APPENDIX: TOPOLOGICAL SPACES

1. Metric spaces **224**

Metric spaces

2. Topological spaces **227**

3. Some basic notions for topological spaces **229**

Homeomorphism
Open maps
Bases
Hausdorff spaces

4. Interiors, closures, accumulation points and limits **231**

5. Constructing new spaces from old **233**

6. Quotient spaces **235**

7. Compactness **238**

8. Connectedness and path-connectedness **240**

9. Mapping spaces **244**

A remark on notation

10. Some random additional exercises in point-set topology **246**

Printed 16/10/2014

1. Metric spaces

The first sections are a brief guide to the concepts of topological spaces, continuous functions, and the other basic aspects of point-set topology which we will need during the course.

Point-set topology is not very interesting to teach; it's a language with which to work, rather than an end in itself. In addition, most of the proofs of the theorems "do themselves": there's only really one way to start, in most cases, and it's just a matter of joining the dots, or more precisely, of linking the relevant definitions. These proofs tend to look complicated when written down, because they involve lots of small steps and lots of notation, rather than a single idea which can be expressed in an English sentence. It is therefore usually easier to construct them oneself than to read them from a book. What I'm getting at is: *I'm going to write down very few proofs in this section. Instead, most things are left as exercises, including many standard results which are worth knowing in their own right.* The easiest way to learn them is by doing the exercises.

The primary object of study in algebraic topology is the *topological space*. It is the most general kind of space in which one can do sensible analysis, by which I mean that the notions of continuity, limit etc. make sense. Let's begin working towards the definition by reciting the time-honoured definition of continuity for a real-valued function of a real variable:

Definition. A function $f : \mathbb{R} \to \mathbb{R}$ is *continuous at* $a \in \mathbb{R}$ if for every $\epsilon > 0$, there exists a $\delta > 0$ such that $|x - a| < \delta$ implies that $|f(x) - f(a)| < \epsilon$.

The intuition is, of course, that the function does not jump about locally: if one looks at a sufficiently small range of values about a, then the values of the function may be confined to be arbitrarily close to $f(a)$. Of course, a function is said without qualification to be *continuous* if it is continuous everywhere, i.e. for at all $a \in \mathbb{R}$.

Metric spaces

Suppose now that we want to generalise to more complicated kinds of function, such as (let's not get carried away) a real-valued function of two real variables. Obviously the correct thing to do is repeat the same definition with the Euclidean distance $\|x - a\|$ replacing $|x - a|$ now that x, a are points of $\mathbb{R}^2$.

In fact the same principle will work to give a sensible definition of continuity of any function between subsets of a Euclidean space $\mathbb{R}^n$; all that is needed is the notion of distance between pairs of points. In this way, one can quite happily start talking about continuous functions between spheres of arbitrary dimensions, because the n-sphere is usually thought of as simply the unit sphere inside the Euclidean space $\mathbb{R}^{n+1}$.

If we want to escape the confines of Euclidean space, it is necessary to abstract away the really essential aspects of Euclidean distance. It turns out that the most important thing is the *triangle inequality*: if you start writing out proofs of the simplest properties of continuous one-variable functions, you will need it pretty quickly.

Let's quickly recall the triangle inequality for $\mathbb{R}^n$ and its proof. We want to prove that for any three vectors x, y, z,

$$\|x - y\| \leq \|x - z\| + \|z - y\|.$$

In other words we need to prove that for every pair of vectors a, b,

$$\|a + b\| \leq \|a\| + \|b\|.$$

Squaring this equation and writing it out in terms of coordinates, it becomes

$$\sum a_i b_i \leq \sqrt{\sum a_i^2}\sqrt{\sum b_i^2}$$

which is just the Cauchy-Schwarz inequality. The proof of this is easy: the quadratic function of λ given by $\sum(a_i - \lambda b_i)^2$ is non-negative, so its discriminant "$b^2 - 4ac$" must be non-positive.

Definition. A *metric space* is a set X equipped with a *metric* function $d : X \times X \to \mathbb{R}_{\geq 0}$ such that

(1). $d(x, y) = 0$ if and only if $x = y$

(2). $d(x, y) = d(y, x)$ for any x, y (symmetry)

(3). $d(x, y) \leq d(x, z) + d(z, x)$, for any x, y, z (triangle inequality).

With this notion, we can make an obvious definition of continuity:

Definition. A function $f : X \to Y$ between metric spaces is *continuous at a* if for every $\epsilon > 0$, there exists a $\delta > 0$ such that $d_X(x, a) < \delta$ implies that $d_Y(f(x), f(a)) < \epsilon$.

Example. Here are some very simple examples of metric spaces and continuous functions.

(1). The basic example is obviously $\mathbb{R}^n$ with the Euclidean distance function, as described above.

(2). The *product* of any two metric spaces becomes a metric space, using the sum of the two metrics:

$$d_{X\times Y}((x_1, y_1), (x_2, y_2)) = d_X(x_1, x_2) + d_Y(y_1, y_2).$$

The "product" $f_1 \times f_2 : X_1 \times X_2 \to Y_1 \times Y_2$ of two continuous functions $f_1 : X_1 \to Y_1$ and $f_2 : X_2 \to Y_2$ is continuous.

(3). For any metric space X, the *identity map* $\mathrm{id}_X : X \to X$ and the *diagonal map* $X \to X \times X$ (given by $x \mapsto x$ and $x \mapsto (x, x)$ respectively) are continuous. The metric itself, as a function $X \times X \to \mathbb{R}$, is continuous.

(4). If $f : X \to Y$ and $g : Y \to Z$ are continuous maps between metric spaces, then their *composite* $g \circ f : X \to Z$ is also continuous.

(5). The functions $+ : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}^n$ and $\cdot : \mathbb{R} \times \mathbb{R}^n \to \mathbb{R}^n$ defining the sum of two vectors and the product of a vector with a scalar, are continuous. The norm function $\| - \| : \mathbb{R}^n \to \mathbb{R}$ is continuous.

(6). Combining all these facts gives slick ways to prove that explicit functions (which we will sometimes need to write down) are continuous. One simply has to factorise them as composites of "elementary" functions, which one knows to be continous. For example, the radial projection function $\mathbb{R}^n - \{0\} \to \mathbb{R}^n - \{0\}$ given by $x \to x/\|x\|$ is continuous: it can be written as the composite

$$x \mapsto (x, x) \mapsto (x, \|x\|) \mapsto (x, \|x\|^{-1}) \mapsto x/\|x\|$$

where we use the diagonal to duplicate x, then the product of the identity with the norm function, then the product of the identity with the inversion function on $\mathbb{R} - \{0\}$, then the scalar multiplication function.

(7). The metric associated to the usual Euclidean norm $\|x\|_2 = (\sum x_i^2)^{1/2}$ on $\mathbb{R}^n$ is not the only way of measuring distance. Analysts often use the metric associated to the ℓ_p*-norm* (for $p \geq 1$) given by

$$\|x\|_p = (\sum |x_i|^p)^{1/p},$$

in particular the cases $p = 1$ (the norm of a vector is just the sum of the absolute values of the coordinates) and $p = \infty$ case, which denotes the limiting case

$$\|x\|_\infty = \max |x_i|.$$

A simple way to get a feel for these norms is to draw their "unit circles" in the case of $\mathbb{R}^2$: for $\| - \|_1$ one gets a "diamond" (a square with vertices at plus and minus the usual basis vectors) and for $\| - \|_\infty$ one gets the square with vertices $(\pm 1, \pm 1)$.

One might expect that continuity of functions into and out of $\mathbb{R}^n$ would depend on which of these ℓ_p-metrics is used. But this is not the case. Two metrics d, d' on a fixed metric space X are said to be *Lipschitz-equivalent* if there exist constants $K, k \geq 1$ such that for all points x, y,

$$\frac{1}{k} d(x, y) \leq d'(x, y) \leq K d(x, y).$$

The condition means that the metrics distort each others' distances by a bounded amount. If this is the case then the identity map of X, considered as a map between the different metric spaces (X, d) and (X, d') or vice versa, is continuous. Then for example, if $f : (X, d) \to Y$ is a continuous map of metric spaces, so is f thought of as a map starting from (X, d'), because this is the composite of the original f with the continuous identity map $(X, d') \to (X, d)$.

It is easy to see that all the ℓ_p metrics on $\mathbb{R}^n$ are Lipschitz equivalent; geometrically this is just the fact that the unit sphere of any of them can be sandwiched between two (positive radius) unit spheres of any other. (One could work out the best possible constants k, K using this picture, though it isn't necessary to do so. They will depend on n as well as the two choices of p.)

A similar example arises if one tries to decide what is the "natural" metric on the sphere S^n. Probably the first thing one thinks of is to use Euclidean distance between points of $\mathbb{R}^{n+1}$, obtaining the *chordal metric*. (It measures the length of a straight-line chord joining the points inside the sphere.) On reflection, this is quite tasteless: to define it we used geometry external to the sphere. A better choice is to use the great-circle distance on the sphere's surface, measuring "as the crow flies". Fortunately the two metrics are Lipschitz equivalent, so if we are only interested in continuity of functions, it is irrelevant whether we have taste or not.

This last exercise highlights a problem with the use of metric spaces as a foundation of the theory of continuity — it shows that the actual metric itself contains *far more information* than we need when simply thinking about continuity. The notion of a *topological space* will be more economical: it will incorporate only what we actually need.

Another reason for being unhappy with metric spaces is that not all the constructions we hope to perform with spaces work well. We can certainly take subspaces and products of metric spaces and get sensible induced metrics. But the notions of *quotient* is hopeless. Even *disjoint union* of two metric spaces is unpleasant: in $X \amalg Y$, we have perfectly sensible ways of measuring distance between pairs of points of X, and between pairs of points of Y. But what should be the distance between a point of X and a point of Y? Of course ad hoc definitions are available, but there is no canonical, choice-free method.

2. Topological spaces

To work towards the definition of a topological space, it helps to rephrase the metric space definition of continuity, avoiding explicit dependence on the metric (which we are trying to get rid of).

Definition. Given a point x of a metric space X and a real number $\epsilon > 0$, let us define the *ball of radius ϵ at x* as

$$B_\epsilon(x) = \{y \in X : d(y, x) < \epsilon\}.$$

Definition. A set N is called a *neighbourhood* of a point $x \in X$ if it contains some ball $B_\epsilon(x)$ of positive radius about x.

Definition. A set U is set to be *open* if it is a neighbourhood of each of its points. Such a set can then be written in the form (check each direction of containment if this seems puzzling!)

$$U = \bigcup_{x \in U} B_{\epsilon(x)}(x).$$

This is a bit of a silly expression, from one point of view: we are writing a set as a union of small balls about all of its points in a very redundant way. However, the intuition that every open set can be expressed as some huge union of special kinds of standard small open sets is a valuable one.

Lemma. (*Local form*) *A function $f : X \to Y$ is continuous at $a \in X$ if and only if, for each neighbourhood N of $f(a)$, the inverse image $f^{-1}(N)$ is a neighbourhood of a.*

(*Global form*) *A function $f : X \to Y$ is continuous if and only if, for each open set U in Y, the inverse image $f^{-1}(U)$ is open in X.*

Proof. It's straightforward to check necessity and sufficiency in each case. □

This lemma then, removes the explicit dependence on the metric, as we desired — the *open sets* of a metric space provide enough information for us to talk about continuity. The conceptual leap to a topological space is then simply the realisation that we may as well only specify these open sets, rather than a metric. Remarkably, a few simple axioms suffice to make the structure behave (for the most part) in the way we have come to expect.

Definition. A *topological space* is a set X together with a set τ_X (its *topology*) of subsets of X, whose elements are referred to as the *open sets*, satisfying the axioms:

(1). the whole space X, and the empty set $\emptyset$ are open

(2). the union of an *arbitrary* family of open sets is again open

(3). the intersection of *finitely many* open sets is again open.

Definition. A function $f : X \to Y$ between topological spaces is *continuous* if for every open U in Y (i.e. $U \in \tau_Y$), the inverse image $f^{-1}(U)$ is open in X (i.e. $f^{-1}(U) \in \tau_X$).

I have chosen to write the "slogan" versions here, instead of emphasising the set-theoretic notation, which require one to be very careful not to confuse the symbols $\in$ and $\subseteq$.

It is convenient to define a *neighbourhood* of a point x in a topological space X to be any set which contains an open set containing x. (When X is a metric space, this coincides with our original definition.) It is then possible to define continuity of a function locally (that is, at a point) in terms of neighbourhoods, just as we did for metric spaces.

These definitions are somewhat frightening, and not just because all the geometry appears to have gone out of the window. The structures involved (topologies) can be absolutely enormous, and the whole apparatus appears unmanageable. Fortunately, the intuition developed by thinking with metric spaces is surprisingly helpful for understanding topological spaces, and after working through analogues of the basic theorems (and playing with some of the standard counterexamples) they begin to seem quite visualisable. As for the amount of structure being carried around — well, the metric on a metric space carries more information than the topology it defines (see below); it's just that it somehow seems "smaller".

Example. Here are some standard examples of topological spaces.

(1). Any metric space (X, d) can be considered as a topological space, letting τ_X be the set of d-open sets in X. The second and third axioms for a topology require checking, and it's worth doing this explicitly to illustrate why one deals with infinitely many sets and the other with finitely many.

If $U = \bigcup U_i$ is any union of d-open sets and $x \in U$, then x lies in at least one of the U_i's. Because this particular U_i is open, it contains some $B_\epsilon(x)$, which therefore lies inside U. This proves that U is a neighbourhood of x, and therefore (because the argument works for each x) is open.

On the other hand, if $U = \bigcap U_i$ is an intersection of open sets $U_1, U_2, \dots, U_n$ and $x \in U$ then we can find a collection of balls $B_{\epsilon_1}(x) \subseteq U_1, B_{\epsilon_2}(x) \subseteq U_2, \dots, B_{\epsilon_1}(x) \subseteq U_1$. The intersection of these balls, which is $B_\epsilon(x)$ where $\epsilon = \min\{\epsilon_1, \epsilon_2, \dots, \epsilon_n\}$ (a positive number), lies inside each U_i and therefore inside U. So again we see that U is open. Note however that if there were infinitely-many U_i's then their associated ϵ_i's might converge to zero, and the intersection of the balls could be just the set $\{x\}$, which wouldn't have to be open.

Two metrics on a set are said to be *equivalent* if they define the same topologies. (Lipschitz equivalence is a special case.)

(2). Any set has a *discrete topology* in which all subsets are open, and an *indiscrete topology* in which only X and $\emptyset$ are open sets. All functions out of a discrete space, and into an indiscrete one, are continuous.

(3). The collection of subsets $\{\emptyset, \{a\}, \{a, b\}, X\}$ of the set $X = \{a, b, c\}$ is a topology. An n-element finite set has 2^n subsets, and therefore at most 2^{2^n} possible topologies. Finding a formula for the number (or better, the number considered up to automorphisms (permutations) of the set) is, as far as I know, a hard and unsolved (but fairly irrelevant) combinatorial problem.

(4). In algebraic geometry one uses the *Zariski topology*, in which the open sets are the complements of subsets defined by algebraic equations. For example in $\mathbb{C}$, the algebraic subsets are just the zero-sets of polynomials, and therefore just the finite subsets. The open sets of the Zariski topology are just the empty set and the (rather big) sets $\mathbb{C} - \{x_1, x_2, \dots, x_n\}$, where the x_i's are points of $\mathbb{C}$.

3. Some basic notions for topological spaces

Homeomorphism

Definition. Two topological spaces are *homeomorphic* if there exists a pair of mutually-inverse continuous maps between them.

The term *isomorphic* would be just as good: an isomorphism between mathematical structures (topological spaces, groups, vector spaces, ...) should *always* be defined as a pair of mutually-inverse "structure-preserving" maps, where structure-preserving is interpreted appropriately: that is, as "homomorphism" in the case of groups, "linear map" in the case of vector spaces, and "continuous map" in the case of topological spaces. The language of *category theory*, which will be explained later, encapsulates this idea neatly.

Example. (1). The open unit interval $(-1, 1)$ and the real line $\mathbb{R}$ are homeomorphic. Just use the map $x \mapsto x/(1 - x^2)$ and its inverse.

(2). Generalising this, we have that the open unit ball $\operatorname{Int} B^n$ and the space $\mathbb{R}^n$ are homeomorphic. The map $x \mapsto x/(1 - \|x\|^2)$ is perhaps the nicest choice of homeomorphism.

(3). The map $t \mapsto e^{2\pi i t}$ is a continuous bijection between the interval $[0, 1)$ and the circle S^1. However, its inverse is *not* continuous, and therefore the circle is not homeomorphic to an interval – which is just as well, as topology would be boring if it were!

Open maps

A function $f : X \to Y$ is called an *open map* if for each open U in X, $f(U)$ is open in Y; notice that the requirement here is on the "pushforwards" $f(U)$ of open sets U in X, rather than on the "pullbacks" $f^{-1}(U)$ of open sets U in Y, as in the definition of continuity. Open maps aren't especially important, but they are useful in constructing homeomorphisms. We often have a situation as in (3) above, where we can construct a continuous bijection $f : X \to Y$ and want to know whether its inverse is continuous. The key point is that when f is a bijection, f^{-1} is continuous *if and only if f is open* (just write down the definitions: they're the same).

The map in example (3) is *not* an open map because for example the open set $[0, \frac{1}{2})$ does not get sent to an open set. An even simpler example of a continuous but non-open map is $x \mapsto x^2$, mapping $\mathbb{R}$ to itself.

Bases

It's worth noting that the idea of *generators* for a group has an analogue in the world of topological spaces, and this is sometimes a convenient time-saving device in proofs. Take any collection ρ of subsets of a set X whose union is all of X. It won't in general be a topology, but it is easy to construct a "smallest" topology (one with the fewest open sets) containing all those subsets, as follows. If we close ρ under finite intersections (by adjoining all sets which are intersections of finitely-many elements of ρ), we obtain a larger collection σ which obviously satisfies the third axiom for a topological space, and also contains the empty set. If we now close σ under arbitrary unions (by – of course! – adjoining all unions of elements of σ), we get a collection τ which satisfies the second axiom, contains X as well as $\emptyset$, and (check) still satisfies the third axiom; it is a genuine topology.

Any collection of open sets such as σ which, when closed under unions, generates τ, is called a *base* for τ; any collection such as ρ, which requires closure under both finite intersections and arbitrary unions to generate τ, is called a *sub-base*. In a metric space, for example, the collection of all open balls $B_\epsilon(x)$ is a base for the topology. It's easy to see that just the balls of radius $1/n$ (for positive integer n) will do. In $\mathbb{R}^n$, we may actually use balls of radius $1/n$ based at rational points (points whose coordinates are all rational). A space such as $\mathbb{R}^n$ with a countable base is called *second countable*; this property is technically part of the definition of a manifold, which we will see later.

Hausdorff spaces

Definition. A topological space is said to be *Hausdorff* if, for any pair of distinct points x, y, one can find disjoint open sets U, V containing x, y respectively.

This definition is part of a family of "separation axioms" dealing with whether points and/or open sets can always be "insulated" from one another by means of larger open sets. Hausdorffness (Hausdorffitude?) is the only one worth bothering with here (at all?), for the following simple reason. Any metric space is automatically Hausdorff: if x, y are distinct then $d(x, y) > 0$, and balls of radius $d(x, y)/3$ at x, y are disjoint, by the triangle inequality. In contrast, the Zariski topology described above has no disjoint (non-empty) open sets at all, so it certainly isn't Hausdorff. Topological spaces, therefore, form a strictly larger class than metric spaces.

4. Interiors, closures, accumulation points and limits

The concept of the *interior* of a subset A of a topological space X is quite a natural one. There are two different formulations of the notion: a local and global one.

Definition. (Local form.) The *interior* $\mathrm{Int}(A)$ of a subset A of X is the subset of all points $a \in A$ such that A is a neighbourhood of a.

Exercise. Show that U is an open set if and only if $U = \mathrm{Int}(U)$.

Exercise. (Global form.) Show that $\mathrm{Int}(A)$ is the union of all open subsets of X contained in A.

A *closed set* in a topological space X is one whose complement is open. This is not a terribly interesting definition, but it does suggest, correctly, that one can could reformulate all statements about topological spaces in terms of closed sets rather than open ones.

Exercise. Give examples of subsets of $\mathbb{R}$ which are open but not closed; closed but not open; both; neither.

Exercise. Show that the intersection of arbitrarily many closed sets is closed, and the union of finitely-many closed sets is closed.

Exercise. Show that a function $f : X \to Y$ between topological spaces is continuous if and only if $f^{-1}(F)$ is closed in X, whenever F is closed in Y.

Exercise. A *closed map* is something which pushes forwards closed sets to closed sets. Show that a continuous bijection which is a closed map is a homeomorphism.

A more interesting characterisation of closed sets comes from considering what the *closure* of a set should be. As with the interior, there are two versions of the definition.

Definition. (Local form.) An *accumulation point* of a set A in a topological space X is any point $x \in X$, each of whose punctured neighbourhoods (things of the form $N - \{x\}$, where N is a neighbourhood of x) contains a point of A. The *closure* $\bar{A}$ of A is the union of A and its set of accumulation points; thus, it is the set of points $x \in X$, each of whose neighbourhoods contains a point of A.

Exercise. What are the accumulation points of the following subsets of $\mathbb{R}$: $\mathbb{Z}$, $\mathbb{Q}$, I, $(0, 1)$?

Exercise. Show that F is a closed set if and only if $F = \bar{F}$.

Exercise. (Global form of definition of closure.) Show that $\bar{A}$ is equal to the intersection of all the closed subsets of X which contain A.

Definition. The *boundary* ∂A of a subset A of X is defined to be its closure minus its interior.

Exercise. Consider the operations of closure, interior, and complement as functions C, I, N which take subsets of $\mathbb{R}$ to subsets of $\mathbb{R}$. Show, by exhibiting a subset A of $\mathbb{R}$ for which $CI(A) \neq IC(A)$, that the operations C and I are not commutative. Prove that $C^2 = C, I^2 = I, N^2 = 1$. Are there any other relations in the semigroup generated by C, I, N? See if you can prove that it is finite, and find its order!

A particularly useful idea in topology is the idea of an *open dense subset* U of a space X: this is an open set whose closure is all of X. This concept typically appears whenever one has some kind of space parametrising geometric configurations, and is intimately tied to the ideas of *genericity* and *stability* of such configurations. As an example, let X be the set of all oriented lines in $\mathbb{R}^2$. It is a 2-dimensional space, in fact homeomorphic to $S^1 \times \mathbb{R}$, because lines can be parametrised by their direction and (signed) distance to the origin. It's easy to check that the set of lines which are not horizontal (parallel to the x-axis) forms an open dense subset: *dense*, because if we have a line parallel to the x-axis, then we can perturb it by an arbitrarily small amount to make it non-parallel; and *open* because if we limit ourselves to small enough perturbations, then non-horizontal lines remain non-horizontal. We could say that horizontality is an *unstable* property, but non-horizontality is *stable*. (Non-horizontal lines are also *generic* in the space of all lines — meaning that a randomly chosen line will be non-horizontal with probability 1 — though here we are using measure-theoretic language which is unwarranted without further work.)

In analysis, the idea of the limit of a sequence of points is just as important as the idea of continuity. It is straightforward to give a definition that works for topological spaces.

Definition. A point x of a topological space X is a *limit point* of the sequence of points $(x_n)_{n\in\mathbb{N}}$ if for any neighbourhood N of x, there exists some integer m such that $x_n \in N$ for all $n \geq m$.

This agrees with the usual definition in a metric space: for example, the limit of the sequence $(x_n = 1/n)$ in $\mathbb{R}$ is 0, while the sequence $(x_n = n)$ has no limit. In any Hausdorff space, the limit of a sequence is unique if it exists. In a non-Hausdorff space this need not be the case, and the intuitive picture breaks down. For example, let X be the "real line with a double zero" made by adding a new point $0'$ to $\mathbb{R}$, and adding open sets which are copies of the existing ones containing 0, but with $0'$ inserted in its place. Then the same sequence $(x_n = 1/n)$ has two limit points!

Although the idea of limit of a sequence appears rather similar to the idea of accumulation point of a set, they are actually rather different. If you try to prove the theorem that the accumulation points of a set are given by the limit points of sequences of elements of that set, you will find that it can't be done (and isn't true) without additional information about the topological space: you need the existence of a countable base of neighbourhoods for each point (something which is true, for example, for a metric space). However this does mean that the topology on a metric space X can be completely determined by saying *which sequences* (x_n) *converge* in X: a subset A will then be closed if and only if every convergent-in-X sequence (a_n) whose points actually lie in A, has a limit which also lies in A. It's quite common for analysts to specify a topology on a set by giving its set of convergent sequences, rather than its topology.

5. Constructing new spaces from old

There are four standard constructions in set theory which can be enhanced to constructions with topological spaces.

(1). If $A \subseteq X$ is a subset of a topological space X, it can be given the *subspace topology*

(2). If X, Y are topological spaces then the set $X \times Y$ can be given the *product topology*

(3.) If $\sim$ is an equivalence relation on a topological space X, then the set of equivalence classes $X/\sim$ may be given the *quotient* (or *identification space*) *topology.* An alternative way to describe this situation is simply to consider that we are given a surjective map of sets $q : X \twoheadrightarrow Y$ and a topology on X; we then produce one on Y. (Any equivalence relation defines a surjection $q : X \twoheadrightarrow X/\sim$, and conversely, any surjection q defines an equivalence relation whose equivalence classes are inverse images of points.)

(4). If X, Y are topological spaces then their *disjoint union* $X \amalg Y$ may be made into a topological space. (This might be termed "coproduct" if one were thinking category-theoretically.)

The precise definitions of these new topologies are easy to understand in terms of the following simple principle. In each of the above constructions, there are certain basic "structural maps" relating the old and new sets, and the new topologies should be chosen so as to make these continuous. Here are the structure maps in the four cases:

(1). An injective inclusion map $i : A \hookrightarrow X$.

(2). Two surjective projection maps $\pi_X : X \times Y \to X, \pi_Y : X \times Y \to Y$.

(3). A surjective quotient map $X \to Y$.

(4). Two injective inclusion maps $i_X : X \hookrightarrow X \amalg Y, i_Y : Y \hookrightarrow X \amalg Y$.

Consider case (1). If we are to make i continuous then we certainly require that all sets of the form $i^{-1}(U)$, where U is open in X, are included in the topology on A. In fact these sets form a topology on A (easy check) and so we are done! Notice that we could have used a *larger* topology on A (the cheat's answer would be to give A the discrete topology, making *all* maps out of it continuous!) but that what we have here is the *minimal* topology on A which makes i continuous.

A similar principle does case (2). We are required to put all sets of the form $U \times Y$ and $X \times V$ (where U, V are open sets of X, Y) into the topology on $X \times Y$. For it to be closed under intersections, we must then add all sets of the form $U \times V$'s. Then to be closed under union, we must add all unions of such "box-shaped" sets. The result of this is the product topology; again, it is the minimal possible choice. Notice that sets of the form $U \times V$ form a base for the product topology.

The third and four cases are duals of the first and second, in the sense that the directions of the structure maps are simply reversed, and injections become surjections. (This sort of duality pervades topology and category theory.) In these cases we aren't being *forced* to add sets to the topology to achieve continuity; rather, since the structure maps go *into* our new spaces, we are being *limited* in how much we can add; there is an *upper bound* to the choice of the topology, rather than a lower bound. We can certainly "cheat" again by using the indiscrete topology on the new spaces, but the sensible thing to do is put in the *maximal* topology which will make the structure maps continuous.

In the case of the quotient space we therefore define U to be open in Y if and only if $q^{-1}(U)$ is open in X. For the disjoint union, the open sets are all sets of the form $U \amalg V$, for open sets U, V of X, Y.

Some of the properties of subspaces, products and disjoint unions are covered by the following exercises. (Quotient spaces are sufficiently important to get the next section all to themselves!)

Exercise. 1. Let X be a topological space. Let Y be a subset of X, equipped with the subspace topology, and let A be a subset of Y. Show that if A is closed in Y and Y is closed in X, then A is closed in X. Show that this statement is still true if both "closeds" are replaced by "opens". Give counterexamples for the two "mixed" cases.

Exercise. If A is a subspace of X, and $f : X \to Y$ is continuous, then the restriction of f to A is continuous.

Exercise. Let $i : A \hookrightarrow X$ be the inclusion of a subspace A in a topological space X. Let Z be another space, and $f : Z \to A$ a function. Show that f is continuous if and only if $i \circ f$ is continuous.

Exercise. Suppose that a topological space X is written as the union of finitely-many closed sets F_i, and that we are given functions $f_i : F_i \to Y$ (to some other space Y) which agree on the overlaps $F_i \cap F_j$. Prove that the function $f : X \to Y$ defined piecewise by the f_i's is a continuous function if and only of the individual f_i's are continuous. (This "gluing lemma" is very handy when dealing with piecewise-defined functions.) What is the corresponding statement when we decompose X into open sets U_i?

Exercise. The product of topological spaces $\mathbb{R} \times \mathbb{R}$ is homeomorphic to the space $\mathbb{R}^2$ (with topology coming from the Euclidean metric).

Exercise. The diagonal map $X \to X \times X$ is always a continuous function, for any space X.

Exercise. A function $f : Z \to X \times Y$ is continuous if and only if its coordinate functions $\pi_X \circ f, \pi_Y \circ f$ are both continuous.

Exercise. Suppose X, Y are topological spaces and $A \subseteq X, B \subseteq Y$ are subsets. Show that the two ways of topologising $A \times B$ (as a subspace of a product, or as a product of subspaces) are homeomorphic.

Exercise. Suppose X, Y are topological spaces and $A \subseteq X, B \subseteq Y$ are subsets. Show that the following Leibniz rule holds:

$$\partial(A \times B) = (\partial A \times \bar{B}) \cup (\bar{A} \times \partial B).$$

Exercise. Suppose $\{X_i\}_{i \in I}$ is an infinite family of topological spaces. What is the correct definition of the product topology on the product of sets $\prod X_i$?

Exercise. A function $X \amalg Y \to Z$ is continuous if and only if its restrictions to X and Y are continuous.

6. Quotient spaces

Quotient spaces are very useful in topology. Recall the definition given above: if X is a space and $\sim$ is an equivalence relation on X, then the set $Y = X/{\sim}$ of equivalence classes becomes a space: if $q : X \to X/{\sim}$ is the natural map taking a point x to its equivalence class $[x]$, then the open sets are those U such that $q^{-1}(U)$ is open in X.

Let's work out an example. Consider the equivalence relation on $\mathbb{R}$ given by $x \sim y$ if $x - y$ is an integer. The equivalence classes are *lattices* of the form $[x] = x + \mathbb{Z}$; they are subsets which are translates of the subset $\mathbb{Z}$ of integers.

The quotient space is meant to parametrise such lattices, in the sense that each point of the quotient "is" a lattice, and two points are "close" in the quotient topology if their lattices are "close" inside $\mathbb{R}$. If one starts with the lattice $\mathbb{Z}$ and pushes it gradually to the right, it returns to its initial position (as a subset, not pointwise) after moving a distance of one unit, having in the process taken on every possible position of a lattice in $\mathbb{R}$. This makes it intuitively clear that the quotient space is a *circle*; let us prove this rigorously, as an example of how to work with quotient spaces.

Let $\mathbb{R}/\mathbb{Z}$ denote the quotient space and $q : \mathbb{R} \to \mathbb{R}/\mathbb{Z}$ be the quotient map $x \mapsto [x] = x + \mathbb{Z}$. Let S^1 be the standard circle of unit complex numbers, equipped with the subspace topology from $\mathbb{C}$. To prove that $\mathbb{R}/\mathbb{Z}$ is homeomorphic to S^1 we need to construct a map $f : \mathbb{R}/\mathbb{Z} \to S^1$ which is a bijection, continuous, and has a continuous inverse.

Simply at the level of *sets* (ignoring continuity), there is a bijective correspondence between the set of functions $f : \mathbb{R}/\mathbb{Z} \to S^1$ and the set of functions $\tilde{f} : \mathbb{R} \to S^1$ having the property that $\tilde{f}(x) = \tilde{f}(y)$ whenever $x \sim y$. The correspondence is given by $\tilde{f} \leftrightarrow f \circ q$, or more explicitly by the formula $f([x]) = \tilde{f}(x)$. (If we're given $\tilde{f}$, this formula can be taken as the *definition* of the corresponding f; we see that the property "$\tilde{f}(x) = \tilde{f}(y)$ whenever $x \sim y$" is precisely what's needed to make $f([x])$ *well-defined*, that is, independent of the choice of element x representing the equivalence class $[x]$. When we put back the topology, there is nothing to worry about: the definition of the quotient topology ensures that under this correspondence, *f is continuous if and only if $\tilde{f}$ is.*

We therefore define $\tilde{f} : \mathbb{R} \to S^1$ to be the map $x \mapsto e^{2\pi i x}$. This satisfies $\tilde{f}(x) = \tilde{f}(y)$ whenever $x - y$ is an integer, and it's obviously continuous. So it induces a well-defined continuous map $f : \mathbb{R}/\mathbb{Z} \to S^1$ by the formula $[x] \mapsto e^{2\pi i x}$.

The map f is surjective because $\tilde{f} = f \circ q$ is. It's easy to check that f is also an injection: if $[x]$ and $[y]$ are two points of $\mathbb{R}/\mathbb{Z}$ such that $f([x]) = f([y])$, then $e^{2\pi i x} = e^{2\pi i y}$, x and y must differ by an integer, and we see that actually $[x] = [y]$.

All that remains is to check that f has a continuous inverse, which we do by showing that it's an open map. (This neat trick avoids having to actually write down the inverse of f, which could be messy and confusing). If U is open in $\mathbb{R}/\mathbb{Z}$ then $f(U)$ can also be written as $f \circ q(q^{-1}(U)) = \tilde{f}(q^{-1}(U))$. Since $q^{-1}(U)$ is open in $\mathbb{R}$, by definition of the quotient topology, and $\tilde{f}$ is an open map (by inspection), $f(U)$ is open, and we're done.

(An alternative trick that's sometimes useful at the end of proofs like this is the "compact space to Hausdorff space" property, number (9) in an exercise in the next section.)

Here is a more sophisticated view of the above example. We say that a group G *acts* on a set X if we are given a homomorphism $\rho : G \to \mathrm{Aut}(X)$. Thus, we associate to each group element $g \in G$ an invertible function $\rho(g) : X \to X$ such that $\rho(gh) = \rho(g) \circ \rho(h)$ and $\rho(1) = \mathrm{id}_X$. Instead of writing $\rho(g)(x)$ for the point to which g's function carries x, we usually just call it gx. In the above example, the group $\mathbb{Z}$ is acting on $\mathbb{R}$; the element n acts by "translation of $\mathbb{R}$ by n".

(In topology we usually want to talk about *continuous* actions of *topological* groups on *topological spaces*, but this is irrelevant for now.)

When a group G acts on a set X, there is a natural equivalence relation on X: define $x \sim y$ if there is an element g such that $gx = y$. The equivalence classes are called *orbits*, and the set of orbits $X/\sim$ is usually written X/G. If X is a topological space (not just a set) then X/G becomes the *space* of orbits via the quotient topology. (This explains the reason for the notation $\mathbb{R}/\mathbb{Z}$ above.)

As a further example, the group $\mathbb{Z}^2$ acts on $\mathbb{R}^2$ by integer translations, and the quotient is the 2-torus $S^1 \times S^1$. (The proof is as in example (1), we would take $\tilde{f} : \mathbb{R}^2 \to S^1 \times S^1$ given by $(x, y) \mapsto (e^{2\pi ix}, e^{2\pi iy})$.) Similarly $\mathbb{Z}^n$ acts on $\mathbb{R}^n$ via translations, and we get the n-torus $T^n = (S^1)^n$.

Quotient spaces arising from group actions are perhaps the nicest "naturally occurring" examples. But the most common kind of quotient in topology is somehow less sophisticated: we simply want to glue together or identify various existing spaces in some way to make a new one, and we define an equivalence relation to achieve this.

For example, consider the unit square $X = I \times I$ with the top edge glued to the bottom and the left edge glued to the right: define $(x, 0) \sim (x, 1)$ for each x and $(0, y) \sim (1, y)$ for each y. The equivalence class of a point in $I \times I$ then contains one, two or four elements according to whether the point is in the interior of the square, interior of an edge, or is one of the corners. You can imagine $X/\sim$ as being like the square, except that you can "go off one side and come back on the opposite side" as in the videogame "Asteroids". Alternatively you can imagine actually pasting the edges of a (stretchy, rubbery) square together: after gluing one pair we'd have a cylinder, and after gluing the remaining pair a torus. We conclude that $X/\sim$ (like the universe in "Asteroids") is a torus.

To actually prove that $X/\sim \cong S^1 \times S^1$ we follow the same method as before. Define a map $\tilde{f} : I \times I \to S^1 \times S^1$ by $(x, y) \mapsto (e^{2\pi ix}, e^{2\pi iy})$; check that it respects the equivalence relation in the right way to induce a map $f : X/\sim \to S^1 \times S^1$; check that this is a bijection, and then show its inverse is continuous.

Here are some further examples of quotient spaces. It would take a lot of effort to describe each of these in the detail it deserves (and also lots pictures, which I am too lazy to do right now) so I will give up and just give you the idea.

(1). Take $I \times I$ and identify $(0, y)$ with $(1, 1 - y)$ for each y. This gives the *Möbius strip.*

(2). Take $I \times I$ and identify $(0, y) \sim (1, 1 - y)$ for each y and also $(x, 0) \sim (x, 1)$ for each x; this gives the *Klein bottle.*

(3). Take S^2 and identify antipodal points. (Or, take the quotient of S^2 by the group Z_2, whose non-trivial element acts via the antipodal map $x \mapsto -x$). The result is the *projective plane* $\mathbb{R}P^2$.

(4). Take a regular octagon and identify its opposite edges in pairs (make them correspond via a translation) just as we did for a torus. The result is a *closed orientable surface of genus* 2, which looks like a torus with two "holes". In fact all connected closed 2-manifolds can be obtained by gluing pairs of edges of polygons in this way.

(5). There is a huge class of spaces which can be built by merely gluing together collections of balls B^n (of varying dimension). These spaces are called *CW complexes* and their behaviour is so nice that they are the main class of spaces with which people actually work in algebraic topology.

(Arbitrary topological spaces can be very "pathological", and one often needs to make additional assumptions in order to make much progress understanding their topology.)

(7). *Real projective n-space* $\mathbb{R}P^n$ is the space of lines through the origin in $\mathbb{R}^{n+1}$. To define it we consider the equivalence relation on $\mathbb{R}^{n+1} - \{0\}$ given by $x \sim y$ if x is a non-zero scalar multiple of y. The equivalence classes are lines through the origin with their zero points missing (if we didn't remove zero, everything would be equivalent to everything else) and so the quotient space is a space whose points correspond to lines through the origin, as required. We could express the same space more nicely as the quotient of $\mathbb{R}^{n+1} - \{0\}$ by the action of the multiplicative group $\mathbb{R}^* = \mathbb{R} - \{0\}$ of dilations.

A slight variation on this gives a different construction of the sphere S^n. If we take $\mathbb{R}^{n+1} - \{0\}$ modulo the action of the multiplicative group $\mathbb{R}_{>0}$ of *positive* dilations, the orbits are half-lines (or *rays*) emanating from (though once more not including) the origin, so the quotient space is the "celestial sphere" (of directions from which light rays can approach the observer at zero) of *directions* in $\mathbb{R}^{n+1}$.

Another variation is to repeat the process with the *complex* vector space $\mathbb{C}^{n+1}$ and the multiplicative group $\mathbb{C}^*$ to give $\mathbb{C}P^n$, the complex projective space. These spaces are of fundamental importance in algebraic geometry.

Exercise. Prove that the product of two Hausdorff space is Hausdorff, and that a subspace of a Hausdorff space is Hausdorff. Give an example to show that a quotient of a Hausdorff space need not be Hausdorff.

Exercise. The group $\mathbb{Z}$ acts on the space $\mathbb{C}^2 - \{0\}$ as a group of dilations, according to the formula

$$n.(z, w)) = (2^n z, 2^n w).$$

Show that the space of orbits is homeomorphic to the space $S^1 \times S^3$.

7. Compactness

In the study of complex vector spaces, one learns that the correct notion of "finiteness" is actually *finite-dimensionality* rather than finiteness as a set. (Only the 0-dimensional space is actually a finite set!) Similarly, when studying groups, rings or modules (vector spaces being a special case), the notion of being *finitely-generated* is useful, in the sense that such "finite" objects have nice properties not shared by their "infinite"relations.

Compactness is in some sense the topological space analogue of being finitely-generated. We will see that for subspaces of Euclidean spaces $\mathbb{R}^n$, it is exactly the same as being closed and bounded, and that intuition about the properties of such sets can often be carried into the general case.

Definition. An *open cover* of a space X is a family of open sets $\{U_i\}_{i \in I}$ whose union is X. A *subcover* of such a cover is any subcollection $\{U_j\}_{j \in J}$, where $J \subseteq I$, whose union is still all of X. A space is *compact* if every open cover has a finite subcover.

Note immediately that this is a *topological property*: that changing a space by a homeomorphism preserves its compactness (or lack of it).

To say that a *subset* A of a space X is compact means that A, when given the subspace topology (and viewed as a space in its own right), is compact in the above sense. Equivalently, we can define an open cover of A to be a family of open sets of X whose union contains A; then A is compact if each of its open covers has a finite subcover.

Example. The open unit ball in any $\mathbb{R}^n$ is not compact, because the family of open balls of radius $1 - 1/k$ (for integers $k \geq 2$) forms a cover, any of whose finite subcovers has a largest element with radius strictly less than 1.

Theorem (Heine-Borel). *The (closed) unit interval I is compact.*

Proof. This is a standard proof by contradiction. Suppose we have a cover with no finite subcover. By restriction, it gives a cover of each of the two half-intervals $[0, \frac{1}{2}]$ and $[\frac{1}{2}, 1]$. At least one of these covers cannot have a finite subcover, or we could combine the two to get a finite subcover of I. So pick this half-interval (if neither has a finite subcover, just choose one of them), and repeat the argument. We obtain a nested sequence of closed intervals of successively halving length whose left endpoints form a bounded, monotonically increasing sequence, and whose right endpoints form a bounded, monotonically decreasing sequence. Each of these sequences therefore has a limit (it supremum/infimum - that these exist is the defining property of the real numbers) but because of the halving lengths they must agree. We obtain therefore that the intersection of the family of halving intervals is a single point $x \in I$. This x must lie in at least one of the sets of the cover, so choose one: because the set is open, it extends a positive distance to either side of x, and must therefore eventually contain the tiny halving intervals of the sequence, contradicting the fact that they do not have finite subcovers. □

Exercise. Prove the following sequence of statements. (At each stage, you can assume the previous ones.)

(1). The product of two compact spaces is compact. (It follows by induction that any finite product of compact spaces is compact, but in fact an *arbitrary* (infinite) product of compact spaces is compact; this is *Tychonoff's theorem.*)

(2). A closed subspace of a compact space is compact.

(3). A closed, bounded subspace of $\mathbb{R}^n$ is compact.

(4). A quotient space of a compact space is compact.

(5). A compact subspace of a Hausdorff space is closed.

(6). A continuous real-valued function on a compact space is bounded and attains its bounds.

(7). A compact subspace of $\mathbb{R}^n$ is closed and bounded.

(8). The union of finitely many compact sets is compact. (Is this true for finite intersections?)

(9). A continuous map from a compact space to a Hausdorff space is a closed map; and consequently that a continuous bijection from a compact space to a Hausdorff space is a homeomorphism.

Exercise (The *Lebesgue lemma*)**.** Given any open cover $\{U_i\}$ of a compact metric space, there is some constant $\delta > 0$ such that for any set of diameter less than δ, one can find one of the U_i's which contains it. (The *diameter* of a subset of a metric space is the supremum of the set of pairwise distances between its points.)

Exercise. An infinite subset of a compact space must have an accumulation point. Deduce that a discrete subset (one which, as a subspace, inherits the discrete topology) of a compact space must be finite.

8. Connectedness and path-connectedness

There are two sensible notions of connectedness in common use. The more obvious one, *path-connectedness*, measures whether any two points of a space may be joined by a continuous path (a continuous image of the unit interval). Thus it studies the space using continuous maps *into* it. It is very common in topology to find a duality between "maps in" and "maps out" notions, and this is no exception; the notion of *connectedness* measures whether there exist continuous maps *out* of a space onto a discrete space with two points.

Definition. A space is *path-connected* if for each pair of points $x, y \in X$, it is possible to find a *path* joining x and y, that is a continuous map $\gamma : I \to X$ such that $\gamma(0) = x, \gamma(1) = y$. I will write such a path as $\gamma : x \to y$, hoping it will not be cause confusion with the notation for functions.

Example. Any Euclidean space $\mathbb{R}^n$ is path-connected; the formula for the straight-line path γ joining two vectors x, y is the "weighted linear combination"

$$\gamma(t) = (1-t)x + ty \quad 0 \le t \le 1,$$

which will come in handy on many occasions. On the other hand, the space $\mathbb{R} - \{0\}$ is not path-connected. This follows from the intermediate value theorem of basic real analysis: any continuous map $I \to \mathbb{R}$ whose initial value is negative and whose final value is positive must take the value 0; therefore there are no continuous paths connecting any negative real with any positive one.

The property of path-connectedness is, like compactness, a topological property: homeomorphism preserves path-connectedness (or lack of it).

Corollary. *The spaces* $\mathbb{R}$ *and* $\mathbb{R}^2$ *are not homeomorphic.*

Proof. If $f : \mathbb{R} \to \mathbb{R}^2$ were a homeomorphism, then f would restrict to a homeomorphism between $\mathbb{R} - \{0\}$ and $\mathbb{R}^2 - \{f(0)\}$. However, the former is not path-connected, whereas the latter one is (easy check), so no such homeomorphism exists. □

Unfortunately this idea does not generalise to distinguish (meaning, to prove non-homeomorphic) Euclidean spaces of dimension greater than 1 from one another. While it is indeed true that $\mathbb{R}^m \not\cong \mathbb{R}^m$ for distinct m, n (this fact being called "Brouwer's invariance of domain"), we need more subtle algebraic-topological tools to show this.

Let us consider now the other notion of connectedness.

Definition. A space is *disconnected* if it is possible to write it as the union of two non-empty disjoint open sets. Therefore, a space X is *connected* if, whenever X is written as a union $X = U_1 \cup U_2$ of disjoint open sets U_i, one of them must be empty. Equivalently, X is connected if the only subsets which are both open and closed are X and $\emptyset$.

This definition is the standard one, but it can be immediately rephrased in a more illuminating way as follows.

Lemma. *A space* X *is disconnected if and only if there exists a continuous surjection* $X \to 2$, *where* $2 = \{0, 1\}$ *denotes the two-point space with the discrete topology. So a space is connected if and only if there is no such surjection.*

Proof. The correspondence should be clear: a map $X \to 2$ is continuous precisely when the preimages $U_0 = f^{-1}(0), U_1 = f^{-1}(1)$ are open, and is a surjection precisely when they are both non-empty. Thus existence of such a function disconnects a space, and conversely a disconnected space has such a function. □

The most common relations between path-connectedness and connectedness are summarised by the following statements.

Lemma. *A path-connected space X is connected.*

Proof. A continuous function $X \to 2$ must be constant along the image of any continuous path $\gamma : I \to X$, by the intermediate value theorem. So for a path-connected space it takes the same value everywhere, and therefore cannot be surjective. □

Example. A connected space need not be path-connected. An example is the subspace X of $\mathbb{R}^2$ consisting of the graph of the function $\sin(1/x)$, for $x \neq 0$, together with a single point at the origin.

To show it is connected: clearly each side of the graph is a path-connected, hence connected, set. So any continuous function $f : X \to 2$ must be constant for $x > 0$ and for $x < 0$. But since there are sequences of points on each side converging to the origin, and by continuity the value $f(0)$ is the limit of each sequence, the two values are equal, and the function f is constant on X. Therefore X is connected.

To show it is not path-connected: suppose there were a path $\gamma : I \to X$ joining the points with x-coordinates ± 1. Let π be the projection from $\mathbb{R}^2$ onto the x-axis. Then $\pi\gamma(I)$ is a subset of the x-axis containing the interval $[-1, 1]$, because of the intermediate value theorem and the fact that it is a path-connected set. Since π restricted to X is injective, $\gamma(I)$ contains all points of X with x-coordinates in $[-1, 1]$. However, this is not a closed subset of $\mathbb{R}^2$, whereas $\gamma(I)$, which is a continuous image of a compact space, must be. This contradiction finishes the proof.

Exercise. A connected subset of a disconnected space $X = U_1 \cup U_2$ lies either completely inside U_1, or completely inside U_2.

Exercise. The product of connected spaces is connected.

Exercise. A quotient of a connected space is connected.

Exercise. If two connected sets intersect non-trivially, then their union is connected.

Exercise. Which of the above statements is true when path-connected replaces connected?

The relation of being joined by a path is easily seen to be an *equivalence relation*: it is reflexive (the constant path $\gamma(t) = x$ joins x to x!), symmetric (define $\gamma^{-1}(t) = \gamma(1 - t)$ to turn a path $\gamma : x \to y$ into $\gamma^{-1} : y \to x$) and, most importantly, transitive: compose paths $\gamma : x \to y$ and $\delta : y \to z$ using the formula below (when composing paths, it seems madness not to write them from left to right, as I do here):

$$(\gamma.\delta)(t) = \begin{cases} \gamma(2t) & 0 \leq t \leq \frac{1}{2} \\ \delta(2t - 1) & \frac{1}{2} \leq t \leq 1. \end{cases}$$

Definition. The *set of path-components* of a space X, denoted $\pi_0(X)$, is the set of equivalence classes under the above relation. Note that there is a map $X \to \pi_0(X)$ sending each point to its equivalence-class, or *path-component*.

Example. An open subset U of $\mathbb{R}^n$ which is connected is path-connected.

Proof. Suppose y is a point in the path-component of x. Since U is open, we may find a small ball $B_\epsilon(y)$ contained in U. Each point of this ball is joined to y by a linear path, and therefore the whole ball lies in the path-component of x; we've therefore shown that any path-component of U is open. Now observe that the fact that the path-components partition U, together with the fact that U is connected, shows that there can be only one. □

It is possible to define *connected components* (or simply *components*) of a space. Like path-components, they partition the space.

Definition. The *connected component* of a point x of X is the union of all connected subsets of X which contain x.

Exercise. If $\{C_i\}$ is any family of connected sets whose intersection is non-empty, then the union $\bigcup C_i$ is connected. Hence "connected components" are (as the name certainly suggests, though in mathematics this kind of logical reasoning can be disastrous) themselves connected!

Exercise. The closure of a connected set is connected. Hence components are closed.

Example. Let X be the subspace of $\mathbb{R}$ consisting of the numbers $\{1/n\}$ (for all $n \in \mathbb{N}$) together with the point $\{0\}$. Each point $\{1/n\}$ is its own component, and therefore (since components partition) so is $\{0\}$. However, $\{0\}$ is not open: this shows that components, though always closed, do *not* have to be open.

Definition. A space X is *locally connected* if for any $x \in X$ and neighbourhood U of x, there is a neighbourhood $V \subseteq U$ of x which is connected.

Remark. In general, a space is said to be *locally something-or-other* if given any point x and neighbourhood U, we can find inside U a smaller neighbourhood V of x which is something-or-other. The most common properties are local compactness, local connectedness, local path-connectedness, and local contractibility.

Exercise. In a space which is locally connected, components are both open and closed.

Exercise. In a space which is locally path-connected, a subset is connected if and only if it is path-connected.

Although the notion of path-connectedness does not seem very subtle or complicated, it lies at the heart of algebraic topology, and for two distinct reasons. Firstly, we will see later how the subject revolves around the families of topological invariants known as *homotopy groups* and *cohomology groups*. There are various possible ways to define these invariants for a space X, but perhaps the slickest (at least conceptually) is in terms of the (path-)connectedness of certain "auxiliary" spaces built naturally out of X. Homotopy classes of maps $X \to Y$, for example, are just the path-components of the mapping space $\mathrm{Map}(X, Y)$, and the homotopy groups $\pi_n(Y)$ result from taking Y to be a sphere S^n.

Secondly, all these invariants are *functors*: they are "machines" which convert topological information (spaces and continuous maps between them) into algebraic information (usually groups, and homomorphisms between groups). The idea of a functor will be explained properly later, but it is worth illustrating here.

Define $\pi_0(X)$ to be the set of path-components of X. If we view it as the quotient of X by the equivalence relation which identifies all pairs of points joined by paths, it's easy to see that a continuous map $f : X \to Y$ induces a well-defined map $\pi_0(f) : \pi_0(X) \to \pi_0(Y)$. So we can view π_0 as a machine which converts not just topological spaces into sets, but also maps between spaces into functions between those sets.

Moreover, we can see that $\pi_0(\mathrm{id}_X) = \mathrm{id}_{\pi_0(X)}$ for any X, and that if $g : Y \to Z$ is also continuous, then $\pi_0(g \circ f) = \pi_0(g) \circ \pi_0(f)$. If we pretend that f, g are elements of a group in which composition is the multiplication, and that $\pi_0(f), \pi_0(g)$ are elements of a different "group", again with composition as multiplication, then these properties make π_0 look like a *homomorphism*. Such a "homomorphism" is actually called a *functor* from the *category of topological spaces* to the *category of sets*.

Exercise. If $f : X \to Y$ is a homeomorphism then $\pi_0(f)$ is a bijection between $\pi_0(X)$ and $\pi_0(Y)$.

A "dual" approach would be to define $H^0(X)$ to be the set of (not necessarily continuous) functions $X \to \mathbb{Z}$ which are constant along paths in X. The natural operation of pointwise addition of functions turns $H^0(X)$ into a group, and fairly clearly it is just the set of integer-valued functions on the set path-components of X (which is why it is in some sense "dual" to $\pi_0(X)$). For example, if X is path-connected then $H^0(X) \cong \mathbb{Z}$. If we're given a continuous function $f : X \to Y$, we can *precompose* it with functions in H^0 so as to obtain an induced map $H^0(f) : H^0(Y) \to H^0(X)$; notice that this map goes *the wrong way*, because functions $Y \to \mathbb{Z}$ are *pulled back* (pre-composed with f) to functions on X. The map $H^0(f)$ is obviously a homomorphism of groups. Because of the wrong-wayness, H^0 is an example of a *contravariant functor* from spaces to sets, whereas a right-way functor like π_0 is said to be *covariant*.

In the next section we will describe the *fundamental group* π_1 and *higher homotopy groups* π_n of a space. They are in a sense just special cases of π_0. Similarly, we will go on to study *singular cohomology groups* H^n of a space. These generalise the zeroth group H^0 above. All these operations π_n, H^n are functors, satisfying laws like the above. The π_n are *covariant* whereas the H^n are *contravariant*.

During the early development of topology (approximately 20s - 40s) many different kinds of cohomology groups for spaces were introduced. Each one can be viewed as generalising a different type of connectedness property. For example, *Čech cohomology* generalises a slightly different definition of $H^0(X)$ as the abelian group of *locally constant* functions $X \to \mathbb{Z}$. (A locally constant function on X is one such that each point possesses a neighbourhood on which the function is constant; when $\mathbb{Z}$ is the target, these are actually the same as continuous functions $X \to \mathbb{Z}$.) If X is locally connected then this group is the same as the group of $\mathbb{Z}$-valued functions on the set of *components* (rather than path-components) of X, but if X is in addition locally path-connected then it agrees with the singular cohomology group $H^0(X)$ we defined above. The higher Čech cohomology also agree with the singular cohomology groups when the space X is nice enough (for example, when it is a CW-complex) but otherwise measure slightly different topological information. This is typical of the variant cohomology theories; they tend only to differ on slightly pathological spaces, and since in the end singular cohomology is perhaps the simplest and most general, it is the one which ultimately became the standard one.

9. Some random additional exercises in point-set topology

1. Consider $\mathbb{C}^2$ with the standard hermitian form

$$\langle (z_1, w_1), (z_2, w_2) \rangle = \bar{z}_1 w_1 + \bar{z}_2 w_2$$

and associated norm $\|(z, w)\|^2 = |z|^2 + |w|^2$. As a metric space, this is the same as $\mathbb{R}^4$ with the Euclidean norm, but the complex coordinates are often more useful; in particular, the 3-sphere S^3 may be regarded as the unit sphere in $\mathbb{C}^2$. Write down a continuous unit tangent vector field on S^3. Can you find another two so that the three form an orthonormal triple at each point? (It might help to return to real coordinates for this part.)

2. Let $M(n, \mathbb{R})$ and $GL(n, \mathbb{R})$ denote the sets of all $n \times n$ matrices, and all *invertible* $n \times n$ matrices, respectively. Give $M(n, \mathbb{R})$ a topology by identifying it with Euclidean space $\mathbb{R}^{n^2}$, and give $GL(n, \mathbb{R})$ the subspace topology.

(a). Show that matrix multiplication is a continuous map $GL(n, \mathbb{R}) \times GL(n, \mathbb{R}) \to GL(n, \mathbb{R})$, and that the inversion map $A \mapsto A^{-1}$ is a continuous map $GL(n, \mathbb{R}) \to GL(n, \mathbb{R})$. (Hint: consider cofactors). These properties make $GL(n, \mathbb{R})$ into a *topological group*.

(b). Show that the subgroup $O(n) \subseteq GL(n, \mathbb{R})$ consisting of all *orthogonal* matrices (ones satisfying $AA^T = 1$) is compact. (Hint: show it is closed and bounded.)

(c). By considering the determinant function, show that $O(n)$ is disconnected.

(d). Let $SO(n) \subseteq O(n)$ be the *special orthogonal* subgroup consisting of matrices with determinant 1. Show that $SO(n)$ is path-connected. (Hint: consider the columns of the matrix as an orthonormal basis; rotate suitably to obtain the standard basis.)

(e). Show that $O(n)$ has two components.

(f). A *lattice* $\Lambda \subseteq \mathbb{R}^n$ is a subset formed by picking a basis of $\mathbb{R}^n$ and then taking all *integral* linear combinations of the basis vectors. A subgroup $\Gamma \subseteq O(n)$ is said to be *crystallographic* if there exists some lattice Λ which is fixed by Γ (this means that $g\Lambda = \Lambda$ for each element $g \in \Gamma$, but not that the individual elements of Λ are necessarily fixed by g). Show that a crystallographic subgroup must be finite.

(g). Suppose instead we had started with $\mathbb{C}^n$, and defined analogously $M(n, \mathbb{C})$, $GL(n, \mathbb{C})$, the unitary group $U(n)$ (consisting of matrices with $A\bar{A}^T = 1$) and the special unitary group $SU(n)$. What would change?

3. Inside $\mathbb{R}^n$ consider the *hyperplanes* (subspaces of codimension 1) H_{ij}, defined for $i \neq j$ by

$$H_{ij} = \{(x_1, x_2, \ldots, x_n) : x_i = x_j\}.$$

Let Δ be the union of all the H_{ij}, and consider the subspace $C = \mathbb{R}^n - \Delta$.

(a). Show that the set $\{(x_1, x_2, \ldots, x_n) : x_1 < x_2 < \cdots < x_n\}$ is path-connected.

(b). Show that C has $n!$ components.

(c). How does this result change if $\mathbb{R}$ is replaced by $\mathbb{C}$?

4. (*The Cantor set.*) Consider the unit interval I, and define a sequence $M_1, M_2, \ldots$ of open subsets by

$$M_n = \bigcup_{0 \leq p \leq n-1} \left(\frac{3p+1}{3^n}, \frac{3p+2}{3^n} \right).$$

Define the *Cantor set* to be $C = I - \bigcup M_n$; we remove the (open) middle third of I, then the middle thirds of what remains, and so on...

(a). Prove that C is a compact metric space which is *totally disconnected*, meaning that each of its points is its own connected component.

(b). Let X be the product of countably many copies of the discrete set $\{0, 1, 2\}$; we will view its elements as sequences $(a_1, a_2, a_3, \ldots)$. Prove that the map $\theta : X \to I$ which takes the sequence (a_i) to $\sum_{i \geq 1} a_i/3^i$ is continuous and onto but not injective.

(c). Let $S \subseteq X$ be the set of sequences which consist only of zeroes and twos. Show that restricting θ to S gives a homeomorphism $\tau : S \to C$: We see that C is homeomorphic to the countable product of discrete two-point sets.

(d). Check that the Cantor set is *perfect* (what a dumb choice of word!), meaning that every point is an accumulation point of C.

It is a fact that any totally disconnected, perfect compact metric space (for example, do a "middle fifths" construction) is homeomorphic to C.

5. (*Space-filling curves.*) Let C be the Cantor set, viewed as the countable product of the discrete set $\{0, 2\}$, and consider the map $\psi : C \to I$ which takes the sequence (a_i) to $\sum_{i \geq 1} \frac{1}{2} a_i/2^i$.

(a). Check that ψ is a continuous surjection.

(b). Check that the map $\Delta : C \to C \times C$ which sends (a_i) to $((a_1, a_3, a_5, \ldots), (a_2, a_4, a_6, \ldots))$ is continuous and hence that $f = (\psi \times \psi) \circ \Delta : C \to I^2$ is a continuous surjection.

(c). Now consider the Cantor set again as a subset of I, via the middle-thirds picture. Extend $f : C \to I^2$ to a map $I \to I^2$ by extending it linearly over the deleted open intervals. Check that the result is a continuous surjection $I \to I^2$, that is a *space-filling curve.* Moral: the concept of "dimension" is not straightforward!