Chapter 1. Basic Point-Set Topology

One way to describe the subject of Topology is to say that it is qualitative geometry. The idea is that if one geometric object can be continuously transformed into another, then the two objects are to be viewed as being *topologically* the same. For example, a circle and a square are topologically equivalent. Physically, a rubber band can be stretched into the form of either a circle or a square, as well as many other shapes which are also viewed as being topologically equivalent. On the other hand, a figure eight curve formed by two circles touching at a point is to be regarded as topologically distinct from a circle or square. A qualitative property that distinguishes the circle from the figure eight is the number of connected pieces that remain when a single point is removed: When a point is removed from a circle what remains is still connected, a single arc, whereas for a figure eight if one removes the point of contact of its two circles, what remains is two separate arcs, two separate pieces.

The term used to describe two geometric objects that are topologically equivalent is *homeomorphic*. Thus a circle and a square are homeomorphic. Concretely, if we place a circle C inside a square S with the same center point, then projecting the circle radially outward to the square defines a function $f:C\to S$, and this function is continuous: small changes in x produce small changes in $f(x)$. The function f has an inverse $f^{-1}:S\to C$ obtained by projecting the square radially inward to the circle, and this is continuous as well. One says that f is a homeomorphism between C and S.

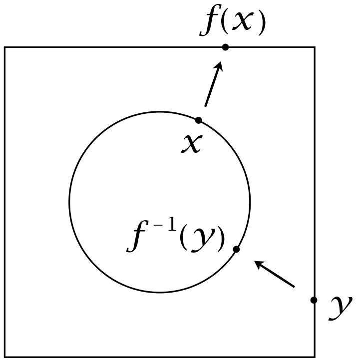

One of the basic problems of Topology is to determine when two given geometric objects are homeomorphic. This can be quite difficult in general.

Our first goal will be to define exactly what the ‘geometric objects’ are that one studies in Topology. These are called *topological spaces*. The definition turns out to be extremely general, so that many objects that are topological spaces are not very geometric at all, in fact.

Topological Spaces

Rather than jump directly into the definition of a topological space we will first spend a little time motivating the definition by discussing the notion of continuity of a function. One could say that topological spaces are the objects for which continuous functions can be defined.

For the sake of simplicity and concreteness let us talk about functions $f:\mathbb{R}\to\mathbb{R}$. There are two definitions of continuity for such a function that the reader may already be familiar with, the $\varepsilon\delta$ definition and the definition in terms of limits. But it is a third definition, equivalent to these two, that is the one we want here. This definition is expressed in terms of the notion of an open set in $\mathbb{R}$, generalizing the familiar idea of an open interval (a,b).

Definition. A subset O of $\mathbb{R}$ is *open* if for each point $x \in O$ there exists an interval (a,b) that contains x and is contained in O.

With this definition an open interval certainly qualifies as an open set. Other examples are:

- $\mathbb{R}$ itself is an open set, as are semi-infinite intervals (a,∞) and $(-\infty,a)$.
- The complement of a finite set in $\mathbb{R}$ is open.
- If A is the union of the infinite sequence $x_n = 1/n$, $n = 1,2,\cdots$, together with its limit 0 then $\mathbb{R} - A$ is open.
- Any union of open intervals is an open set. The preceding examples are special cases of this. The converse statement is also true: every open set O is a union of open intervals since for each $x \in O$ there is an open interval (a_x,b_x) with $x \in (a_x,b_x) \subset O$, and O is the union of all these intervals (a_x,b_x).
- The empty set $\varnothing$ is open, since the condition for openness is satisfied vacuously as there are no points x where the condition could fail to hold.

Here are some examples of sets which are not open:

- A closed interval $[a,b]$ is not an open set since there is no open interval about either a or b that is contained in $[a,b]$. Similarly, half-open intervals $[a,b)$ and $(a,b]$ are not open sets when $a < b$.
- A nonempty finite set is not open.

Now for the nice definition of a continuous function in terms of open sets:

Definition. A function $f:\mathbb{R}\to\mathbb{R}$ is *continuous* if for each open set O in $\mathbb{R}$ the inverse image $f^{-1}(O) = \{\, x \in \mathbb{R} \mid f(x) \in O \,\}$ is also an open set.

To see that this corresponds to the intuitive notion of continuity, consider what would happen if this condition failed to hold for a function f. There would then be an open set O for which $f^{-1}(O)$ was not open. This means there would be a point $x_0 \in f^{-1}(O)$ for which there was no interval (a,b) containing x_0 and contained in $f^{-1}(O)$. This is equivalent to saying there would be points x arbitrarily close to x_0

that are in the complement of $f^{-1}(O)$. For x to be in the complement of $f^{-1}(O)$ means that $f(x)$ is not in O. On the other hand, x_0 was in $f^{-1}(O)$ so $f(x_0)$ is in O. Since O was assumed to be open, there is an interval (c,d) about $f(x_0)$ that is contained in O. The points $f(x)$ that are not in O are therefore not in (c,d) so they remain at least a fixed positive distance from $f(x_0)$. To summarize: there are points x arbitrarily close to x_0 for which $f(x)$ remains a fixed positive distance away from $f(x_0)$. This certainly says that f is discontinuous at x_0.

This reasoning can be reversed. A reasonable interpretation of discontinuity of f at x_0 would be that there are points x arbitrarily close to x_0 for which $f(x)$ stays at least a fixed positive distance away from $f(x_0)$. Call this fixed positive distance ε. Let O be the open set $(f(x_0)-\varepsilon, f(x_0)+\varepsilon)$. Then $f^{-1}(O)$ contains x_0 but it does not any points x for which $f(x)$ is not in O, and we are assuming there are such points x arbitrarily close to x_0, so $f^{-1}(O)$ is not open since it does not contain all points in some interval (a,b) about x_0.

The definition we have given for continuity of functions $\mathbb{R}\to\mathbb{R}$ can be applied more generally to functions $\mathbb{R}^n\to\mathbb{R}^n$ and even $\mathbb{R}^m\to\mathbb{R}^n$ once one has a notion of what open sets in $\mathbb{R}^n$ are. The natural definition generalizing the case $n=1$ is to say that a set O in $\mathbb{R}^n$ is open if for each $x\in O$ there exists an open ball containing x and contained in O, where an open ball of radius r and center x_0 is defined to be the set of points x of distance less than r from x_0. Here the distance from x to x_0 is measured as in linear algebra, as the length of the vector $x-x_0$, the square root of the dot product of this vector with itself.

This definition of open sets in $\mathbb{R}^n$ does not depend as heavily on the notion of distance in $\mathbb{R}^n$ as might appear. For example in $\mathbb{R}^2$ where open balls become open disks, we could use open squares instead of open disks since if a point $x\in O$ is contained in an open disk contained in O then it is also contained in an open square contained in the disk and hence in O, and conversely, if x is contained in an open square contained in O then it is contained in an open disk contained in the open square and hence in O. In a similar way we could use many other shapes besides disks and squares, such as ellipses or polygons with any number of sides.

After these preliminary remarks we now give the definition of a topological space.

Definition. A *topological space* is a set X together with a collection $\mathcal{O}$ of subsets of X, called *open sets*, such that:

(1) The union of any collection of sets in $\mathcal{O}$ is in $\mathcal{O}$.
(2) The intersection of any finite collection of sets in $\mathcal{O}$ is in $\mathcal{O}$.

(3) Both $\varnothing$ and X are in $\mathcal{O}$.

The collection $\mathcal{O}$ of open sets is called a *topology* on X.

All three of these conditions hold for open sets in $\mathbb{R}$ as defined earlier. To check that (1) holds, suppose that we have a collection of open sets O_α where the index α ranges over some index set I, either finite or infinite. A point $x \in \bigcup_\alpha O_\alpha$ lies in some O_α, which is open so there is an interval (a,b) with $x \in (a,b) \subset O_\alpha$, hence $x \in (a,b) \subset \bigcup_\alpha O_\alpha$ so $\bigcup_\alpha O_\alpha$ is open. To check (2) it suffices by induction to check that the intersection of two open sets O_1 and O_2 is open. If $x \in O_1 \cap O_2$ then x lies in open intervals in O_1 and O_2, and there is a smaller open interval in the intersection of these two open intervals that contains x. This open interval lies in $O_1 \cap O_2$, so $O_1 \cap O_2$ is open. Finally, condition (3) obviously holds for open sets in $\mathbb{R}$.

In a similar fashion one can check that open sets in $\mathbb{R}^2$ or more generally $\mathbb{R}^n$ also satisfy (1)–(3).

Notice that the intersection of an infinite collection of open sets in $\mathbb{R}$ need not be open. For example, the intersection of all the open intervals $(-1/n, 1/n)$ for $n = 1, 2, \cdots$ is the single point $\{0\}$ which is not open. This explains why condition (2) is only for finite intersections.

It is always possible to construct at least two topologies on every set X by choosing the collection $\mathcal{O}$ of open sets to be as large as possible or as small as possible:

- The collection $\mathcal{O}$ of *all* subsets of X defines a topology on X called the *discrete topology*.
- If we let $\mathcal{O}$ consist of just X itself and $\varnothing$, this defines a topology, the *trivial topology*.

Thus we have three different topologies on $\mathbb{R}$, the usual topology, the discrete topology, and the trivial topology. Here are two more, the first with fewer open sets than the usual topology, the second with more open sets:

- Let $\mathcal{O}$ consist of the empty set together with all subsets of $\mathbb{R}$ whose complement is finite. The axioms (1)–(3) are easily verified, and we leave this for the reader to check. Every set in $\mathcal{O}$ is open in the usual topology, but not vice versa.
- Let $\mathcal{O}$ consist of all sets O such that for each $x \in O$ there is an interval $[a,b)$ with $x \in [a,b) \subset O$. Properties (1)–(3) can be checked by almost the same argument as for the usual topology on $\mathbb{R}$, and again we leave this for the reader to do. Intervals $[a,b)$ are certainly in $\mathcal{O}$ so this topology is different from the usual topology on $\mathbb{R}$. Every interval (a,b) is in $\mathcal{O}$ since it can be expressed as a union of an increasing sequence of intervals $[a_n, b)$ in $\mathcal{O}$. It follows that $\mathcal{O}$ contains all

sets that are open in the usual topology since these can be expressed as unions of intervals (a, b).

These examples illustrate how one can have two topologies $\mathcal{O}$ and $\mathcal{O}'$ on a set X, with every set that is open in the $\mathcal{O}$ topology is also open in the $\mathcal{O}'$ topology, so $\mathcal{O} \subset \mathcal{O}'$. In this situation we say that the topology $\mathcal{O}'$ is *finer* than $\mathcal{O}$ and that $\mathcal{O}$ is *coarser* than $\mathcal{O}'$. Thus the discrete topology on X is finer than any other topology and the trivial topology is coarser than any other topology. In the case $X = \mathbb{R}$ we have interpolated three other topologies between these two extremes, with the finite-complement topology being coarser than the usual topology and the half-open-interval topology being finer than the usual topology. Of course, given two topologies on a set X, it need not be true that either one is finer or coarser than the other.

Here is another piece of basic terminology:

Definition. A subset A of a topological space X is *closed* if its complement $X - A$ is open.

For example, in $\mathbb{R}$ with the usual topology a closed interval $[a, b]$ is a closed subset. Similarly, in $\mathbb{R}^2$ with its usual topology a closed disk, the union of an open disk with its boundary circle, is a closed subset.

Instead of defining a topology on a set X as a collection of open sets satisfying the three axioms, one could equally well consider the collection of complementary closed sets, and define a topology on X to be a collection of subsets called closed sets, such that the intersection of any collection of closed sets is closed, the union of any finite collection of closed sets is closed, and both the empty set and the whole set X are closed. Notice that the role of intersections and unions is switched compared with the original definition. This is because of the general set theory fact that that the complement of a union is the intersection of the complements, and the complement of an intersection is the union of the complements.

Interior, Closure, and Boundary

Consider an open disk D in the plane $\mathbb{R}^2$, consisting of all the points inside a circle C. We would like to assign precise meanings to certain intuitive statements like the following:

- C is the boundary of the open disk D, and also of the closed disk $D \cup C$.
- D is the interior of the closed disk $D \cup C$, and $D \cup C$ is the closure of the open disk D.

The key distinction between points in the boundary of the disk and points in its interior is that for points in the boundary, every open set containing such a point also contains points inside the disk and points outside the disk, while each point in the interior of the disk lies in some open set entirely contained inside the disk.

With this observation in mind let us consider what happens in general. Given a subset A of a topological space X, then for each point $x \in X$ exactly one of the following three possibilities holds:

(1) There exists an open set O in X with $x \in O \subset A$.
(2) There exists an open set O in X with $x \in O \subset X - A$.
(3) Every open set O with $x \in O$ meets both A and $X - A$.

Points x such that (1) holds form a subset of A called the *interior* of A, written $\mathrm{int}(A)$. The points where (2) holds then form $\mathrm{int}(X - A)$. Points x where (3) holds form a set called the *boundary* or *frontier* of A, written ∂A. The points x where either (1) or (3) hold are the points x such that every open set O containing x meets A. Such points are called *limit points* of A, and the set of these limit points is called the *closure* of A, written $\overline{A}$. Note that $A \subset \overline{A}$, so we have $\mathrm{int}(A) \subset A \subset \overline{A} = \mathrm{int}(A) \cup \partial A$, this last union being a disjoint union. We will use the symbol $\amalg$ to denote union of disjoint subsets when we want to emphasize the disjointness, so $A = \mathrm{int}(A) \amalg \partial A$ and $X = \mathrm{int}(A) \amalg \partial A \amalg \mathrm{int}(X - A)$.

As an example, in $\mathbb{R}$ with the usual topology the intervals (a, b), $[a, b]$, $[a, b)$, and $(a, b]$ all have interior (a, b), closure $[a, b]$ and boundary $\{a, b\}$. Similarly, in $\mathbb{R}^2$ with the usual topology, if A is the union of an open disk D with any subset of its boundary circle C then $\mathrm{int}(A) = D$, $\overline{A} = D \cup C$, and $\partial A = C$. For a somewhat different type of example, let $A = \mathbb{Q}$ in $X = \mathbb{R}$ with the usual topology on $\mathbb{R}$. Then $\mathrm{int}(A) = \varnothing$ and $\overline{A} = \partial A = \mathbb{R}$.

Proposition 1.1. *For every $A \subset X$ the following statements hold:*

(a) $\mathrm{int}(A)$ *is open.*
(b) $\overline{A}$ *is closed.*
(c) A *is open if and only if* $A = \mathrm{int}(A)$.
(d) A *is closed if and only if* $A = \overline{A}$.

Proof. (a) If x is a point in $\mathrm{int}(A)$ then there is an open set O_x with $x \in O_x \subset A$. We have $O_x \subset \mathrm{int}(A)$ since for each $y \in O_x$, O_x is an open set with $y \in O_x \subset A$ so $y \in \mathrm{int}(A)$. It follows that $\mathrm{int}(A) = \bigcup_x O_x$, the union as x ranges over all points of $\mathrm{int}(A)$. This is a union of open sets and hence open.

(b) Since $X = \operatorname{int}(A) \amalg \partial A \amalg \operatorname{int}(X - A)$, we have $\overline{A}$ as the complement of $\operatorname{int}(X - A)$, so $\overline{A}$ is closed, being the complement of an open set by part (a).

(c) If $A = \operatorname{int}(A)$ then A is open by (a). Conversely, if A is open then every $x \in A$ is in $\operatorname{int}(A)$ since we can take $O = A$ in condition (1). Thus $A \subset \operatorname{int}(A)$. The opposite inclusion always holds, so $A = \operatorname{int}(A)$.

(d) If $A = \overline{A}$ then A is closed by (b). Conversely, if A is closed then $X - A$ is open, so each point of $X - A$ is contained in an open set $X - A$ disjoint from A, which means that no point of $X - A$ is a limit point of A, or in other words we have $\overline{A} \subset A$. We always have $A \subset \overline{A}$, so $A = \overline{A}$. □

A small caution: Some authors use the term 'limit point' in a more restricted sense than we are using it here, requiring that every open set containing x contains points of A *other than x itself*. Other names for this more restricted concept that one sometimes finds are 'point of accumulation' and 'cluster point'.

One might have expected the definition of a limit point to be expressed in terms of convergent sequences of points. In an arbitrary topological space X it is natural to define $\lim_{n \to \infty} x_n = x$ to mean that for every open set O containing x the points x_n lie in O for all but finitely many values of n, or in other words there exists an $N > 0$ such that $x_n \in O$ for all $n > N$. It is obvious that $\lim_{n \to \infty} x_n = x$ implies that x a limit point of the subset of X formed by the sequence of points x_n. However there exist topological spaces in which a limit point of a subset need not be the limit of any convergent sequence of points in the subset. For subspaces X of $\mathbb{R}^n$ this strange behavior does not occur since if x is a limit point of a subset $A \subset X$ then for each $n > 0$ there is a point $x_n \in A$ in the open ball of radius $1/n$ centered at x, and for this sequence of points x_n we have $\lim_{n \to \infty} x_n = x$.

Basis for a Topology

Many arguments with open sets in $\mathbb{R}$ reduce to looking at what happens with open intervals since open sets are defined in terms of open intervals. A similar statement holds for $\mathbb{R}^2$ and $\mathbb{R}^n$ with open disks and balls in place of open intervals. In each case arbitrary open sets are unions of the special open sets given by open intervals, disks, or balls. This idea is expressed by the following terminology:

Definition. A collection $\mathcal{B}$ of open sets in a topological space X is called a *basis* for the topology if every open set in X is a union of sets in $\mathcal{B}$.

A topological space can have many different bases. For example, in $\mathbb{R}^2$ another basis besides the basis of open disks is the basis of open squares with edges parallel to the coordinate axes. Or we could take open squares with edges at 45 degree angles to the coordinate axes, or all open squares without restriction. Many other shapes besides squares could also be used.

If $\mathcal{B}$ is a basis for X and Y is a subspace of X, then we can obtain a basis for Y by taking the collection $\mathcal{B}_Y$ of intersections $Y \cap B$ as B ranges over all the sets in $\mathcal{B}$. This gives a basis for Y because an arbitrary open set in the subspace topology on Y has the form $Y \cap (\bigcup_\alpha B_\alpha) = \bigcup_\alpha (Y \cap B_\alpha)$ for some collection of basis sets $B_\alpha \in \mathcal{B}$. In particular this says that for any subspace X of $\mathbb{R}^n$, a basis for the topology on X is the collection of open sets $X \cap B$ as B ranges over all open balls in $\mathbb{R}^n$. For example, for a circle in $\mathbb{R}^2$ the open arcs in the circle form a basis for its topology.

If $\mathcal{B}$ is a basis for a topology on X, then $\mathcal{B}$ satisfies the following two properties:

(1) Every point $x \in X$ lies in some set $B \in \mathcal{B}$.

(2) For each pair of sets B_1, B_2 in $\mathcal{B}$ and each point $x \in B_1 \cap B_2$ there exists a set B_3 in $\mathcal{B}$ with $x \in B_3 \subset B_1 \cap B_2$.

The first statement holds since X is open and is therefore a union of sets in $\mathcal{B}$. The second statement holds since $B_1 \cap B_2$ is open and hence is a union of sets in $\mathcal{B}$.

Proposition 1.2. *If $\mathcal{B}$ is a collection of subsets of a set X satisfying* (1) *and* (2) *then $\mathcal{B}$ is a basis for a topology on X.*

The open sets in this topology have to be exactly the unions of sets in $\mathcal{B}$ since $\mathcal{B}$ is a basis for this topology.

Proof. Let $\mathcal{O}$ be the collection of subsets of X that are unions of sets in $\mathcal{B}$. Obviously the union of any collection of sets in $\mathcal{O}$ is in $\mathcal{O}$. To show the corresponding result for finite intersections it suffices by induction to show that $O_1 \cap O_2 \in \mathcal{O}$ if $O_1, O_2 \in \mathcal{O}$. For each $x \in O_1 \cap O_2$ we can choose sets $B_1, B_2 \in \mathcal{B}$ with $x \in B_1 \subset O_1$ and $x \in B_2 \subset O_2$. By (2) there exists a set $B_3 \in \mathcal{B}$ with $x \in B_3 \subset B_1 \cap B_2 \subset O_1 \cap O_2$. The union of all such sets B_3 as x ranges over $O_1 \cap O_2$ is $O_1 \cap O_2$, so $O_1 \cap O_2 \in \mathcal{O}$.

Finally, X is in $\mathcal{O}$ by (1), and $\varnothing \in \mathcal{O}$ since we can regard $\varnothing$ as the union of the empty collection of subsets of $\mathcal{B}$. $\square$

Terminology: Neighborhoods

We have frequently had to deal with open sets O containing a given point x. Such an open set is called a 'neighborhood' of x. Actually, it is useful to use the following broader definition:

Definition. A *neighborhood* of a point x in a topological space X is any set $A \subset X$ that contains an open set O containing x.

The more restricted kind of neighborhood can then be described as an open neighborhood.

As an example of the usefulness of this terminology, we can rephrase the condition for a point x to be a limit point of a set A to say that every open neighborhood of x meets A. The word 'open' here can in fact be omitted, for if every neighborhood of x meets A then in particular every open neighborhood meets A, and conversely, if every open neighborhood meets A then so does every other neighborhood since every neighborhood contains an open neighborhood.

Similarly, a boundary point of A is a point x such that every neighborhood of x meets both A and $X - A$.

Metric Spaces

The topology on $\mathbb{R}^n$ is defined in terms of open balls, which in turn are defined in terms of distance between points. There are many other spaces whose topology can be defined in a similar way in terms of a suitable notion of distance between points in the space. Here is the ingredient needed to do this:

Definition. A *metric* on a set X is a function $d : X \times X \to \mathbb{R}$ such that

(1) $d(x,y) \ge 0$ for all $x,y \in X$, with $d(x,x) = 0$ and $d(x,y) > 0$ if $x \ne y$.

(2) $d(x,y) = d(y,x)$ for all $x,y \in X$.

(3) $d(x,y) \le d(x,z) + d(z,y)$ for all $x,y,z \in X$.

This last condition is called the 'triangle inequality' because for the usual distance function in the plane it says that length of one side of a triangle is always less than or equal to the sum of the lengths of the other two sides.

Given a metric on X one defines the open ball of radius r centered at x to be the set $B_r(x) = \{ y \in X \mid d(x,y) < r \}$.

Proposition. *The collection of all balls $B_r(x)$ for $r > 0$ and $x \in X$ forms a basis for a topology on X.*

A space whose topology can be obtained in this way via a basis of open balls with respect to a metric is called a *metric space*.

Proof. First a preliminary observation: For a point $y \in B_r(x)$ the ball $B_s(y)$ is contained in $B_r(x)$ if $s \le r - d(x,y)$, since for $z \in B_s(y)$ we have $d(z,y) < s$ and hence $d(z,x) \le d(z,y) + d(y,x) < s + d(x,y) \le r$.

Now to show the condition to have a basis is satisfied, suppose we are given a point $y \in B_{r_1}(x_1) \cap B_{r_2}(x_2)$. Then the observation in the preceding paragraph implies that $B_s(y) \subset B_{r_1}(x_1) \cap B_{r_2}(x_2)$ for any $s \le \min\{r_1 - d(x_1,y), r_2 - d(x_2,y)\}$. □

Different metrics on the same set X can give rise to different bases for the same topology. For example, in $\mathbb{R}^2$ we have the usual metric defined in terms of lengths of vectors by $d(x,y) = |x-y|$ has 'balls' which are disks, but another metric whose 'balls' are squares is $d(x,y) = \max\{|x_1 - y_1|, |x_2, y_2|\}$ for $x = (x_1, y_1)$ and $y = (y_1, y_2)$. We will leave is to the reader to verify that this satisfies the properties of a metric. Another metric is $d(x,y) = |x_1 - y_1| + |x_2 - y_2|$, which has balls $B_r(X)$ that are also squares, but rotated 45 degrees from the squares in the previous metric.

Many important spaces in the subject of Topology are metric spaces, but the point of view of Topology is to ignore any particular choice of metric as much as possible, and just focus on the open sets, on the topology itself.

Subspaces

We turn now to a topic which ought to be simple, and seems simple enough at first glance, but turns out to be a source of many headaches until one finally becomes comfortable with it.

Given a topology $\mathcal{O}$ on a space X and a subset $A \subset X$, we would like to use the topology on X to define a topology $\mathcal{O}_A$ on A. There is an easy way to do this: Just define a set $O \subset A$ to be in $\mathcal{O}_A$ if there exists an open set O' in $\mathcal{O}$ such that $O = A \cap O'$. Axiom (1) holds since if $O_\alpha = A \cap O'_\alpha$ for sets $O'_\alpha \in \mathcal{O}$ then $\bigcup_\alpha O_\alpha = \bigcup_\alpha (A \cap O'_\alpha) = A \cap (\bigcup_\alpha O'_\alpha)$, so $\bigcup_\alpha O_\alpha$ is in $\mathcal{O}_A$. Axiom (2) is similar since $\bigcap_\alpha O_\alpha = \bigcap_\alpha (A \cap O'_\alpha) = A \cap (\bigcap_\alpha O'_\alpha)$, which is in $\mathcal{O}_A$ if there are just finitely many indices α. Axiom (3) is obvious.

The topology $\mathcal{O}_A$ on A is called the *subspace topology*, and A with this topology is called a *subspace* of X. For example, if we take X to be $\mathbb{R}^2$ with its usual topology, then every subset of $\mathbb{R}^2$ becomes a topological space. In particular, geometric figures such as circles and polygons can now be viewed as topological spaces. Likewise, geometric figures in $\mathbb{R}^3$ such as spheres and polyhedra become topological spaces, with the subspace topology from the usual topology on $\mathbb{R}^3$.

In case the space X is a metric space, any subset $A \subset X$ becomes a metric space by restricting the metric $X \times X \to \mathbb{R}$ to $A \times A$, since the three defining properties of a metric obviously still hold for the restricted distance function. The following Proposition gives some strong evidence that the subspace topology is a natural topology to use on subsets.

Proposition. *The metric topology on a subset A of a metric space X is the same as the subspace topology.*

Proof. Observe first that for a ball $B_r(x)$ in X, the intersection $A \cap B_r(x)$ consists of all points in A of distance less than r from x, so this is a ball in A regarded as a metric space in itself. For a collection of such balls $B_{r_\alpha}(x_\alpha)$ we have

$$A \cap \Big(\bigcup_\alpha B_{r_\alpha}(x_\alpha)\Big) = \bigcup_\alpha \big(A \cap B_{r_\alpha}(x_\alpha)\big)$$

The left side of this equation is a typical open set in A with the subspace topology, and the right side is a typical open set in the metric topology, so the two topologies coincide. □

A subspace $A \subset X$ whose subspace topology is the discrete topology is called a *discrete subspace* of X. This is equivalent to saying that for each point $x \in A$ there is an open set in X whose intersection with A is just x. For example, $\mathbb{Z}$ is a discrete subspace of $\mathbb{R}$, but $\mathbb{Q}$ is not discrete. The sequence $1/2, 1/3, 1/4 \cdots$ without its limit 0 is a discrete subspace of $\mathbb{R}$, but with 0 it is not discrete.

For a subspace $A \subset X$, a subset of A which is open or closed in A need not be open or closed in X. However, there are times when this is true:

Lemma. *For a subspace $A \subset X$ which is open in X, a subset $B \subset A$ is open in A if and only if it is open in X. This is also true when 'open' is replaced by 'closed' throughout the statement.*

Proof. If $B \subset A$ is open in A, it has the form $A \cap O$ for some open set O in X. This intersection is open in X if A is open in X. Conversely, if $B \subset A$ is open in X then $A \cap B = B$ is open in A. (Note that the converse does not use the assumption that A is open in X.) The argument for closed sets is just the same. □

Closures behave nicely with respect to subspaces:

Lemma. *Given a space X, a subspace Y, and a subset $A \subset Y$, then the closure of A in the space Y is the intersection of the closure of A in X with Y.*

This amounts to saying that a point $y \in Y$ is a limit point of A in Y (i.e. using the subspace topology on Y) if and only if y is a limit point of A in X.

Proof. For a point $y \in Y$ to be a limit point of A in X means that every open set O in X that contains y meets A. Since $A \subset Y$, this is equivalent to $O \cap Y$ meeting A, or in other words, that every open set in Y containing y meets A. $\square$

The analogous statement for interiors is not true. For example, if A is a line segment in the x-axis in $\mathbb{R}^2$, then the interior of A in the x-axis is an open interval, but the interior of A in $\mathbb{R}^2$ is empty.

Continuity and Homeomorphisms

Recall the definition: A function $f : X \to Y$ between topological spaces is continuous if $f^{-1}(O)$ is open in X for each open set O in Y. For brevity, continuous functions are sometimes called *maps* or *mappings*. (A map in the everyday sense of the word is in fact a function from the points on the map to the points in whatever region is being represented by the map.)

Lemma. *A function $F : X \to Y$ is continuous if and only if $F^{-1}(C)$ is closed in X for each closed set C in Y.*

Proof. An evident set-theory fact is that $f^{-1}(Y - A) = X - f^{-1}(A)$ for each subset A of Y. Suppose now that f is continuous. Then for any closed set $C \subset Y$, we have $Y - C$ open, hence the inverse image $f^{-1}(Y - C) = X - f^{-1}(C)$ is open in X, so its complement $f^{-1}(C)$ is closed. Conversely, if the inverse image of every closed set is closed, then for O open in Y the complement $Y - O$ is closed so $f^{-1}(Y - O) = X - f^{-1}(O)$ is closed and thus $f^{-1}(O)$ is open, so f is continuous. $\square$

Here is another useful fact:

Lemma. *Given a function $f : X \to Y$ and a basis $\mathcal{B}$ for Y, then f is continuous if and only if $f^{-1}(B)$ is open in X for each $B \in \mathcal{B}$.*

Proof. One direction is obvious since the sets in $\mathcal{B}$ are open. In the other direction, suppose $f^{-1}(B)$ is open for each $B \in \mathcal{B}$. Then any open set O in Y is a union $\bigcup_\alpha B_\alpha$

of basis sets B_α, hence $f^{-1}(O) = f^{-1}(\bigcup_\alpha B_\alpha) = \bigcup_\alpha f^{-1}(B_\alpha)$ is open in X, being a union of the open sets $f^{-1}(B_\alpha)$. $\square$

Lemma. *If $f : X \to Y$ and $g : Y \to Z$ are continuous, then their composition $gf : X \to Z$ is also continuous.*

Proof. This uses the easy set-theory fact that $(gf)^{-1}(A) = f^{-1}(g^{-1}(A))$ for any $A \subset Z$. Thus if f and g are continuous and A is open in Z then $g^{-1}(A)$ is open in Y so $f^{-1}(g^{-1}(A))$ is open in X. This means gf is continuous. $\square$

Lemma. *If $f : X \to Y$ is continuous and A is a subspace of X, then the restriction $f|_A$ of f to A is continuous as a function $A \to Y$.*

Proof. For an open set $O \subset Y$ we have $(f|_A)^{-1}(O) = f^{-1}(O) \cap A$, which is an open set in A since $f^{-1}(O)$ is open in X. $\square$

Definition. A continuous map $f : X \to Y$ is a *homeomorphism* if it is one-to-one and onto, and its inverse function $f^{-1} : Y \to X$ is also continuous.

[To be added: some examples of homeomorphisms, e.g., an open interval (a, b) is homeomorphic to $\mathbb{R}$, an open ball in $\mathbb{R}^n$ is homeomorphic to $\mathbb{R}^n$.]

Product Spaces

Given two sets X and Y, their product is the set $X \times Y = \{\, (x, y) | x \in X \text{ and } y \in Y \,\}$. For example $\mathbb{R}^2 = \mathbb{R} \times \mathbb{R}$, and more generally $\mathbb{R}^m \times \mathbb{R}^n = \mathbb{R}^{m+n}$. If X and Y are topological spaces, we can define a topology on $X \times Y$ by saying that a basis consists of the subsets $U \times V$ as U ranges over open sets in X and V ranges over open sets in Y. The criterion for a collection of subsets to be a basis for a topology is satisfied since $(U_1 \times V_1) \cap (U_1 \times V_2) = (U_1 \cap U_2) \times (V_1 \cap V_2)$. This is called the *product topology* on $X \times Y$. The same topology could also be produced by taking the smaller basis consisting of products $U \times V$ where U ranges over a basis for the topology on X and V ranges over a basis for the topology on Y. This is because $(\bigcup_\alpha U_\alpha) \times (\bigcup_\beta V_\beta) = \bigcup_{\alpha,\beta} (U_\alpha \times V_\beta)$.

For example, a basis for the product topology on $\mathbb{R} \times \mathbb{R}$ consists of the open rectangles $(a_1, b_1) \times (a_2, b_2)$. This is also a basis for the usual topology on $\mathbb{R}^2$, so the product topology coincides with the usual topology.

More generally one can define the product $X_1 \times \cdots \times X_n$ to consist of all ordered n-tuples $(x_1, \cdots, x_n)$ with $x_i \in X_i$ for each i. A basis for the product topology on

$X_1 \times \cdots \times X_n$ consists of all products $U_1 \times \cdots \times U_n$ as each U_i ranges over open sets in X_i, or just over a basis for the topology on X_i. Thus $\mathbb{R}^n$ with its usual topology is also describable as the product of n copies of $\mathbb{R}$, with basis the open 'boxes' $(a_1, b_1) \times \cdots \times (a_n, b_n)$.

Example. If we view points in the unit circle S^1 in $\mathbb{R}^2$ as angles θ, then polar coordinates give a homeomorphism $f: S^1 \times (0,\infty) \to \mathbb{R}^2 - \{0\}$ defined by $f(\theta, r) = (r\cos\theta, r\sin\theta)$. This is one-to-one and onto since each point in $\mathbb{R}^2$ other than the origin has unique polar coordinates (θ, r). To see that f is a homeomorphism, just observe that it takes a basis set $U \times V$, where U is an open interval (θ_0, θ_1) of θ values and V is an open interval (r_0, r_1) of r values, to an open 'polar rectangle' and such rectangles form a basis for the topology on $\mathbb{R}^2 - \{0\}$ as a subspace of $\mathbb{R}^2$. By restricting f to a product $S^1 \times [a,b]$ for $0 < a < b$ we obtain a homeomorphism from this product to a closed annulus in $\mathbb{R}^2$, the region between two concentric circles.

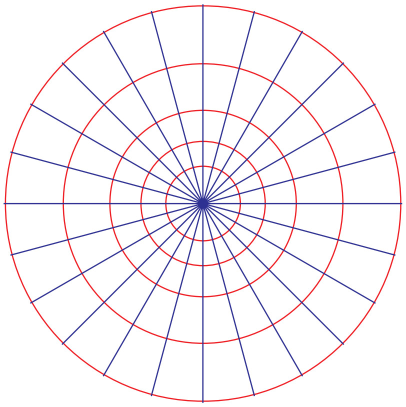

More generally, $\mathbb{R}^n - \{0\}$ is homeomorphic to $S^{n-1} \times (0,\infty)$ where S^{n-1} is the unit sphere in $\mathbb{R}^n$. Using vector notation, a homeomorphism $f: S^{n-1} \times (0,\infty) \to \mathbb{R}^n - \{0\}$ is given by $f(v, r) = rv$, with inverse $f^{-1}(v) = (v/|v|, |v|)$. The continuity of f and f^{-1} can be deduced from explicit algebraic formulas for them, as we will see later in this section.

Example. A product $S^1 \times [a,b]$ is homeomorphic to a cylinder as well as to an annulus. If we use cylindrical coordinates (r, θ, z) in $\mathbb{R}^3$ then a cylinder is specified by taking r to be a constant r_0, letting θ range over the circle S^1, and restricting z to an interval $[a,b]$.

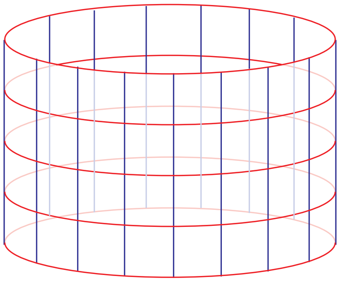

Example. The product $S^1 \times S^1$ is homeomorphic to a torus, say the torus T in $\mathbb{R}^3$ obtained by taking a circle C in the yz-plane disjoint from the z-axis and rotating this circle about the z-axis. We can parametrize points on T by a pair of angles (θ_1, θ_2) where θ_1 is the angle through which the yz-plane has been rotated and θ_2 is the angle between the horizontal radial vector of C pointing away from the z-axis and the radial vector to a given point of C. One can think of θ_1 and θ_2 as longitude and latitude on T. A basic open set $U \times V$ in $S^1 \times S^1$ is a product of two open arcs, and this corresponds to an open curvilinear rectangle on T. Such rectangles form a basis for the topology on T as a subspace of $\mathbb{R}^3$, so it follows that T is homeomorphic to $S^1 \times S^1$.

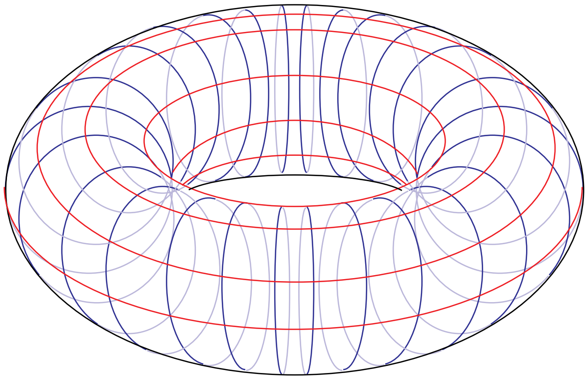

A product space $X \times Y$ has two projection maps $p_1 : X \times Y \to Y$ and $p_2 : X \times Y \to Y$ defined by $p_1(x, y) = x$ and $p_2(x, y) = y$. These maps are continuous since if $U \subset X$ is open then so is $p_1^{-1}(U) = U \times Y$, and if $V \subset Y$ is open then so is $p_2^{-1}(V) = X \times V$.

For each $y \in Y$ there is an inclusion map $i_y : X \to X \times Y$ given by $i_y(x) = (x, y)$. This is continuous because $i_y^{-1}(U \times V)$ is either U if $y \in V$ or $\emptyset$ if $y \notin V$. The map i_y is a homeomorphism onto its image $X \times \{y\}$ since it has a continuous inverse, the restriction of the projection p_1 to $X \times \{y\}$. One can think of $X \times Y$ as the union of the family of subspaces $X \times \{y\}$, each homeomorphic to X, with one such subspace for each $y \in Y$. The situation is of course symmetric with respect to interchanging X and Y, so for each $x \in X$ there is a continuous inclusion map $i_x : Y \to X \times Y$ which is a homeomorphism onto its image $\{x\} \times Y$, and $X \times Y$ is the union of these copies of the space X, one for each point of Y.

A function $f : Z \to X \times Y$ has the form $f(z) = (f_1(z), f_2(z))$. A basic property of the product topology on $X \times Y$ is:

Proposition. *A function $f : Z \to X \times Y$ is continuous if and only if its component functions $f_1 : Z \to X$ and $f_2 : Z \to Y$ are both continuous.*

Proof. We have $f_1 = p_1 f$ and $f_2 = p_2 f$ so f_1 and f_2 are continuous if f is continuous. For the converse, note that $f^{-1}(U \times V) = f_1^{-1}(U) \cap f_2^{-1}(V)$, so this will be open if U and V are open and f_1 and f_2 are continuous. $\square$

As an application, we can give a topological proof that a function $\mathbb{R}^n \to \mathbb{R}$ given by a polynomial in n variables is continuous. The first step is the following fact:

- If two functions $f, g : \mathbb{R}^n \to \mathbb{R}$ are continuous, then so also are the sum function $f+g$ and the product function $f \cdot g$. Namely, we can view $f+g$ as the composition $\mathbb{R}^n \to \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ where the first map is $x \mapsto (f(x), g(x))$ and the second map is $(x, y) \mapsto x + y$. We know the first map is continuous if f and g are continuous, and it is easy to check that the second map is continuous by seeing directly that the inverse image of an open interval is open. For $f \cdot g$ the argument is similar, replacing the second map by the product map $(x, y) \mapsto xy$.

A general polynomial in n variables is built up using repeated addition and multiplication from constant functions, which are certainly continuous, and the coordinate functions x_i, which are nothing but the projections of $\mathbb{R}^n$ onto its n $\mathbb{R}$ factors so are continuous as well.