $K_1 = V_m = K_1 - K_m$. But this means that $K_m = \emptyset$, which is a contradiction.

Finally, we conclude with some other nice properties of compact spaces.

Theorem

Let X be a metric space. Then the following are true: - If Y is compact in X , and $Z \subset Y$, then Z is compact in Y iff Z is closed. - If Y_1, \dots, Y_n are a finite collection of compact sets, then their union is compact. - Every finite subset of X is compact.

The first theorem tells us that in compact spaces, compactness is equivalent to being closed, whereas previously we only proved that compactness implies completeness (which is a strictly stronger condition than being closed). We can think of this as analogous to the previous theorem that in complete spaces, complete sets are compact and vice versa.