

Topology and Topological Spaces

Mathematical spaces such as vector spaces, normed vector spaces (Banach spaces), and metric spaces are generalizations of ideas that are familiar in \mathbb{R} or in \mathbb{R}^n . For example, the various norms in \mathbb{R}^n , and the various metrics, generalize from the Euclidean norm and Euclidean distance. We generalize further, from \mathbb{R}^n to other sets, such as the set of all bounded real sequences, ℓ^∞ , or the set of continuous functions on the unit interval, $C([0, 1])$. We always generalize in such a way that the original version of the idea is a special case of the more general version. We make up our definition of the idea's general version to capture what seem to be the essential properties of the original, special version. This is exactly what we did in our definitions of vector space, norm, and metric, for example.

Let's focus, in particular, on the notion of a metric space, which is a set, together with a metric defined on the set. This concept is motivated by the more special notion of a norm. The norm is a concept of length that's defined only in vector spaces, but not for sets that aren't endowed with a vector space structure. For any norm, we also have the associated concept of distance in the vector space: $d(x, x') = \|x - x'\|$. But when we distilled the essential properties of distance, which we wrote as (D1)-(D4), and used them as the *definition* of a metric, we could then define metrics without any reference to a norm, and could therefore define them on any set.

A very good example is the discrete metric on any set X — *i.e.*, $d(x, x') = 1$ for all distinct elements x and x' in X . If X doesn't have a vector space structure, then there is no such thing as a norm on X , so the discrete metric in such a case certainly doesn't come from a norm. And even if X is a vector space, such as \mathbb{R}^n , the discrete metric is still well defined, and it's still not generated by any norm on \mathbb{R}^n . (Why not?)

Now let's consider the notions of open and closed sets, convergence of sequences, and continuity of functions. All these notions have familiar definitions in \mathbb{R}^n , in terms of norms and open balls. (And these are already generalizations from \mathbb{R} , where these concepts are defined in terms of absolute value and open intervals.) Once we have the more general notion of a metric, and a metric space, we're able to make essentially the same definitions of open and closed sets, convergence, and continuity in metric spaces — and thereby deal with all kinds of new situations, often ones in which (as with the discrete metric) the norm has no meaning, so if we hadn't generalized to metric spaces, we would not have been able to define the concepts of open balls, open sets, convergence, and continuity in these situations.

The idea of topology, and a topological space, is to generalize many of the ideas from metric spaces to situations where we don't have a metric, but where open sets, convergence, and continuity can still be defined — a parallel to defining these ideas when we don't have a norm but we do have a metric. When we're working with a set X that does have a metric, we defined the open sets in X in terms of open balls: a set is open if it contains an open ball around each of its points. We then proved the following theorem about the open sets:

Theorem: In a metric space (X, d) :

(O1) \emptyset and X are open sets.

(O2) If S_1, S_2, \dots, S_n are open sets, then $\bigcap_{i=1}^n S_i$ is an open set.

(O3) Let A be an arbitrary set. If S_α is an open set for each $\alpha \in A$, then $\bigcup_{\alpha \in A} S_\alpha$ is an open set. In other words, the union of any collection of open sets is open.

Suppose, for a given metric space (X, d) , we denote the collection of all the open sets in X by \mathcal{T} — *i.e.*, $\mathcal{T} := \{S \subseteq X \mid S \text{ is open}\}$. Then we can rewrite the properties (O1)-(O3) as

(O1) \emptyset and X are in \mathcal{T} .

(O2) If S_1, S_2, \dots, S_n are in \mathcal{T} , then $\bigcap_{i=1}^n S_i$ is in \mathcal{T} .

(O3) Let A be an arbitrary set. If S_α is in \mathcal{T} for each $\alpha \in A$, then $\bigcup_{\alpha \in A} S_\alpha$ is in \mathcal{T} . In other words, the union of any collection of sets in \mathcal{T} is itself in \mathcal{T} .

Now, for any set X , without any reference to a metric on X , if we have a collection \mathcal{T} of subsets of X that has the three properties (O1)-(O3), then we call \mathcal{T} a *topology* for X , and we call (X, \mathcal{T}) a *topological space*, and we call the elements of \mathcal{T} the open sets.

Definition: Let X be a set. A **topology** for X is a collection \mathcal{T} of subsets of X that satisfies (O1)-(O3); the pair (X, \mathcal{T}) is called a **topological space**.

Notice the exact parallel to our definition of a metric: we first identified the properties (D1)-(D4) of “norm-distance,” and then defined a metric to be any function that has those properties, whether it comes from a norm or not. Here we identify the three fundamental properties (O1)-(O3) of open sets in a metric space, and then define a collection of subsets of *any* set X to be open sets if the collection has those three properties, whether the sets are defined by a metric on X or not.

With this definition in hand, the above theorem can be restated as follows:

Theorem: In a metric space (X, d) , the collection of all the open sets is a topology for X .

Just as we earlier found that a normed vector space is also a metric space (if we define the metric from the norm in a standard way), and that there are also many metric spaces that are *not* normed spaces, here we find that a metric space is also a topological space (if we define the open sets from the metric in a standard way), and that there are also many topological spaces that are not metric spaces.

The last part of the preceding sentence — “there are also many topological spaces that are not metric spaces” — needs some clarification. What we mean when we say that a topological space (X, \mathcal{T}) is not a metric space is that there is no metric on the set X for which the set of open sets is \mathcal{T} . But the sentence says there are *many* such “non-metric” topological spaces — and we haven’t actually seen *any* such spaces. So here are two examples:

Example: The Indiscrete Topology

For any set X , let \mathcal{T} consist of just the two sets \emptyset and X — *i.e.*, $\mathcal{T} = \{\emptyset, X\}$. This topology cannot be generated by a metric. The example in which this is easiest to see is a two-point set, $X = \{a, b\}$. The only metric for this set X is the discrete metric. (More precisely, every metric on X has the form $d(a, b) = d(b, a) = c$ for some positive real number c , and all these metrics are easily seen to be equivalent to one another.) For the discrete metric, then, each of the four subsets of X is open (and also closed) — *i.e.*, the topology associated with the discrete metric is $\{\emptyset, \{a\}, \{b\}, X\}$ — and this is obviously different than the indiscrete topology, in which only \emptyset and X are open. Incidentally, it’s easy to see, using our open-sets definitions of convergence and continuity, that for any set X with the indiscrete topology, every sequence in X converges to every point of X ; and any function $f : S \rightarrow X$ with target space X is continuous (if its domain S is any topological space). This is dramatically different than the situation with metric spaces (and their associated topological spaces).

Example: The Lexicographic Topology

Let $X = [0, 1]^2$, the unit square in \mathbb{R}^2 , and let \succsim be the lexicographic order on X . Note that \succsim is a total order. Define an open interval in X to be any set of the form $\langle a, b \rangle := \{x \in X \mid a \prec x \prec b\}$. Let \mathcal{T} be the collection of all unions of intervals; *i.e.*, a subset of X is in \mathcal{T} , and thus an open set, if and only if it’s a (possibly infinite) union of intervals. It’s clear that \mathcal{T} satisfies the conditions (O1)-(O3), so it’s a topology for X , called the lexicographic topology, and (X, \mathcal{T}) is a lexicographic topological space.

The proof that there is no metric that yields the lexicographic topology is virtually identical to the proof we gave earlier that there is no utility function representation of the lexicographic preorder \succsim . We again proceed by contradiction: suppose that $d : X \times X \rightarrow \mathbb{R}_+$ is a metric for this topology — *i.e.*, the open sets defined by the d -open balls are exactly the sets in \mathcal{T} . For each $x \in X$, let $a(x) = d(\mathbf{0}, (x, 0))$ and $b(x) = d(\mathbf{0}, (x, 1))$. Then $a(x)$ and $b(x)$ are real numbers, and $a(x) < b(x)$, because for any real numbers r and \tilde{r} , the open ball $B(\mathbf{0}, r)$ is a proper subset of $B(\mathbf{0}, \tilde{r})$ if and only if $r < \tilde{r}$. As in our earlier, utility-function proof, there must be a rational number between $a(x)$ and $b(x)$, and we denote this rational number by $q(x)$. Now we have a distinct rational number $q(x)$ associated with every real number $x \in [0, 1]$, which we know is impossible (because the rational numbers are countable and the set $[0, 1]$ is not). Therefore our assumption that d is a metric cannot be true, which completes the proof.

Incidentally, note that in the lexicographic topology, *any* utility function that represents the lexicographic order will be continuous. (This topology is actually the smallest topology for which that's true, which is also easy to see.) But this fact is of no use, since we already know that there can't be a utility function representation of the lexicographic order.

As we've said above, we can use our metric-space open-sets definitions of convergence and continuity in topological spaces as well, just the same as we've done in metric spaces. We can do the same for compactness, but things don't work quite the same for compactness. An alternative definition of compactness is as follows:

Definition: Let (X, \mathcal{T}) be a topological space. An **open cover** of a subset $S \subseteq X$ is a collection \mathcal{C} of open sets the union of which **covers** S — *i.e.*, $S \subseteq \cup_{V \in \mathcal{C}} V$. A **subcover** of \mathcal{C} is a subset of \mathcal{C} — *i.e.*, a collection of some of the sets in \mathcal{C} . A subcover is finite if it consists of only a finite number of the sets in \mathcal{C} .

Definition: Let (X, \mathcal{T}) be a topological space. A subset $S \subseteq X$ is **compact** if every open cover of S contains a finite subcover.

Note that the space X itself is compact if every open cover of X has a finite subcover. This property is sometimes called the **Heine-Borel Property**. In metric spaces, we have the following theorem.

Heine-Borel Theorem: A metric space is compact (*i.e.*, it has the Heine-Borel Property) if and only if it has the Bolzano-Weierstrass Property.

Spaces with the Bolzano-Weierstrass Property are also called **sequentially compact**. So the Heine-Borel theorem says that in metric spaces, compactness is the same as sequential compactness. It's pretty straightforward to prove that compact metric spaces are sequentially compact, as we'll do below. It's somewhat more complicated to prove that sequentially compact metric spaces are compact.

Since we will be working only with metric spaces, the two properties are equivalent for us, and either one can be taken as the definition of compactness, with the other one as a characterization of compactness.

Here's a proof of the easy half of the Heine-Borel Theorem — if a metric space has the Heine-Borel Property, then it also has the Bolzano-Weierstrass Property, *i.e.*, compactness implies sequential compactness in a metric space. We omit the more complicated proof of the other half of the Heine-Borel Theorem.

Proof:

Let (X, d) be a metric space, and assume that X has the Heine-Borel Property. Let $\{x_n\}$ be a sequence in X ; we must show that $\{x_n\}$ has a convergent subsequence (equivalently, that it has a cluster point). Suppose it doesn't; then no point in X is a cluster point of $\{x_n\}$, so for every point $x \in X$, there is an open ball $B(x, \epsilon_x)$ about x that contains no more than a finite number of terms of $\{x_n\}$. The collection of all these open balls obviously covers X , since each $x \in X$ is at the center of one of the balls. Therefore there is a finite subcover — *i.e.*, X is covered by only a finite number of the balls. But each ball contains only a finite number of terms of $\{x_n\}$, so altogether this finite collection of balls (and therefore X itself) contains only a finite number of the terms of $\{x_n\}$. This is impossible, because $\{x_n\}$ has an infinite number of terms — a contradiction that establishes that $\{x_n\}$ must have a cluster point after all. \square

Here's an application of these concepts to decision theory — a simple and elegant proof that under relatively weak assumptions, a preference relation will have a maximal element on any compact set.

Definition: An irreflexive binary relation \succ on a set X is **acyclic** if it has no cycles — *i.e.*, no n -tuples (x_1, \dots, x_n) of elements of X that satisfy $x_1 \succ x_2 \succ x_3 \succ \dots \succ x_{n-1} \succ x_n \succ x_1$.

Note that acyclicity is weaker than transitivity: a transitive relation is obviously acyclic.

Bergstrom-Walker Theorem: Let X be a compact subset of a topological space and let \succ be an acyclic binary relation on X . If the \succ -lower-contour sets are all open, then there is a \succ -maximal element of X .

Proof:

For each $x \in X$, let $L(x)$ denote the lower-contour set of x , $\{z \in X \mid x \succ z\}$, and note that a subset $S \subseteq X$ has a maximal element if and only if the collection of all lower-contour sets, $\{L(x) \mid x \in S\}$, does not cover S . Suppose that X has no \succ -maximal element; then the collection of all lower-contour sets, $\{L(x) \mid x \in X\}$, is an open cover of X . Since X is compact, there is a finite subset $S \subseteq X$ such that $\{L(x) \mid x \in S\}$ covers X and *a fortiori* covers S . Hence the finite set S has no maximal element. It's easy to see, however, that for an acyclic relation \succ on X , every finite subset of X *must* have a \succ -maximal element. This contradiction establishes that X itself must have a \succ -maximal element. \square

We know that if \succsim is complete and transitive (a complete preorder), and if all of its strict upper- and lower-contour sets are open, then it can be represented by a continuous utility function $u(\cdot)$, and the Weierstrass Theorem applied to $u(\cdot)$ therefore ensures that \succsim will have a maximum element on any compact set. However, the theorem we've just proved tells us that under much weaker assumptions on \succsim — and even if \succsim is not representable — we can still be assured that \succsim will have a maximal element on any compact set.

Since a metric space is automatically a topological space, and since Euclidean space is a metric space, the theorem applies to compact subsets of \mathbb{R}^n , such as a typical budget set, and to lots of other feasible sets as well.