Lecture notes on topology

Paolo Capriotti

March 2019

Contents

1 Introduction

Topology is the study of those properties of "geometric objects" that are invariant under "continuous transformations".

In these notes, we will make the above informal description precise, by introducing the axiomatic notion of a *topological space*, and the appropriate notion of *continuous function* between such spaces.

Today, topology is used as a base language underlying a great part of modern mathematics, including of course most of geometry, but also analysis and algebra. As a discipline of its own, one could say that topology is mainly concerned with the *classification* of spaces up to *homeomorphism*, which is the notion of "isomorphism" of spaces resulting from the above choice of maps.

Intuitively, we can think of homeomorphisms as procedures that take a space and "deform" it by twisting, bending, stretching or compressing it, but without ever creating "tears" or puncturing "holes". The standard example is turning a torus (fig. 5) into (the surface of) a coffee mug.[1]

The formalism described in these notes will allow us to turn the visual intuition underlying transformations such as the one above into rigorous mathematical proofs. Since it is not too easy to parameterise a coffee mug with a nice simple formula, we will not actually deal with this particular example, but we will see a few similar ones, which should hopefully serve to convince the reader that all these constructions can actually be made precise.

On the other hand, we will see that most of the difficulty in topology is not in constructing homeomorphisms, where our geometric intuition can guide us, but instead in showing that no homeomorphism *can* exist between a certain pair of spaces. This is where the theory developed in these notes will really help. We will define certain "attributes" of topological spaces that can be proved to be preserved by homeomorphisms, and this will allow us to make sure that what we regard as different spaces are indeed different as far as topology is concerned.

Section 2 will be concerned with reviewing the probably familiar notion of *metric space*. The theory of metric spaces can be used a setting for developing topology, but the rigidity of the metric structure makes certain constructions (e.g. quotients) not possible. Furthermore, the notion of distance in a metric space is not invariant under the transformations we are interested in, meaning that metric spaces are in some sense "over-specified", or, in other words, carry "too much" information.

In section 3 we will therefore introduce the notion of *topological space*, which will be the focus of the remainder of these notes. Metric spaces will provide a very important class of examples of topological spaces, but we will see that even in those examples it is often best to ignore the metric structure.

Section 4 will examine notions of convergence in spaces, connecting them to the familiar case of sequences in metric space. We will see that convergence of sequences is often not general enough to capture the topological structure of a space, and we will remedy this using filters and the corresponding notion of convergence.

After that, we will turn our attention to topological properties of spaces, which will allow us to take our first steps towards distinguishing pairs of non-homeomorphic spaces. First is the notion of connectedness in section 6, which formalises the intuitive idea of a space that is made up of a single "piece". Next is section 7, dealing with *separation* and *countability* axioms, which are technical well-behavedness conditions, are useful to transport some of the techniques used in metric spaces to more general topological spaces. Section 8 is instead devoted to *compactness*, which generalises to arbitrary topological spaces the properties of closed and bounded subsets of $\mathbb{R}^n$.

Finally, section 9 will provide a taste of algebraic topology, by developing what is probably the most basic of the algebraic invariants of a space: the *fundamental*

[1] See `https://en.wikipedia.org/wiki/Homeomorphism` for an animation.

group. It will allow us to distinguish more pairs of spaces, as well as to prove the classical fixed point theorem by Brouwer.

Throughout these notes, we will encounter several examples of spaces, some of which may be already familiar. We will use those spaces to refine our understanding of the topological properties that we will study, as well as to demonstrate how the techniques developed here can help us establish whether two spaces are homeomorphic.

2 Review of metric spaces

Definition 2.0.1. A *metric* on a set X is a function $d : X \times X \to \mathbb{R}_+$ satisfying:

(i) $d(x, y) = 0$ if and only if $x = y$;
(ii) $d(x, y) = d(y, x)$ for all $x, y \in X$ (*symmetry*);
(iii) $d(x, z) \leq d(x, y) + d(y, z)$ for all $x, y, z \in X$ (*triangle inequality*).

A *metric space* is a set X equipped with a metric d.

2.1 Examples

Let $d : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}_+$ the familiar *Euclidean distance*, defined by:

$$d(x, y) := \|x - y\|, \quad \text{where} \quad \|x\| := \sqrt{\sum_{i=1}^{n} |x_i|^2}.$$

It is immediate to verify that d satisfies all the properties of Definition 2.0.1. Therefore, it defines a metric on $\mathbb{R}^n$. In the following, we will always regard $\mathbb{R}^n$ as a metric space by implicitly equipping it with the Euclidean metric, and we will refer to it as the *Euclidean space* of dimension n.

Similarly, there is a Euclidean distance on $\mathbb{C}^n$, given by the same formula, where now $|-|$ denotes the absolute value (or magnitude) of a complex number.

A more trivial example is that of a *discrete metric space*. Let X be any set, and define

$$d(x, y) = \begin{cases} 1 & \text{if } x \neq y \\ 0 & \text{if } x = y. \end{cases}$$

All the properties of Definition 2.0.1 are easily verified.

2.2 Balls and open sets

We can directly generalise the notion of a *ball* in Euclidean space to a metric space:

Definition 2.2.1. Let X be a metric space with metric d. Given a point $p \in X$, and a real number $r > 0$, define a subset of X as follows:

$$B_r(p) = \{x \in X \mid d(x, p) < r\}.$$

We will refer to $B_r(p)$ as the (*open*) *ball* of centre p and radius r.

When we want to make the metric explicit, we will write $B_r^d(p)$ for the open ball of centre p and radius r with respect to the metric d.

Definition 2.2.2. A subset U of a metric space X is said to be *open* if for every point $x \in U$ there exists $\varepsilon > 0$ such that $B_\varepsilon(x) \subseteq U$.

Exercise 1. Prove that $U \subseteq X$ is open if and only if it is a union of open balls. Deduce in particular that open balls are open.

Lemma 2.2.3. *Open subsets of a metric space X are stable under finite intersections and arbitrary unions.*

Proof. Let $(U_i)_{i\in I}$ be a family of open sets. If $x \in \bigcup_{i\in I} U_i$, then there exists $i \in I$ such that $x \in U_i$. Consequently, there exists $\varepsilon > 0$ such that $B_\varepsilon(x) \subseteq U_i \subseteq \bigcup_{i\in I} U_i$, so open sets are stable under unions.

As for intersections, let $x \in \bigcap_{i\in I} U_i$. For each i, let ε_i be such that $B_{\varepsilon_i}(x) \subseteq U_i$. If I is finite, then the ε_i have a minimum ε, and the ball of centre x and radius ε is contained in all of the $B_{\varepsilon_i}(x)$, hence in the intersection of the U_i. Therefore, open sets are closed under finite intersections. $\square$

Definition 2.2.4. Let d and d' be two metrics on a set X. We say that d and d' are *equivalent* for all points $x \in X$, and all $r > 0$ there exist $r', r'' > 0$ such that $B_{r'}^{d'}(p) \subseteq B_r^d(p)$ and $B_{r''}^{d}(p) \subseteq B_r^{d'}(p)$.

Lemma 2.2.5. *Two metrics d and d' are equivalent if and only if they determine the same collection of open sets.*

Proof. Let d and d' be equivalent metrics. Since the definition of equivalence of metrics is symmetric, it is enough to show that all the open sets with respect to the metric d are also open with respect to the metric d'.

So let U be open with respect to d. For $x \in U$, let $\varepsilon > 0$ be such that $B_\varepsilon^d(x) \subseteq U$. By the definition of equivalence, there exists $\varepsilon' > 0$ such that $B_{\varepsilon'}^{d'}(x) \subseteq B_\varepsilon^d(x)$, hence in particular $B_{\varepsilon'}^{d'}(x) \subseteq U$.

Conversely, suppose d and d' have the same open sets. Fix $x \in X$ and $r > 0$. Then $B_r^d(x)$ is open with respect to d by Exercise 1, hence with respect to d'. In particular, there exists $r' > 0$ such that $B_{r'}^{d'}(x) \subseteq B_r^d(x)$. The other half of the definition of equivalence can be proved similarly. It follows that d and d' are equivalent. $\square$

Exercise 2. Let p be a real number with $p \geq 1$. Define the distance d_p on $\mathbb{R}^n$ by:

$$d_p(x, y) := \|x - y\|_p,$$

where $\| - \|_p$ is the p-norm, defined by:

$$\|x\|_p = \left(\sum_{i=1}^{n} |x_i|^p \right)^{1/p}.$$

This definition can be extended to $p = \infty$, by setting:

$$\|x\|_\infty = \max_{i=1,\dots,n} |x_i|.$$

Show that all the distances d_p are pairwise equivalent. [*Hint: compare* $\|x\|_p$ *and* $\|x\|_\infty$]

2.3 Continuous functions on metric spaces

Definition 2.3.1. Let X and Y be metric spaces, and $f : X \to Y$ be a function between them. If $x \in X$, we say that f is *continuous at* x if for all $\varepsilon > 0$ there exists $\delta > 0$ such that $f(B_\delta(x)) \subseteq B_\varepsilon(f(x))$. We say that f is *continuous* if it is continuous at every point of X.

Note that a continuous function f is not necessarily *distance-preserving*, i.e. it is not necessarily the case that $d(f(x), f(y)) = d(x, y)$. However, the converse holds, and it is an immediate consequence of the following exercise.

Exercise 3. Let $f : X \to Y$ be a function satisfying

$$d(f(x), f(y)) \leq d(x, y).$$

Prove that f is continuous.

Exercise 4. Let $p \in X$ be a point of a metric space X. Show that the function $f : X \to \mathbb{R}$ given by $f(x) = d(x, p)$ is continuous.

Definition 2.3.1 generalises the well-known notion of continuity of real and complex functions to metric spaces. A key observation is that continuity of functions depends only on the equivalence class of the metric, since it can in fact be characterised by the open sets only.

Proposition 2.3.2. *A function $f : X \to Y$ between metric spaces is continuous if and only if for all open subsets U of Y, the inverse image $f^{-1}(U)$ is an open subset of X.*

Proof. Let f be continuous at every point of X, and U an open subset of Y. Let $x \in f^{-1}(U)$. Since $f(x) \in U$, there exists $\varepsilon > 0$ such that $B_\varepsilon(f(x)) \subseteq U$, hence by continuity there exists $\delta > 0$ with $f(B_\delta(x)) \subseteq B_\varepsilon(f(x))$. It follows that $B_\delta(x) \subseteq f^{-1}(U)$, hence $f^{-1}(U)$ is open.

Conversely, suppose that inverse images of open sets along f are open, fix a point $x \in X$, and let $\varepsilon > 0$. Since $f^{-1}(B_\varepsilon(f(x)))$ is an open subset of X, there exists $\delta > 0$ such that $B_\delta(x) \subseteq f^{-1}(B_\varepsilon(f(x)))$, and this is exactly equivalent to the definition of continuity at x. □

3 Topological spaces

We have seen that one can use metric spaces as a setting for studying continuous transformations, but that metric spaces are equipped with "too much information", since continuity of functions does not see the difference between equivalent metrics, and is only sensitive to the open sets of a space.

This suggests that the appropriate setting for studying continuous functions is some kind of space where the open sets are the primary notion. It turns out that a very direct approach to this idea works very well. Namely, we can study sets equipped with a collection of subsets satisfying some of the properties that open sets of a metric space satisfy. This is captured by the following definition:

Definition 3.0.1. A *topology* on a set X is a collection $\tau \subseteq \mathbb{P}(X)$ of subsets of X that is stable under finite intersections and arbitrary unions. A set equipped with a topology is called a *topological space*.

Note that since empty intersections and empty unions are also included, a topology is required, in particular, to contain the whole set X and the empty set. We will sometimes refer to a topological space simply as a "space".

The following is a restatement of Lemma 2.2.3:

Corollary 3.0.2. *The open sets of a metric space form a topology.*

Therefore, if X is a metric space, we will implicitly regard it as a topological space with the topology provided by its open sets. We say that this topology is *induced* from the metric. As observed before, equivalent metrics on a space determine the same open sets, which means that they induce the same topology.

By analogy with the case of metric spaces, we will call a subset of a topological set X *open* if it belongs to the chosen topology on X.

3.1 Basic examples

It follows from Corollary 3.0.2 that any metric space provides an example of topological space, hence in particular $\mathbb{R}^n$ and $\mathbb{C}^n$ can canonically be regarded as topological spaces using the topology induced from the Euclidean metric.

There are many situations, however, where a topology can be obtained from a metric, but it more convenient and less ad hoc to define it directly in terms of open sets, rather than defining a metric first. Furthermore, certain topologies are just not induced from a metric, as we will see later.

To define a topology on a set X, we simply *declare* some of its subsets to be open. Provided that open sets satisfy the conditions of Definition 3.0.1, i.e. that they are stable under finite intersections and binary unions, such a choice uniquely determines a topology on X. We now introduce two canonical topologies that can be defined for *any* set.

Definition 3.1.1. Let X be a set. The *chaotic* topology on X (sometimes also called *indiscrete* or *codiscrete*) is defined by declaring only the empty set and X to be open. Conversely, the *discrete* topology on X is defined by declaring *every* subset of X to be open.

It is clear that the chaotic and discrete topologies on X are respectively the minimum and maximum among the collection of all topologies on X, ordered by inclusion. We will sometimes use the notation ΔX (resp. ∇X) to denote the topological space obtained by equipping X with the discrete (resp. chaotic) topology.

Exercise 5. Show that the *discrete metric* on a set X defined in section 2 (i.e. $d(x, y) = 1$ for $x \neq y$) induces the discrete topology.

Exercise 6. Show that if the chaotic topology on X is induced by a metric, then X has at most one point.

Let 2 denote a set with two elements, namely $2 = \{0, 1\}$. We have already seen that there are at least two topologies on 2, the discrete and chaotic ones. However, there is another interesting topology on this set, namely the one where we declare $\{1\}$ to be open, but not $\{0\}$ (and, of course, also the one obtained by reversing the roles of 0 and 1). The set 2, equipped with this topology, is called the *Sierpinski space*, and denoted $\mathbb{S}$.

3.2 Continuous functions

Inspired by Proposition 2.3.2, we can now give a definition of *continuous function* that works at the level of generality of topological spaces.

Definition 3.2.1. Let X and Y be topological space, and $f : X \to Y$ a function. We say that f is continuous if for all open subsets U of Y, the inverse image $f^{-1}(U)$ is an open set of X.

Proposition 3.2.2. *The identity function on a topological space X is continuous. If $f : X \to Y$ and $g : Y \to Z$ are continuous functions between topological spaces, then $g \circ f$ is continuous.*

Proof. If U is an open subset of X, then $\mathsf{id}^{-1}(U) = U$, hence it is open.

As for composition, just observe that $(g \circ f)^{-1}(U) = f^{-1}(g^{-1}(U))$. Now, if U is open, so is $V := g^{-1}(U)$ by continuity of g, and therefore also $f^{-1}(V)$ by continuity of f. □

Continuous functions play the role of the "morphisms" of topological spaces, just like group homomorphisms for groups, linear functions for vector spaces, and so on. Technically, Proposition 3.2.2 can be summarised by saying that topological spaces and continuous functions form a *category*.

Just like in groups and vector spaces, topological spaces come with a notion of isomorphism, which is more specifically called *homeomorphism*:

Definition 3.2.3. A *homeomorphism* between topological spaces X and Y is a continuous function $f : X \to Y$ such that there exists a continuous function $g : Y \to X$ satisfying $g \circ f = \mathsf{id}$ and $f \circ g = \mathsf{id}$. We say that X and Y are homeomorphic if there exists a homeomorphism between them.

Homeomorphic spaces are completely indistinguishable from the topological point of view. This is because every property of topological spaces can be expressed in terms of open sets, and therefore can be transported along a homeomorphism. We will write $X \cong Y$ to mean that X and Y are homeomorphic.

Note that a bijective continuous function is not necessarily a homeomorphism. A trivial example is the identity function $\Delta 2 \to \mathbb{S}$ (recall that $\mathbb{S}$ denotes the Sierpinski space, i.e. the set 2, with the topology where $\{1\}$ is open, and $\{0\}$ is not). We will see more "geometrical" examples later (Exercise 56).

Exercise 7. Let X be a set. Recall that ΔX and ∇X denote respectively the discrete and chaotic spaces with underlying set X. Prove that the identity function $\Delta X \to \nabla X$ is continuous and bijective, but it is a homeomorphism if and only if X contains at most one point.

Definition 3.2.4. A function $f : X \to Y$ between topological space is said to be *open* if it maps open subsets of X to open subsets of Y.

Exercise 8. Give an example of a continuous function that is not open. [*Hint: try a constant function*]

Proposition 3.2.5. *Let $f : X \to Y$ be a continuous and open bijection between topological spaces. Then f is a homeomorphism.*

Proof. Let $g : Y \to X$ be the inverse of f. We only need to show that g is continuous. If $U \subseteq X$ is an open set, then $g^{-1}(U) = f(U)$, hence it is open in Y. □

Spaces with the discrete and chaotic topologies enjoy certain so-called *universal properties*, which can be stated in terms of continuous functions respectively to and from them.

Proposition 3.2.6. *If Y is a topological space with the chaotic topology and X is any topological space, then all functions $X \to Y$ are continuous. Dually, if X has the discrete topology and Y is any topological space, then all functions $X \to Y$ are continuous.*

Proof. A function $f : X \to Y$ is continuous if and only if the inverse images along f of the open sets of Y are open subsets of X. But if the topology on Y is chaotic, the open sets of Y are only $\emptyset$ and Y, and their inverse images are $\emptyset$ and X, which are open in X.

If instead X is discrete, then every subset of X is open, hence the condition of the definition of continuity for a function $f : X \to Y$ holds trivially. □

The Sierpinski space $\mathbb{S}$ also possesses a kind of universal property, as the following exercise shows.

Exercise 9. Show that the open sets of a topological space X are in bijective correspondence with continuous functions $X \to \mathbb{S}$.

3.3 Closed sets

Definition 3.3.1. A subset C of a topological space X is said to be *closed* if its complement $X \setminus C$ is open.

Note that the notion of closed is in the appropriate sense *dual* to the notion of open, but it is not opposite. In particular, a set can be both open and closed (for instance, the whole space X is both open and closed, as well as the empty subset), or neither open nor closed.

Exercise 10. Show that a half-open interval $[a, b[$, with $a < b$, is neither open nor closed in $\mathbb{R}$.

Closed sets satisfy properties which are dual to those satisfied by open sets. Here by dual we mean that they are formally obtained by replacing intersections with unions and vice versa. More explicitly:

Proposition 3.3.2. *Closed sets in a topological space are stable under finite unions and arbitrary intersections.*

Proof. The complement of a finite union of closed sets is equal to the intersection of the complements, which are open. Since open sets are stable under finite intersections, the intersection is open, which means that its complement is closed. The proof of stability under intersections is completely analogous. $\square$

Since closed sets determine and are completely determined by open sets, they contain exactly the same information as open sets. In particular, the notion of topology could have been defined by axiomatising the properties of closed sets, and then open sets could have been defined accordingly as those sets whose complement is closed. This is summarised by the following.

Proposition 3.3.3. *Let χ be a collection of subsets of a set X that is stable under finite unions and arbitrary intersections. Then there exists a unique topology τ on X such that χ is the collection of closed sets of X.*

We can use Proposition 3.3.3 whenever it happens to be more natural to define a topology in terms of closed sets. A common example is the following.

Definition 3.3.4. Let X be any set. The *cofinite topology* on X is defined by declaring a set to be closed if it is either finite, or the whole space X.

Since finite sets are clearly stable under finite unions and non-empty intersections, Definition 3.3.4 does indeed define a topology by virtue of Proposition 3.3.3. It is called *cofinite* because the open sets are the complements of finite sets (plus of course the empty set). In the special case where $X = k$

is a field (or, more generally, a commutative ring), the cofinite topology on k coincides with the so-called *Zariski topology* of the affine line, used in algebraic geometry.

Exercise 11. Prove that the cofinite topology on X coincides with the discrete topology if and only if X is finite.

Many concepts defined in terms of open sets admit characterisations in terms of closed sets. For example:

Proposition 3.3.5. *A function $f : X \to Y$ is continuous if and only if the inverse image along f of any closed subset of Y is a closed subset of X.*

Proof. If f is continuous, let $C \subseteq Y$ be a closed set. Then $f^{-1}(C) = X \setminus f^{-1}(Y \setminus C)$. Since $Y \setminus C$ is open, it follows that $f^{-1}(Y \setminus C)$ is open, hence $f^{-1}(C)$ is closed, as required. The other direction is analogous. $\square$

Definition 3.3.6. Let $f : X \to Y$ be a function between topological spaces. We say that f is *closed* if it maps closed subsets of X to closed subsets of Y.

Exercise 12. Let $\pi : \mathbb{R}^2 \to \mathbb{R}$ be the projection onto the x axis. Show that π is open, but not closed. [*Hint: consider the set of points $(x, y) \in \mathbb{R}^2$ such that* $xy = 1$].

In analogy with Definition 3.2.4, we have the following:

Proposition 3.3.7. *Let $f : X \to Y$ be a continuous closed bijection. Then f is a homeomorphism.*

Proof. Let $g : Y \to X$ be the inverse of f. If $C \subseteq X$ is closed, then $g^{-1}(C) = f(C)$ is closed too, because f is closed. By Proposition 3.3.5, g is continuous, hence f is a homeomorphism. $\square$

3.4 Interior and closure

Given a subset of a topological space, there is a canonical way to "turn it into an open set", which is expressed by the following:

Definition 3.4.1. Let S be a subset of a topological space X. The *interior* of S, denoted S°, is the union of all the open sets contained in S.

Proposition 3.4.2. *Let S, T be subsets of a topological space X. Then:*

(i) $S^{\circ} \subseteq S$;
(ii) S° *is open;*
(iii) S° *is the largest open set contained in* S;
(iv) *if* $S \subseteq T$, *then* $S^{\circ} \subseteq T^{\circ}$;
(v) S *is open if and only if* $S^{\circ} = S$;
(vi) $(S^{\circ})^{\circ} = S^{\circ}$.
(vii) $(S \cap T)^{\circ} = S^{\circ} \cap T^{\circ}$.

Proof. Property (i) is immediate from the definition. Property (ii) follows from the fact that open sets are stable under union, and S^{o} is defined as a union of open sets. It then follows that S^{o} is an open set contained in S, and since it is the union of all of those sets, it must be the largest such, hence we have (iii).

As for (iv), if $S \subseteq T$, then it follows from (i) that $S^{\mathrm{o}} \subseteq T$, so S^{o} is an open set contained in T, hence $S^{\mathrm{o}} \subseteq T^{\mathrm{o}}$ from property (iii).

Now assume that S is open. Then clearly S is the largest open set contained in S, therefore $S^{\mathrm{o}} = S$. The converse is obvious since S^{o} is open. This proves (v). Also, since S^{o} is open, property (vi) follows directly from (v).

Finally, to show property (vii), we show the two inclusions. First, from $S \cap T \subseteq T$ and property (iv), we get that $(S \cap T)^{\mathrm{o}} \subseteq T^{\mathrm{o}}$, and symmetrically $(S \cap T)^{\mathrm{o}} \subseteq S^{\mathrm{o}}$, therefore $(S \cap T)^{\mathrm{o}} \subseteq S^{\mathrm{o}} \cap T^{\mathrm{o}}$. Conversely, since $S^{\mathrm{o}} \subseteq S$ and $T^{\mathrm{o}} \subseteq T$, we have $S^{\mathrm{o}} \cap T^{\mathrm{o}} \subseteq S \cap T$, hence by (iii) $S^{\mathrm{o}} \cap T^{\mathrm{o}} \subseteq (S \cap T)^{\mathrm{o}}$. □

In particular, property (v) of Proposition 3.4.2 shows that open sets are completely determined by the operation of taking the interior of a subset of X. In fact, it is possible to give an equivalent definition of topological space purely based on an axiomatisation of the interior operation, as the following proposition shows.

Proposition 3.4.3. *Let X be a set equipped with an operation $i : \mathbb{P}(X) \to \mathbb{P}(X)$ satisfying the following axioms:*

(i) $i(S) \subseteq S$;
(ii) $i(i(S)) = i(S)$;
(iii) $i(X) = X$;
(iv) $i(S \cap T) = i(S) \cap i(T)$.

Then there exists a unique topology on X such that $i(S) = S^{\mathrm{o}}$.

Proof. Uniqueness is clear, since if there is a topology for which i is the interior operator, then it must follow from Proposition 3.4.2 that a set is open if and only if $i(S) = S$, which uniquely determines the topology.

To show existence, then, we declare a set S open if $i(S) = S$, and show the properties of Definition 3.0.1. Observe first of all that if $S \subseteq T$, then $i(S) = i(S \cap T) = i(S) \cap i(T)$, hence $i(S) \subseteq i(T)$. Now let $(S_j)_{j \in J}$ be a family of open sets and let U be their union. We know from axiom (i) that $i(U) \subseteq U$. To show the converse, let j be any index in J. Since $S_j \subseteq U$, it follows from the observation above that $S_j = i(S_j) \subseteq i(U)$. Therefore, by taking the union over j, we get that $U \subseteq i(U)$, as required.

As for stability under finite intersections, it is enough to show that X is open, and that open sets are stable under binary intersections, since all other finite intersections follow easily by induction. Clearly X is open by axiom (iii), so let S and T be open sets. Then $i(S \cap T) = i(S) \cap i(T) = S \cap T$, which means that $S \cap T$ is also open, completing the proof that the above choice of open sets defines a topology on X.

It remains to show that with this choice of topology, $i(S) = S^{\mathrm{o}}$. First of all, axiom (ii) implies that every set of the form $i(S)$ is open. In particular, $i(S)$ is an open subset of S. Let U be any open subset of S. From the definition of open, we get $U = i(U) \subseteq i(S)$. Therefore, $i(S)$ is the maximum open subset of S, which implies that $i(S) = S^{\mathrm{o}}$, as required. $\square$

Analogous to the interior operation (and *dual* in an appropriate sense), there is a *closure* operation on any topological space.

Definition 3.4.4. Let S be a subset of a topological space X. The *closure* of S, denoted $\overline{S}$, is the intersection of all the closed sets containing S. A subset S of X is said to be *dense* if $\overline{S} = X$.

Exercise 13. A subset S of a topological space X is dense if and only if every non-empty open subset of X meets S.

Exercise 14. Let S be a dense subset of a topological space X, and $f : X \to Y$ a surjective continuous function. Then $f(S)$ is dense in Y.

We leave it to the reader to formulate and prove the results about closure analogous to Proposition 3.4.2 and Proposition 3.4.3.

The duality between open and closed sets implies a similar relationship between interior and closure. For example, we have the following result:

Proposition 3.4.5. *For any subset S of a topological space X,*

$$X \setminus \overline{S} = (X \setminus S)^{\mathrm{o}}.$$

Proof. Just a direct calculation using the definitions:

$$X \setminus \overline{S} = X \setminus \bigcap_{C \supseteq S} C = \bigcup_{C \supseteq S} X \setminus C = \bigcup_{U \subseteq X \setminus S} U = (X \setminus S)^{\mathrm{o}},$$

where C ranges over closed sets, U ranges over open sets, and we have used the fact that the open sets contained in $X \setminus S$ are exactly the complements of the closed sets containing S. $\square$

For any subset S, we have of course that $S^{\mathrm{o}} \subseteq \overline{S}$, since the former is contained in S, and the latter contains S. This motivates the following definition:

Definition 3.4.6. The *boundary* of a set S is the difference $\overline{S} \setminus S^{\mathrm{o}}$.

Exercise 15. Prove that the boundary of any set S is closed.

Exercise 16. Let S be a subset of $\mathbb{R}$. If S has a supremum m, then $m \in \partial S$.

3.5 Neighbourhoods and continuity at a point

Continuity can also be defined *locally*, i.e. around any given point, generalising the definition of continuity at a point for functions of metric spaces. In order to do that, we first establish the notion of *neighbourhood*, which will be of central importance later.

Definition 3.5.1. Let x be a point of a topological space X. A *neighbourhood* of x is a subset N such that there exists an open set U with $x \in U$ and $U \subseteq N$.

Proposition 3.5.2. *Let x be a point of a topological space X. The following properties hold:*

(i) every neighbourhood of x contains x;
(ii) every open set containing x is a neighbourhood of x;
(iii) if N is a neighbourhood of x, and $N \subseteq N'$, then N' is a neighbourhood of x;
(iv) if N, N' are neighbourhoods of x, then so is $N \cap N'$.

Proof. Properties (i) and (ii) are obvious from the definition. As for (iii), if N is a neighbourhood of x, then $x \in U$, for some open set U contained in N. But then $U \subseteq N \subseteq N'$, hence N' is also a neighbourhood of x.

Finally, let N, N' be neighbourhoods of x. Let U, U' be open sets containing x, with $U \subseteq N$ and $U' \subseteq N'$. Then $U \cap U'$ is also open, it contains x, and $U \cap U' \subseteq N \cap N'$, hence $N \cap N'$ is a neighbourhood of x. ☐

We will see later that the collection of neighbourhoods of a point is the primary example of a *filter*, and this statement is essentially the content of Proposition 3.5.2.

Similarly to what we have already seen for closed sets, interior and closure, neighbourhoods contain enough information to reconstruct the whole topology on a space.

Proposition 3.5.3. *A subset U of a topological space X is open if and only if it is a neighbourhood of all its points.*

Proof. If U is open, then it is clearly a neighbourhood of all its points. Conversely, suppose that U is a neighbourhood of all its points. For all $x \in U$, let V_x be an open set containing x and such that $V_x \subseteq U$. Now define:

$$U' := \bigcup_{x \in U} V_x.$$

If $x \in U$, then we know that $x \in V_x$, hence in particular $x \in U'$, therefore $U \subseteq U'$. On the other hand, $V_x \subseteq U$ for all $x \in U$, hence by taking the union we also get $U' \subseteq U$. It follows that $U = U'$, so U is a union of open sets, hence open. ☐

Using neighbourhoods, it is possible to formulate a notion of continuity at a specific point.

Definition 3.5.4. Let X and Y be topological spaces, $f : X \to Y$ a continuous function, and $x \in X$ a point. We say that f is *continuous at* x if the inverse image along f of any neighbourhood of $f(x)$ is a neighbourhood of x.

Exercise 17. Prove that a function $f : X \to Y$ is continuous at x if and only if for all open neighbourhoods U of $f(x)$, the inverse image through f of U is a neighbourhood of x.

As expected, continuity at a point is the correct local formulation of Definition 3.2.1.

Proposition 3.5.5. *Let $f : X \to Y$ be a function between topological spaces. Then f is continuous if and only if it is continuous at every point of x.*

Proof. First suppose that f is continuous, and let $x \in X$ be any point. By Exercise 17, it is enough to show that for all open sets U containing $f(x)$, the inverse image $f^{-1}(U)$ is a neighbourhood of x. But $f^{-1}(U)$ is open by continuity of f, and it clearly contains x, so we are done.

Conversely, suppose f is continuous at every point of X. Let U be any open subset of Y. To show that $f^{-1}(U)$ is open, we prove that it is a neighbourhood of all its points (Proposition 3.5.3). Let $x \in f^{-1}(U)$ be any point. Then $f(x) \in U$, which implies that U is a neighbourhood of $f(x)$, and therefore by continuity at x we have that $f^{-1}(U)$ is a neighbourhood of x, as required. □

The following characterisations of interior and closure in terms of neighbourhoods are often useful.

Proposition 3.5.6. *Let S be a subset of a topological space X. A point $x \in X$ is in the interior of S if and only if S is a neighbourhood of x.*

Proof. It is enough to show that $x \in S^\circ$ if and only if there is an *open* neighbourhood of x contained in S, which is just a restatement of the definition of interior. □

Proposition 3.5.7. *Let S be a subset of a topological space X. A point $x \in X$ is in the closure of S if and only if every neighbourhood of x meets S.*

Proof. By negating the equivalence, it is enough to show that $x \notin \overline{S}$ if and only if there exists a neighbourhood of x that does not intersect S. This follows immediately from Proposition 3.4.5 and Proposition 3.5.6. □

4 Convergence and bases

4.1 Sequences

In metric spaces, it is possible to give equivalent characterisations of many topological notions using convergence of sequences. Let us a review the basic definitions:

Definition 4.1.1. A *sequence* in a set X is simply a function $x : \mathbb{N} \to X$.

If $n \in \mathbb{N}$ is a natural number, we will use the notation x_n to denote the value of the sequence x at the point n. Often, we will define sequences by writing their generic element in brackets with the index as a subscript, like so:

$$(x_n)_n.$$

Given a subset $A \subseteq X$, we will sometimes say that a sequence x in X is *eventually* in A if there exists $n_0 \in \mathbb{N}$ such that $x_n \in A$ for all $n \geq n_0$. Similarly, we will say that x is *frequently* in A if for all $n_0 \in \mathbb{N}$ there exists $n \geq n_0$ such $x_n \in A$. Note that x is frequently in A if and only if it is not eventually in the complement of A.

Definition 4.1.2. A sequence x in a topological space X is said to *converge* to a point ℓ if for all neighbourhoods N of ℓ there exists a natural number n_0 such that $x_n \in N$ for all $n \geq n_0$. The point ℓ is called a *limit* of the sequence x.

In other words, a sequence x converges to ℓ if for all neighbourhoods N of ℓ, x is eventually in N.

Exercise 18. Prove that for all points ℓ of a topological space X, the constant sequence $(\ell)_n$ converges to ℓ.

We will write $x \to \ell$ to mean that the sequence x converges to ℓ. It is clear that limits of sequences may not exist (take for example the sequence $((-1)^n)_n$ in $\mathbb{R}$). What may be slightly more surprising is that, at the level of generality of topological spaces, limits of sequences may not be unique, as the following exercise shows.

Exercise 19. Let X be a topological space with the chaotic topology. Then any sequence in X converges to any point of X.

One fundamental property of convergence of sequences is that it is preserved by continuous functions between any topological spaces. First note that, if $f : X \to Y$ is a function, and x is a sequence in X, the composed function $f \circ x$ is a sequence in Y.

Proposition 4.1.3. *Let $f : X \to Y$ be a continuous function between topological spaces. For all sequences x in X, if $x \to \ell$, then $f \circ x \to f(\ell)$.*

Proof. Let x be a sequence in X converging to ℓ and let N be any neighbourhood of $f(\ell)$. By Proposition 3.5.5, we know that $f^{-1}(N)$ is a neighbourhood of ℓ, therefore there exists $n_0 \in \mathbb{N}$ such that $x_n \in f^{-1}(N)$ for all $n \geq n_0$, which means that $f(x_n) \in N$ for all $n \geq n_0$. We have thus shown that $f \circ x \to f(\ell)$, as required. □

In metric spaces, we characterise closed sets as those such that convergent sequences in them cannot have a limit outside. More precisely:

Proposition 4.1.4. *Let X be a metric space. The closure of a subset S of X is the set of limits of sequences in S. In particular, a subset C of X is closed if and only if for all sequences x, if $x_n \in C$ for all n, and $x \to \ell$, then $\ell \in C$.*

Proof. If x is a sequence in S and $x \to \ell$, then every neighbourhood of ℓ contains at least one of the x_n (infinitely many, as a matter of fact), therefore ℓ is in the closure of S by Proposition 3.5.7. Conversely, let $\ell \in \overline{C}$. For all $n \in \mathbb{N}$, we know that the open ball $B_{2^{-n}}(\ell)$ is a neighbourhood of ℓ, hence there exists a point $x_n \in B_{2^{-n}}(\ell) \cap C$. We want to show that the sequence $(x_n)_n$ converges to ℓ. Let N be any neighbourhood of ℓ. By definition of the topology induced by a metric, there must be some $\varepsilon > 0$ such that $B_\varepsilon(\ell) \subseteq N$. If $n_0 \in \mathbb{N}$ is such that $2^{-n_0} \leq \varepsilon$, we get that for all $n \geq n_0$, $B_{2^{-n}}(\ell) \subseteq N$, hence in particular $x_n \in N$, as required by the definition of convergence.

The second statement follows immediately. □

Exercise 20. Show that one of the directions of Proposition 4.1.4 can be generalised to topological spaces, namely: for any subset S of a topological space X, if x is a sequence in S that converges to a point $\ell \in X$, then $\ell \in \overline{S}$.

Unfortunately, it is not immediately possible to generalise the other direction of Proposition 4.1.4 to topological spaces.

To construct a counterexample, let X be any set, and declare a subset in X *closed* if it is either countable or the whole space X (cf. Definition 3.3.4). Recall that a set A is said to be *countable* if there exists a surjection $\mathbb{N} \to A$. Since finite unions and non-empty intersections of countable subsets are countable, it is clear that this determines a topology on X (using Proposition 3.3.3). This topology is referred to as the *cocountable topology* on X, by analogy with the cofinite topology (Definition 3.3.4).

It turns out that, using convergence of sequences, we cannot distinguish a space with the cocountable topology from a discrete space:

Proposition 4.1.5. *Let X be a space with the cocountable topology. Then every convergent sequence in X is eventually constant.*

Proof. Let x be a sequence, and assume that $x \to \ell$. Let S be the *image* of the sequence, i.e. the set of points $y \in X$ such that there exists $n \in \mathbb{N}$ with $x_n = y$. Since $\mathbb{N}$ is countable, S is countable as well, hence so is $S' = S \setminus \{\ell\}$. Now, $X \setminus S'$ is a neighbourhood of ℓ, hence x is eventually in the complement S', which implies that x is eventually in $\{\ell\}$, i.e. it is eventually constant. □

It then follows from Proposition 4.1.5 that Proposition 4.1.4 cannot be generalised to uncountable spaces with the cocountable topology.

Corollary 4.1.6. *Let X be an uncountable set (for example the set $\mathbb{R}$ of real numbers) equipped with the cocountable topology, and $x \in X$. Let $S = X \setminus \{x\}$. Then $x \in \overline{S}$, but there is no sequence in S that converges to x.*

Proof. Since X is uncountable, so is S. Therefore, the only closed set containing S is X itself, hence $\overline{S} = X$, from which the first assertion follows. The second is an immediate consequence of Proposition 4.1.5. □

4.2 Filters

To address the fundamental inability of sequences to correctly capture topological properties in the general case, we will now introduce a new, more general, type of convergence. The objects that will replace sequences are called *filters*:

Definition 4.2.1. Let X be a set. A *filter* on X is a collection $\mathcal{F}$ of subsets of X, satisfying the following properties:

(i) $\emptyset \notin \mathcal{F}$;
(ii) the collection $\mathcal{F}$ is stable under finite intersections;
(iii) if $A \in \mathcal{F}$, and $A \subseteq B$, then $B \in \mathcal{F}$.

As in the definition of topology, the requirement of $\mathcal{F}$ to be stable under finite intersections implicitly contains the fact that $X \in \mathcal{F}$, since the whole set X can be regarded as an empty intersection. In particular, a filter is always non-empty.

Note that some authors omit condition (i), and instead call a filter *proper* if (i) is satisfied. Note that (i) is equivalent to requiring that the filter is not the entire powerset of X. In these notes, we will not need to make this distinction, hence we simply assume (i) whenever we speak about filters. Condition (iii) is often expressed by saying that $\mathcal{F}$ is *upwards closed*.

One way to think about filters is that they give a notion of "largeness" for subsets of a given set X. The axioms of Definition 4.2.1 capture properties that "large" sets ought to satisfy, for any given interpretation of the word "large". The connection with sequences is given by the following definition.

Definition 4.2.2. Let x be a sequence on a set X. The *filter associated to x*, denoted $\mathsf{Ev}(x)$, is defined to be the collection of subsets A of X such that x eventually belongs to A.

More explicitly, $A \in \mathsf{Ev}(x)$ if and only if there exists $n_0 \in \mathbb{N}$ such that $x_n \in A$ for all $n \geq n_0$. Clearly $\mathsf{Ev}(x)$ does not contain the empty set, and it is easy to see that it is indeed stable under finite intersections and upwards closed, hence $\mathsf{Ev}(x)$ is a filter.

The reason why filters can be used instead of sequences is that there is a notion of convergence for them.

Definition 4.2.3. Let X be a topological space, $\mathcal{F}$ a filter on X, and $\ell \in X$ a point. We say that $\mathcal{F}$ converges to ℓ, written $\mathcal{F} \to \ell$, if $\mathcal{F}$ contains all the neighbourhoods of ℓ.

The notion of convergence of filters directly generalises that of convergence of sequence, as the following proposition shows.

Proposition 4.2.4. *Let X be a topological space, and x a sequence in X. Then x converges to ℓ if and only if $\mathsf{Ev}(x)$ converges to ℓ.*

Proof. By definition of convergence of filters, $\mathsf{Ev}(x)$ converges to ℓ if and only if every neighbourhood of ℓ is in $\mathsf{Ev}(x)$, i.e. if and only if for all neighbourhoods N of ℓ, x eventually belongs to N. The latter is simply a restatement of the definition of convergence of x to ℓ. □

A useful example of convergent filter that does not arise as the filter associated to a sequence is given by the neighbourhoods of a point.

Proposition 4.2.5. *Let x be a point in a topological space X. The collection of neighbourhoods of x is a filter on X, called the* filter of neighbourhoods *of X, and denoted $\mathcal{N}(x)$.*

Proof. Immediate from Proposition 3.5.2. □

Exercise 21. Show that $\mathcal{N}(x)$ converges to x.

4.3 Topology in terms of filters

As we have seen in Corollary 4.1.6, it is not possible to characterise closed (and hence open) sets in terms of convergence of sequences. One of the main reasons why filters are interesting in topology is that this can be remedied by replacing sequences with filters.

When considering sequences, it is useful sometimes to restrict one's attention to sequences whose elements all belong to a given subset. In order to generalise this idea to filters, we need to be able to regard a filter on a subset as a filter on the whole set. More generally, we define an operation to transport a filter along any function.

Definition 4.3.1. Let $f : X \to Y$ be any function, and $\mathcal{F}$ a filter on X. The filter $f_*(\mathcal{F})$ is defined as:

$$f_*(\mathcal{F}) = \{B \in Y \mid f^{-1}(B) \in \mathcal{F}\}.$$

We will refer to $f_*(\mathcal{F})$ as the *image* of $\mathcal{F}$ along f.

The proof that $f_*(\mathcal{F})$ is indeed a filter is left as an exercise for the reader. As a special case of Definition 4.3.1, if A is a subset of X, and $i^A : A \to X$ denotes the inclusion function, any filter $\mathcal{F}$ on the subset A determines a filter $i^A_*(\mathcal{F})$ on the whole set X. A natural question then arises: if $\mathcal{G}$ is a filter on X, under which conditions is X the image of a filter on a subset? The following proposition gives a partial answer.

Proposition 4.3.2. *Let $f : X \to Y$ be a function, and $\mathcal{G}$ a filter on Y. There exists a filter $\mathcal{F}$ on Y such that $f_*(\mathcal{F}) \supseteq \mathcal{G}$ if and only if, for all $B \in \mathcal{G}$, the inverse image $f^{-1}(B)$ is non-empty.*

Proof. If $f_*(\mathcal{F}) \supseteq \mathcal{G}$ for some $\mathcal{F}$, then for all $B \in \mathcal{G}$ we have that $f^{-1}(B) \in \mathcal{F}$ by definition of image, hence in particular it is non-empty.

Conversely, let $\mathcal{F}$ be defined as the collection of subsets of the form $f^{-1}(B)$, where B ranges over the elements of $\mathcal{G}$. The assumption on $\mathcal{G}$ says that $\mathcal{F}$ does not contain the empty set. It is also clear that $\mathcal{F}$ is upwards closed and stable under finite intersections, so $\mathcal{F}$ is a filter. Finally, if $B \in \mathcal{G}$, that means that $f^{-1}(B) \in \mathcal{F}$, hence $B \in f_*(\mathcal{F})$, therefore $f_*(\mathcal{F}) \supseteq \mathcal{G}$, as required. □

We are now ready to prove the generalisation of Proposition 4.1.4 to topological spaces, obtained by replacing sequences with filters.

Proposition 4.3.3. *Let X be a topological space, and S a subset of X. A point ℓ is in the closure of S if and only if there exists a filter $\mathcal{F}$ on S with $i^S_*(\mathcal{F}) \to \ell$.*

Proof. By Proposition 3.5.7, $\ell \in \overline{S}$ if and only if every neighbourhood of ℓ meets S. By Proposition 4.3.2 this is the same as saying that there exists a filter $\mathcal{F}$ on S such that $i^S_*(\mathcal{F}) \supseteq \mathcal{N}(\ell)$, or in other words that there exists a filter $\mathcal{F}$ on S such that its image converges to ℓ, as required. □

Corollary 4.3.4. *A subset C of a topological space X is closed if and only if for all filters $\mathcal{F}$ on C, if $i^C_*(\mathcal{F}) \to \ell$, then $\ell \in C$.*

We can establish a similar characterisation for open sets in terms of filters.

Proposition 4.3.5. *Let X be a topological space, and S a subset of X. A point ℓ is in the interior of S if and only if all the filters that converge to ℓ contain U.*

Proof. By Proposition 3.5.6, $\ell \in S^{\circ}$ if and only if S is a neighbourhood of ℓ. Clearly, a set N is a neighbourhood of a point x if and only if all filters that converge to x contain N. The required equivalence then follows. □

Finally, we show how continuity of functions can also be expressed in terms of convergence of filters.

Lemma 4.3.6. *Let $f : X \to Y$ be a function between topological spaces, and $x \in X$ any point. The function f is continuous at x if and only if the filter $f_*(\mathcal{N}(x))$ converges to $f(x)$.*

Proof. By definition, f is continuous at x if and only if for all $N \in \mathcal{N}(f(x))$, we have that $f^{-1}(N) \in \mathcal{N}(x)$. This condition can be equivalently stated as $\mathcal{N}(f(x)) \subseteq f_*(\mathcal{N}(x))$, which is exactly the statement that $f_*(\mathcal{N}(x))$ converges to $f(x)$. □

Corollary 4.3.7. *A function $f : X \to Y$ between topological spaces is continuous if and only if it preserves converges of filters, i.e. for all filters $\mathcal{F}$ on X, if $\mathcal{F}$ converges to some point x, then $f_*(\mathcal{F})$ converges to $f(x)$.*

Proof. By Proposition 3.5.5, f is continuous if and only if it is continuous at every point, which by Lemma 4.3.6 is equivalent to preserving convergence of filters at every point. $\square$

4.4 The set of topologies as a lattice

Definition 4.4.1. Let τ and ρ be topologies on a set X. We say that τ is *coarser* than ρ if $\tau \subseteq \rho$ as subsets of the powerset of X. Equivalently, we can say that ρ is *finer* than τ in this case.

Definition 4.4.1 introduces some suggestive terminology for the natural partial order on topologies. Saying that τ is *coarser* than ρ means that fewer sets are open according to τ than according to ρ. The fact that this is a partial order order follows immediately from the fact that it is just the restriction on topologies of the usual order relation for subsets of the powerset of X.

Therefore, we will henceforth regard the set of topologies on a set X as a *partially ordered set* (or *poset*, for short). Note that the chaotic (resp. discrete) topology on X is the minimum (resp. maximum) element of this poset.

Proposition 4.4.2. *The poset of topologies on X is a complete lattice, i.e. every set of topologies on X has a least upper bound and a greatest lower bound.*

Proof. We first show that every set of topologies has a greatest lower bound. Let S be a set of topologies on X, and define τ to be their intersection. In other words, a subset A of X belongs to τ if and only if it belongs to all the topologies in S. It is then a simple verification to deduce the stability properties of τ from those of the topologies of S. Therefore, τ is a topology on X, and it is clearly coarser than every topology in S. To show that τ is the greatest lower bound of S, let ρ be any topology coarser than every element of S. If A is an element of ρ, than by the assumption on ρ we get that A belongs to every topology of S, which implies that it belongs to τ. We have then proved that ρ is coarser than τ, as required.

It is a general fact that a poset that has greatest lower bounds of arbitrary subsets has also least upper bounds. We will prove it here for the special case of the posets of topologies, but the proof is directly generalisable.

Let S be an arbitrary set of topologies on X. Let S' be the set of topologies that are finer than every topology in S (in other words, S' is the set of upper bounds of S). By the previous part of the proof, S' has a greatest lower bound τ. To show that τ is a least upper bound for S, it is enough to show that it is an upper bound, because it follows immediately from its definition that it is coarser than every upper bound. So let ρ be a topology in S, and let us show that τ is finer than ρ. Since τ is the greatest lower bound of S', it is enough to show that ρ is also a lower bound of S'. But if σ is a topology in S', then by definition of S' we have that σ is finer than ρ. That concludes the proof. $\square$

Using the terminology of Definition 4.4.1, we can say that the least upper bound of a set S of topologies is the *coarsest topology which is finer than every topology in S*. Dually, the greatest lower bound of S is the *finest topology that is coarser than every topology in S*.

Note that, although greatest lower bounds of topologies can simply be computed by intersection (just like in the containing poset of subsets of the powerset of X), least upper bounds are not just unions. The simplest counterexample is the chaotic topology, which can be regarded as the least upper bound of the empty set of topologies, and which is not simply the empty set.

Exercise 22. Find two topologies on a set X such that their union is not a topology.

4.5 Bases and subbases of topologies

We can make use of the order structure on the set of topologies to *generate* topologies from arbitrary collections of subsets.

Definition 4.5.1. Let $\mathcal{A}$ be any subset of the powerset of X, and let τ be a topology on X. We say that τ is *generated* by $\mathcal{A}$ if it is the coarsest topology containing $\mathcal{A}$.

Exercise 23. Prove that the chaotic topology on X is generated by $\mathcal{A} = \emptyset$, while the discrete topology is generated by $\mathcal{B} = \{\{x\} \mid x \in X\}$.

It turns out that every collection of subsets generates a topology, and it is possible to give an explicit description of it.

Proposition 4.5.2. *Let $\mathcal{A}$ be a subset of the powerset of X. Then the topology generated by $\mathcal{A}$ exists and it consists exactly of unions of finite intersections of elements of $\mathcal{A}$.*

Proof. Let τ be the set of unions of finite intersections of elements of $\mathcal{A}$. It is clear that τ is stable under unions. To show that τ is a topology, it is then enough to show that it is closed under finite intersections. Let U_i be a family of elements of τ, with $i = 1, \dots, m$. For all such i, we can write:

$$U_i = \bigcup_{j \in J_i} \bigcap_{k=1}^{n_{ij}} A_{ijk},$$

for some sets J_i, natural numbers n_{ij} and elements A_{ijk} of $\mathcal{A}$. Then, by distributivity of intersections across unions, we get:

$$\bigcap_{i=1}^{m} U_i = \bigcup_{t \in \prod_{i=1}^{m} J_i} \bigcap_{i=1}^{m} \bigcap_{k=1}^{n_{i,t_i}} A_{i,t_i,k},$$

which is also a union of finite intersections of elements of $\mathcal{A}$.

Therefore, τ is a topology that contains $\mathcal{A}$. To show that it is the coarsest such, let ρ be any topology that contains $\mathcal{A}$. Since ρ must be stable under finite intersections, ρ contains in particular all finite intersections of elements of $\mathcal{A}$. Furthermore, ρ is also stable under unions, hence it contains all unions of finite intersections of elements of $\mathcal{A}$, i.e. it contains τ, and we are done. □

If τ is the topology generated by a collection $\mathcal{A}$, we say that $\mathcal{A}$ is a *subbase* of τ. If $\mathcal{A}$ satisfies the condition that every finite intersection of elements of $\mathcal{A}$ is a union of elements of $\mathcal{A}$, then it follows from Proposition 4.5.2 that the topology τ generated by $\mathcal{A}$ consists simply of unions of elements of $\mathcal{A}$. In that case, we say that $\mathcal{A}$ is a *base* for the topology τ. It is clear that if $\mathcal{A}$ is stable under finite intersections, then in particular it is a base for the topology it generates.

Exercise 24. Show that the least upper bound of a set of topologies is the topology generated by their union.

Bases and subbases are useful, because they allow us to check continuity more easily, as the following proposition shows.

Proposition 4.5.3. *Let $f : X \to Y$ be a function between topological spaces, and let $\mathcal{A}$ be a subbase for the topology on Y. Then f is continuous if and only if for all $U \in \mathcal{A}$, $f^{-1}(U)$ is open in X.*

Proof. If f is continuous, then clearly $f^{-1}(U)$ is open for all $U \in \mathcal{A}$, since such subsets U are open in Y. Conversely, suppose that the inverse images along f of elements of $\mathcal{A}$ are open in X, let τ be the topology on Y, and let ρ be the collection of subsets S of Y such that $f^{-1}(S)$ is open in X. Then by assumption ρ contains $\mathcal{A}$, and it is easy to see that ρ is a topology. It follows that $\tau \subseteq \rho$, hence f is continuous. □

We now define initial and final topologies, which will be used several times for constructing topological spaces in section 5.

Definition 4.5.4. Let X be a set, and $f_i : X \to Y_i$ a collection of functions into topological spaces Y_i, where i ranges over an arbitrary set I. The *initial topology* on X induced by the f_i is the coarsest topology that makes all the f_i continuous.

Definition 4.5.5. Let Y be a set, and $f_i : X_i \to Y$ a collection of functions from topological spaces X_i, where i ranges over an arbitrary set I. The *final topology* on Y induced by the f_i is the finest topology that makes all the f_i continuous.

A priori, initial and final topologies might not exist, since not every collection of topologies admits a coarsest and a finest one. However, their existence follows from the following characterisations.

Proposition 4.5.6. *With the notation of Definition 4.5.4, the initial topology induced by the f_i is the topology generated by subsets of the form $f_i^{-1}(U)$, where U ranges over all the open sets of Y_i.*

Proof. If τ denotes the topology generated by the $f_i^{-1}(U)$, then clearly τ makes all the f_i continuous. Furthermore, any topology ρ that makes f_i continuous for all $i \in I$ must contain those subsets, hence it must be finer than τ. It follows that τ is the coarsest topology making the f_i continuous. $\square$

Proposition 4.5.7. *With the notation of Definition 4.5.5, the final topology induced by the f_i consists exactly of those subsets U such that $f_i^{-1}(U)$ is open for all $i \in I$.*

Proof. Let τ_i denote the collection of subsets U of Y such that $f_i^{-1}(U)$ is open in X_i. It is easy to see that τ_i is a topology and that f_i is continuous with respect to τ_i. Therefore, the greatest lower bound of all the τ_i makes *all* of the f_i continuous, and is the finest such by construction. Since greatest lower bounds are given by intersections, τ contains exactly those subsets U such that $f_i^{-1}(U)$ is open for all $i \in I$, as claimed. $\square$

Corollary 4.5.8. *Let Y be equipped with the final topology induced by the f_i. Then a set $C \subseteq Y$ is closed if and only if $f_i^{-1}(C)$ is closed for all $i \in I$.*

Proof. Immediate, because inverse images preserve complements. $\square$

One reason why initial and final topologies are interesting is that they enjoy the following universal properties.

Proposition 4.5.9. *With the notation of Definition 4.5.4, let Z be a topological space, and $g : Z \to X$ a function. Then g is continuous with respect to the initial topology on X induced by the f_i if and only if for all $i \in I$, the function $f_i \circ g$ is continuous.*

Proof. By Proposition 4.5.3 and Proposition 4.5.6, g is continuous if and only if $g^{-1}(f_i^{-1}(U))$ is open in Z for all $i \in I$ and U open in Y_i. Since $g^{-1} \circ f_i^{-1} = (f_i \circ g)^{-1}$, the latter condition is equivalent to $f_i \circ g$ being continuous for all $i \in I$. $\square$

Proposition 4.5.10. *With the notation of Definition 4.5.5, let Z be a topological space, and $g : Y \to Z$ a function. Then g is continuous with respect to the final topology on Y induced by the f_i if and only if for all $i \in I$, the function $g \circ f_i$ is continuous.*

Proof. By Proposition 4.5.7, g is continuous if and only if $f_i^{-1}(g^{-1}(U))$ is continuous for all $i \in I$ and all U open in Z. Now, $f_i^{-1} \circ g^{-1} = (g \circ f_i)^{-1}$, hence the latter condition is equivalent to all the functions $g \circ f_i$ being continuous. $\square$

Exercise 25. Show that the chaotic and discrete topologies are respectively the initial and final topologies induced by the empty family of functions.

Exercise 26. Let $f : X \to Y$ a continuous open (resp. closed) surjection between topological spaces. Show that Y has the final topology induced by the single map f.

4.6 Generating filters

Definition 4.6.1. Let $\mathcal{A} \subseteq \mathbb{P}(X)$ be a collection of subsets of a set X. We say that $\mathcal{A}$ has the *finite intersection property* if all finite subcollections of $\mathcal{A}$ have non-empty intersection.

Proposition 4.6.2. *Let $\mathcal{A} \subseteq \mathbb{P}(X)$ be a collection of subsets of a set X. Then $\mathcal{A}$ has the finite intersection property if and only if $\mathcal{A}$ is contained in a filter.*

Proof. It is clear that filters have the finite intersection property, so if $\mathcal{A} \subseteq \mathcal{F}$, where $\mathcal{F}$ is a filter, it follows that $\mathcal{A}$ also has the finite intersection property.

Conversely, let $\mathcal{A}$ have the finite intersection property. Define $\mathcal{F}$ to be the collection of all subsets S of X such that there exists a finite subcollection $\mathcal{A}_0 \subseteq \mathcal{A}$ with $\bigcap \mathcal{A}_0 \subseteq S$. It is clear that $\mathcal{F}$ is stable under finite intersections and upwards closed. If $\emptyset \in \mathcal{F}$, then there must be a finite subcollection of $\mathcal{A}$ with empty intersection, which contradicts the assumption. Therefore, $\mathcal{F}$ is a filter that contains $\mathcal{A}$. □

Proposition 4.6.3. *The collection of filters is stable under non-empty intersections.*

Proof. Let $(\mathcal{F}_i)_{i \in I}$ be a family of filters, with $I \neq \emptyset$, and let $\mathcal{G} = \bigcap_{i \in I} \mathcal{F}_i$. It is clear that $\mathcal{G}$ is upwards closed and stable under finite intersections, since all the $\mathcal{F}_i$ are filters by assumption. If $\emptyset \in \mathcal{G}$, then $\emptyset \in \mathcal{F}_i$ for *all* $i \in I$, which, since I is non-empty, implies that there exists *some* $i \in I$ with $\emptyset \in \mathcal{F}_i$, contradicting the fact that $\mathcal{F}_i$ is a filter. □

Note that the assumption of non-emptiness in Proposition 4.6.3 is necessary, since the empty intersection of filters is the collection of all subsets, which is not a filter.

Proposition 4.6.4. *Let $\mathcal{A}$ be a collection of subsets of a set X with the finite intersection property. Then there exists a minimum filter containing $\mathcal{A}$.*

Proof. By Proposition 4.6.2, there exists at least a filter containing $\mathcal{A}$. Therefore, the intersection $\mathcal{F}$ of all filters containing $\mathcal{A}$ is a filter by Proposition 4.6.3, and it is clear that $\mathcal{F}$ is the minimum such filter. □

We will say that $\mathcal{F}$ is *generated* by $\mathcal{A}$ if $\mathcal{F}$ is the minimum filter containing $\mathcal{A}$.

Let us now turn our attention again to filters in a topological spaces.

Proposition 4.6.5. *Let $\mathcal{F}$ be a filter on a topological space X, and $\ell \in X$ a point. Then $\mathcal{F}$ is contained in a filter that converges to ℓ if and only if, for all $A \in \mathcal{F}$ and all neighbourhoods U of ℓ, we have that $A \cap U \neq \emptyset$.*

Proof. The condition is necessary, for if $\mathcal{F} \subseteq \mathcal{G}$, and $\mathcal{G} \to \ell$, then given $A \in \mathcal{F}$ and $U \in \mathcal{N}(\ell)$, one has $A \in \mathcal{G}$ and $U \in \mathcal{G}$, hence $A \cap U \neq \emptyset$.

Conversely, assume that every element of $\mathcal{F}$ meets every neighbourhood of ℓ, and let $\mathcal{A} = \mathcal{F} \cup \mathcal{N}(\ell)$. It follows that $\mathcal{A}$ has the finite intersection property, therefore it is contained in a filter $\mathcal{G}$ by Proposition 4.6.2. Then $\mathcal{F} \subseteq \mathcal{G}$ by construction, and also $\mathcal{G} \to \ell$, since it contains all the neighbourhoods of ℓ. $\square$

5 Constructing topological spaces

So far, we have developed some basic theory of topological spaces and continuous functions, but we have not seen many examples. In this section, we will introduce a number of constructions that will allow us to define new spaces from existing ones, greatly expanding our ability to exhibit examples. All of the constructions will be based on initial and final topologies, as defined in section 4.5.

5.1 Subspaces

As a special - but very important - case of initial topology, let us consider a subset A of a topological space X, together with its inclusion function $i : A \to X$. The initial topology induced by i is called the *subspace topology* on A.

Through the subspace topology, we can regard every subset of a topological space X as a topological space in its own right. We will often implicitly regard any such subset as a topological space, and often we will speak of a *subspace* of X when doing so.

Proposition 5.1.1. *Let A be a subspace of a topological space X. A set $S \subseteq A$ is open (resp. closed) in (the subspace topology of) A if and only if there exists an open (resp. closed) set $\widetilde{S} \subseteq X$ such that $\widetilde{S} \cap A = S$. In particular, if A is itself open in X, then the open sets of A are exactly those open sets of X that are contained in A.*

Proof. Let $i : A \to X$ be the inclusion function. If $\widetilde{S}$ is a subset of X, then $i^{-1}(\widetilde{S}) = \widetilde{S} \cap A$. It follows from Proposition 4.5.6 that the initial topology induced by i is generated by the collection τ of subsets of the form $U \cap A$ with U ranging over the open subsets of X. Furthermore, it is immediate to verify that τ is a topology, hence it coincides with the topology generated by it.

The second assertion follows immediately from the fact that open sets in X are stable under binary intersections. $\square$

It is important to note that the open sets of a subspace are *not* necessarily open in the containing topological space. As an extreme example, a subspace $A \subseteq X$ need not be open when regarded as a subset of X, but it is clearly always open as a subset of A itself (since topologies always contain the whole space).

Exercise 27. Let $\mathbb{Q}$ be the space of rational numbers regarded as a subspace of $\mathbb{R}$. Describe the open sets of $\mathbb{Q}$.

Exercise 28. Let X be a metric space, and $A \subseteq X$ any subset. Regard A as a metric space by equipping it with the restriction of the metric of X. Show that the topology on A induced by this metric is equal to the subspace topology.

Exercise 29. Prove that any two open (resp. closed) intervals of $\mathbb{R}$ are homeomorphic. Prove that any open interval of $\mathbb{R}$ is homeomorphic to $\mathbb{R}$ itself.

It may be sometimes difficult to define a continuous function on a whole space X, often because it is hard to find a definition that is valid everywhere on X. In such cases, it is convenient to be able to "patch" together multiple functions valid on subspaces.

Definition 5.1.2. A *cover* of a topological space X is a family of subsets of X whose union is X. An *open cover* (resp. *closed cover*) is a cover consisting of open (resp. closed) sets. We say that a cover is *finite* if the indexing set is finite.

Proposition 5.1.3. *Let $(U_i)_{i \in I}$ be an open cover of a topological space X, and let $f_i : U_i \to Y$ be continuous functions to a topological space Y. If for all pair of indices $i, j \in I$ we have that $f_i | U_i \cap U_j = f_j | U_i \cap U_j$, then there exists a unique $g : X \to Y$ such that $g | U_i = f_i$ for all $i \in I$.*

Proof. It is clear that there exists a unique function $g : X \to Y$ such that $g | U_i = f_i$ for all $i \in I$, so it remains to show that g is continuous. If V is an open set in Y, then clearly:

$$g^{-1}(V) = \bigcup_{i \in I} f_i^{-1}(V).$$

By continuity of f_i, the set $f_i^{-1}(V)$ is open in U_i. Since U_i itself is open in X, it follows that all the $f_i^{-1}(V)$ are open in X, hence so is their union. $\square$

Note that Proposition 5.1.3 cannot be generalised to arbitrary covers. An extreme counterexample is a cover consisting of singleton subsets $\{x\}$ for all $x \in X$. Since singleton spaces are discrete, a family of functions over this cover defines an *arbitrary* function on X, which clearly cannot be shown to be continuous.

However, we have a limited version of Proposition 5.1.3 for *finite* closed covers.

Proposition 5.1.4. *Let $(C_i)_{i \in I}$ be a finite closed cover of a topological space X, and let $f_i : C_i \to Y$ be continuous functions to a topological space Y. If for all pair of indices $i, j \in I$ we have that $f_i | C_i \cap C_j = f_j | C_i \cap C_j$, then there exists a unique $g : X \to Y$ such that $g | C_i = f_i$ for all $i \in I$.*

The proof of Proposition 5.1.4 is entirely analogous to that of Proposition 5.1.3, and is left as an exercise for the reader.

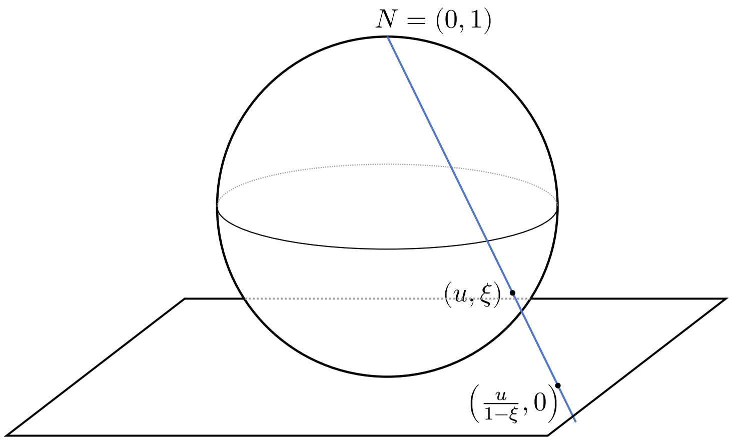

Figure 1: The stereographic projection

5.1.1 Examples

Many common topological spaces are naturally defined as subspaces. One of the most important examples are the *spheres*.

Definition 5.1.5. Let n be a natural number. The *n-sphere* S^n is the subspace of $\mathbb{R}^{n+1}$ consisting of all those points x such that $\|x\| = 1$.

Note that S^0 is a discrete space consisting of exactly the two points 1 and -1 in $\mathbb{R}$. More generally, if we regard $\mathbb{R}^{n+1}$ as $\mathbb{R}^n \times \mathbb{R}$, we can distinguish the two points $N = (0, 1)$ and $S = (0, -1)$ on S^n, called respectively *north pole* and *south pole*. The one-dimensional sphere S^1 is referred to as the *circle*.

Exercise 30. Show that S^n is a closed subset of $\mathbb{R}^{n+1}$.

One of the reasons why spheres are important in topology is that they can be thought of as the result of "adding one point" to a Euclidean space. We will make this intuition completely precise later when discussing compactification, but for now, we can partially justify it by exhibiting a homeomorphism between a "punctured" sphere and Euclidean space.

Definition 5.1.6. The *stereographic projection* $\phi : S^n \setminus \{N\} \to \mathbb{R}^n$ is the map defined by:

$$\phi(u, \xi) = \frac{u}{1 - \xi},$$

where again $\mathbb{R}^{n+1}$ is regarded as $\mathbb{R}^n \times \mathbb{R}$.

Since $1 - \xi$ only vanishes at the north pole, ϕ is a well-defined continuous function. The geometric idea behind ϕ is very simple (fig. 1): fix a point (u, ξ) on the sphere, and draw a line starting from the north pole $(0, 1)$. Such a line can be parameterised as $(tu, 1 - t + t\xi)$, so it will intersect the hyperplane of points with the last coordinate equal to 0 exactly when $t = 1/(1 - \xi)$, and the intersection point is therefore $\phi(u, \xi) = tu = u/(1 - \xi)$. In other words, we are projecting the sphere onto a hyperplane from the north pole. This explains why we cannot include the north pole itself in the projection.

Exercise 31. Show that the stereographic projection $\phi : S^n \setminus \{N\} \to \mathbb{R}^n$ is a homeomorphism by exhibiting a continuous inverse. [*Hint: reverse the projection by starting with a point on the hyperplane, drawing a line to the north pole, and intersecting it with the sphere*].

Another important family of spaces in topology is given by *disks*.

Definition 5.1.7. The n-disk D^n is the subspace of $\mathbb{R}^n$ consisting of all those points x such that $\|x\| \leq 1$.

Exercise 32. Show that D^n is a closed subset of $\mathbb{R}^n$, and that the boundary of D^{n+1} is S^n.

Exercise 33. Let $\pi : \mathbb{R}^{n+1} \to \mathbb{R}^n$ be the projection function that "forgets" the last coordinate. Show that π restricts to a homeomorphism between the *upper hemisphere* S^n_+ (i.e. the subspace of S^n consisting of the points $(u, \xi) \in \mathbb{R}^{n+1} \times \mathbb{R}$ where $\xi \geq 0$) and the disk D^n.

Exercise 34. Construct a continuous function $D^n \to S^n$ that maps S^{n-1} to a point and is injective in the interior of D^n.

5.2 Products

Let us recall that, given a family of sets $(X_i)_{i \in I}$, the *product* of the family is the set

$$P = \prod_{i \in I} X_i,$$

consisting of families $x = (x_i)_{i \in I}$, with $x_i \in X_i$. If we fix an index $i \in I$, we can define a *projection function* $\pi_i : P \to X_j$ that maps a family x to its value x_j at the index i.

Products of sets satisfy the following universal property, whose proof is left as an exercise for the reader.

Proposition 5.2.1. *If $(X_i)_{i \in I}$ is a family of sets, Y is any set, and $f_i : Y \to X_i$ is a family of functions, there exists a unique function*

$$g : Y \to \prod_{i \in I} X_i$$

such that $\pi_i \circ g = f_i$ for all $i \in I$.

We will sometimes denote the function g given by Proposition 5.2.1 as $\langle f_i \rangle_{i \in I}$.

We now wish to obtain a similar construction for topological spaces. Namely, given a family $(X_i)_{i \in I}$ of topological spaces, we wish to define a *product* space P, together with continuous functions $\pi_i : P \to X_i$, satisfying a universal property analogous to Proposition 5.2.1.

The solution is very simple: we build a space starting with the set-theoretic product construction, and equipping it with a topology that makes the projection functions continuous. This is made explicit by the following definition.

Definition 5.2.2. Let $(X_i)_{i \in I}$ be a family of topological spaces. The *product space* of the X_i is defined to be the set

$$P = \prod_{i \in I} X_i,$$

equipped with the initial topology induced by the projection functions $\pi_i : P \to X_i$.

To show that Definition 5.2.2 is the "correct" definition, we prove a corresponding universal property mirroring the one for products of sets.

Proposition 5.2.3. *Let $(X_i)_{i \in I}$ be a family of topological spaces. For any topological space Y, and any family of continuous functions $f_i : Y \to X_i$, there exists a unique continuous function $g : Y \to \prod_{i \in I} X_i$ such that $\pi_i \circ g = f_i$.*

Proof. Let $P = \prod_{i \in I} X_i$. Thanks to Proposition 5.2.1, we know that there exists a unique function $g : Y \to P$, so the uniqueness part is clearly satisfied. As for existence, it suffices to show that the function g is continuous. By Proposition 4.5.9, g is continuous as long as $\pi_i \circ g$ is continuous for all $i \in I$. But $\pi_i \circ g = f_i$, which is continuous by assumption. □

As a degenerate example of product, consider the empty family of spaces. Since there is exactly one empty family, the product of the empty family is a space consisting of a single point. We will refer to this space as the *terminal space*, and denote it 1. Its topology is both the discrete and the chaotic topology on the underlying set.

One of the reason why universal properties are interesting and useful is that constructions satisfying universal properties are usually uniquely determined up to the appropriate notion of "isomorphism". For us, this notion is given by homeomorphism (Definition 3.2.3). More precisely, let us say that a space P, together with continuous functions $p_i : P \to X_i$ *satisfies the universal property of the product of the* X_i if for all spaces Y, and all continuous functions $f : Y \to X_i$, there exists a unique continuous function $g : Y \to P$ such that $p_i \circ g = f_i$.

Exercise 35. Without referring to the construction of the product above, prove that given two spaces P and P', with continuous functions $p_i : P \to X_i$ and $p'_i : P \to X_i$, both satisfying the universal property of the product of the X_i, there exists a unique homeomorphism $h : P \to P'$ such that $p'_i \circ h = p_i$.

Although we have defined the product of a family of spaces using an explicit construction, we are also allowed to think of the universal property itself as an *indirect* definition of the product, in the sense that any space P equipped with continuous functions $p_i : P \to X_i$ satisfying the universal property of the product given above could equally well be regarded as *a product* of the X_i, regardless of how *the product* is defined. From this point of view, the construction given in Definition 5.2.2, together with the proof of Proposition 5.2.3, could be interpreted as a proof that *there exists a product* of any family of spaces.

Because of the considerations above, we will often say that a certain space P *is the product* of a family $(X_i)_{i\in I}$ when it satisfies the corresponding universal property. Thanks to Exercise 35, this is equivalent to saying that P is homeomorphic to the product of the X_i.

Exercise 36. Show that $\mathbb{R}^n$ is the product of n copies of $\mathbb{R}$. Show that $\mathbb{C}$ is the product of two copies of $\mathbb{R}$. Deduce that $\mathbb{C}^n \cong \mathbb{R}^{2n}$.

Exercise 37. Prove that $D^n \times D^m \cong D^{n+m}$. [*Hint: define a function* $\phi : D^n \times D^m \to D^{n+m}$ *as* $\phi(u,v) = (u,v)\sqrt{\|u\|^2 + \|v\|^2}/\mathsf{max}(\|u\|, \|v\|)$ *and prove that* ϕ *is continuous in* $(0,0)$.]

Exercise 38. Construct a homeomorphism between $\mathbb{R}^{n+1} \setminus \{0\}$ and $S^n \times \mathbb{R}$ [*Hint: polar coordinates*].

The topology on the product of two topological spaces can also be characterised directly in terms of bases.

Proposition 5.2.4. *Let* X, Y *be topological spaces. Let* $\mathcal{B}$ *be the collection of subsets of* $X \times Y$ *of the form* $U \times V$*, where* U *is an open set in* X *and* V *an open set in* Y*. Then* $\mathcal{B}$ *is a base for the topology of* $X \times Y$*.*

Proof. First observe that $\mathcal{B}$ is stable under finite intersections, so all we need to show is that it generates the topology on $X \times Y$.

Let $\pi_0 : X \times Y \to X$ and $\pi_1 : X \times Y \to Y$ denote the two projections. If U is open in X, then by definition of product space $\pi_0^{-1}(U)$ is an open set of $X \times Y$. But $\pi_0^{-1}(U) = U \times Y \in \mathcal{B}$, and similarly $\pi_1^{-1}(V) = X \times V$, hence the topology generated by $\mathcal{B}$ is finer than the product topology. On the other hand, $U \times V = \pi_0^{-1}(U) \cap \pi_1^{-1}(V)$, hence the elements of $\mathcal{B}$ are open in the product topology. It follows that $\mathcal{B}$ generates exactly the product topology. $\square$

Note that Proposition 5.2.4 cannot directly be generalised to arbitrary products. In particular, given a family of topological spaces $(X_i)_{i\in I}$, a product of $\prod_{i\in I} U_i$ where each U_i is an open set in X_i is not guaranteed to be open in the product topology.

Exercise 39. Let D be a discrete space. Show that $D^{\mathbb{N}} := \prod_{n\in\mathbb{N}} D$ is discrete if and only if D is the terminal space. Deduce that an infinite product of open sets is not necessarily open in the product topology.

Note that the product of a family of the form $(X)_{i\in I}$, i.e. consisting of a copy of the same space X for each element of a set I, can be thought of as the set of functions $I \to X$, which is also denoted X^I.

Exercise 40. Prove that a filter $\mathcal{F}$ converges to $x \in \prod_{i \in I} X_i$, if and only if it does so *pointwise*, i.e. if for all $i \in I$, the image filter $(\pi_i)_*(\mathcal{F})$ converges to x_i.

5.3 Coproducts

While completely dual to products, coproducts of sets might be unfamiliar to the reader, hence we review them here. Our goal, as in the previous subsections, will then be to establish a construction for topological spaces that satisfies an analogous universal property.

The *coproduct* of sets is another name for their *disjoint union*. Given a family of sets $(X_i)_{i \in I}$, if the X_i happen to be pairwise disjoint (i.e. $X_i \cap X_j = \emptyset$ for all $i, j \in I$), then their coproduct could be defined as simply their union $\bigcup_{i \in I} X_i$.

In the general case, we need to somehow force all the X_i to be disjoint before taking their union. The "trick" is to pair every element of an X_i with the corresponding index i, i.e. to replace X_i with $X_i' := \{i\} \times X_i$. This way, even if an element x belongs to more than one set, say $x \in X_i$ and $x \in X_j$, the pairs (i, x) and (j, x) are still distinct, hence the modified sets X_i' are *always* disjoint. The construction is summarised by the following definition.

Definition 5.3.1. Let $(X_i)_{i \in I}$ be a family of sets. The *coproduct* of the X_i, denoted

$$\coprod_{i \in I} X_i,$$

is the set of pairs (i, x), where $i \in I$ and $x \in X_i$.

There are obvious maps $\iota_i : X_i \to \coprod_{i \in I} X_i$ defined by $\iota_i(x) = (i, x)$. We will refer to the ι_i as *coproduct injections*. A coproduct injection ι_i is indeed injective, hence it allows us to regard X_i as a subset of the coproduct, with which we will often implicitly identify it.

Exercise 41. Let $f : A \to B$ be a function. For all $b \in B$, let F_b be the *fibre* of f over b, i.e. the inverse image of the singleton $\{b\}$. Establish a bijection between A and $\coprod_{b \in B} F_b$.

The reason for introducing the term "coproduct" instead of simply referring to this construction as a disjoint union is twofold:

1. it emphasises its dual nature to products (as demonstrated below)
2. it generalises to other *categories* where coproducts are not disjoint unions.

The duality with products manifests itself in the universal property for coproducts, whose statement can be obtained purely *formally* by reversing all the arrows in Proposition 5.2.1.

Proposition 5.3.2. *If $(X_i)_{i \in I}$ is a family of sets, Y is any set, and $f_i : X_i \to Y$ is a family of functions, there exists a unique function*

$$g : \coprod_{i \in I} X_i \to Y$$

such that $g \circ \iota_i = f_i$ for all $i \in I$.

Proof. If such a g exists, then $g(i, x) = g(\iota_i(x)) = f_i(x)$, hence g is unique. On the other hand, we can take this equation as the definition of g, which shows existence as well. □

Exercise 42. Show that the empty set is the coproduct of the empty family of sets.

Just like in the case of products, we will often consider *binary* coproducts, denoted $X + Y$ for the family consisting of the two spaces X and Y. The elements of $X + Y$ are all those pairs $(0, x)$ and $(1, y)$, with $x \in X$ and $y \in Y$. The reason for using the symbol $+$ to denote binary coproducts is explained by the following exercise.

Exercise 43. If X and Y are finite sets of cardinalities n and m respectively, show that $X + Y$ has cardinality $n + m$.

Thanks to the duality between coproduct and products, there is an obvious candidate for the construction of the coproduct of topological spaces, which we now define.

Definition 5.3.3. Let $(X_i)_{i \in I}$ be a family of topological spaces. The *coproduct space* of the X_i is the coproduct of the X_i as sets:

$$C = \coprod_{i \in I} X_i,$$

equipped with the final topology induced by the coproduct injections $\iota_i : X_i \to C$.

By Proposition 4.5.7, a set U is open in C if and only if $U \cap X_i$ is open in X_i for all $i \in I$. Here we are identifying X_i with the corresponding subset $\iota_i(X_i)$ of C. In particular, every X_i is always open when regarded as a subset of C.

Exercise 44. Prove that all coproduct injections ι_i are homeomorphisms with their image, i.e. that the induced maps $X_i \to \iota_i(X_i)$ are homeomorphisms.

Adapting the universal property of coproducts to topological spaces is straightforward:

Proposition 5.3.4. *Let $(X_i)_{i \in I}$ be a family of topological spaces. For any topological space Y, and any family of continuous functions $f_i : X_i \to Y$, there exists a unique continuous function $g : \coprod_{i \in I} X_i \to Y$ such that $g \circ \iota_i = f_i$.*

Proof. We know from Proposition 5.3.2 that there exists a unique such function g, so all we need to show is that g is continuous, but that follows directly from Proposition 5.1.3 and the fact that every X_i is open in the coproduct space. □

Similar considerations to those about products apply here. In particular, we can define what it means for a space to have the *universal property of the coproduct* of a family of spaces, and any two such spaces are homeomorphic with a unique homeomorphism that is compatible with the coproduct injections. The details can be obtained by formally *dualising* those for products, and are left to the reader to work out explicitly.

Exercise 45. Let U, V be disjoint open sets of a topological space X. Show that the subspace $U \cup V$ is homeomorphic to the coproduct $U + V$. Does the statement also hold for closed sets? Find a counterexample to the statement when only U is assumed to be open.

Coproducts of spaces have a direct geometric interpretation: the coproduct of two spaces X and Y is what we get if we consider both of them at the same time as a single space. A "picture" of $X + Y$ simply looks like a picture of X next to a picture of Y.

Exercise 46. Show that a discrete space X is the coproduct of the family $(1)_{x \in X}$, i.e. a coproduct of copies of the terminal space 1, one for each point in x.

Exercise 47. Let $C = \coprod_{i \in I} X$ be the coproduct of a *constant* family of spaces, i.e. of a family consisting of copies of a single space, one for each element of the indexing set I. Show that $C \cong I \times X$, where I is regarded as a discrete topological space.

5.4 Quotients

If $\sim$ is an equivalence relation on a set X, recall that the *quotient* $X/{\sim}$ of X by $\sim$ is defined to be the set of equivalence classes of $\sim$. There is a function $\pi : X \to X/{\sim}$ that maps every element of x to its equivalence class. Quotients also satisfy a universal property, whose proof is left to the reader.

Proposition 5.4.1. *Let $\sim$ be an equivalence relation on a set X, and let $f : X \to Y$ be a function such that $x \sim x'$ implies $f(x) = f(x')$. Then there exists a unique function $g : X/{\sim} \to Y$ such that $g \circ \pi = f$.*

The condition on the function f in Proposition 5.4.1 can be expressed by saying that f *is compatible with the equivalence relation* $\sim$.

As usual, we are looking to adapt the quotient construction to topological spaces. Again, the definition is straightforward.

Definition 5.4.2. Let X be a topological space and $\sim$ an equivalence relation on X. The *quotient space* $X/{\sim}$ is defined to be the quotient of X by $\sim$ as a set, equipped with the final topology induced by the single function $\pi : X \to X/{\sim}$.

In other words, a subset U is open in $X/{\sim}$ if and only if its inverse image along π is open in X. In particular π is continuous. The fact that the topology on $X/{\sim}$ is defined as a final one makes it easy to show the corresponding universal property.

Proposition 5.4.3. *Let X be a topological space and $\sim$ an equivalence relation on X. Let $f : X \to Y$ be a continuous function that is compatible with $\sim$. Then* *there exists a unique continuous function $g : X/{\sim} \to Y$ such that $g \circ \pi = f$.*

Proof. We know from Proposition 5.4.1 that there is a unique such function g, so we only need to show that g is continuous, and that follows immediately from Proposition 4.5.10 and the fact that f is continuous. $\square$

5.4.1 Example: gluing disks

Quotients are extremely important in topology, as they allow us to *glue* together spaces. To illustrate the idea, we will work through the example of gluing two disks by their boundaries.

Consider two copies of the disk D^n, i.e. the coproduct space $X = D^n + D^n$. We can picture X as two parallel disks in $\mathbb{R}^{n+1}$, lying one above the other (this is easier to visualise when $n \leq 2$).

Now we want to glue together the boundaries of the two disks, by stitching together pairs of points on the respective boundaries that lie directly above each other (i.e. when their first n coordinates are equal).

To make this intuition precise, first recall that by Exercise 47 X can be regarded as the product $2 \times X$, where $2 = \{0, 1\}$ is a discrete space. We define an equivalence relation $\sim$ on X by declaring that $(i, x) \sim (i', x')$ precisely when $x = x'$ and at least one of the following two conditions holds: $i = i'$ or $x \in \partial D^n$.

In other words, we are identifying two points when they are the same point, or when their second components are equal. The first condition is needed to make sure that $\sim$ is an equivalence relation, while the second formalises the idea that we are prescribing that corresponding points in the two boundaries should be stitched together.

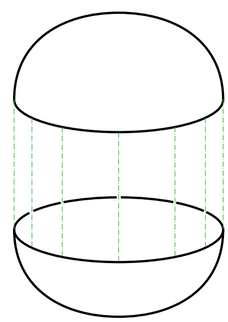

Figure 2: Two disks being glued

Exercise 48. Show that $\sim$ defined above is indeed an equivalence relation on X.

We can now take the quotient space $X/{\sim}$. We can visualise this space as the result of actually applying the stitches prescribed by $\sim$. If we make the stitches airtight, and inflate the inside, what we get is a sphere. Of course, this is only a vague intuition, but thankfully, we can apply the theory developed so far to make it completely precise.

Proposition 5.4.4. *The quotient space $X/{\sim}$ is homeomorphic to S^n.*

Proof. Regard S^n as a subset of $\mathbb{R} \times \mathbb{R}^n$, and denote by S^n_+ (resp. S^n_-) the set of points (ξ, u) such that $\xi \geq 0$ (resp $\xi \leq 0$). It is clear that S^n_+ and S^n_- are closed subsets of S^n, and that their union is S^n.

Now define a map $\phi_+ : S^n_+ \to X/{\sim}$ as $\phi_+(\xi, u) = \pi(1, u)$, where $\pi : X \to X/{\sim}$ is the projection into the quotient. Similarly, define $\phi_- : S^n_- \to X/{\sim}$ by $\phi_-(\xi, u) =$

$\pi(0,u)$. It is clear that ϕ_+ and ϕ_- are continuous.

The intersection $E = S^n_+ \cap S^n_-$ consists of points of the form $(0,u)$, where $\|u\| = 1$, hence such a u belongs to the boundary of the disk D^n. It follows that $\pi(0,u) = \pi(1,u)$ by the definition of $\sim$, hence ϕ_+ and ϕ_- coincide on E.

Therefore, we can apply Proposition 5.1.4 to obtain a continuous function $\phi : S^n \to X/\sim$ such that $\phi(\xi,u) = \pi(i,u)$, where $i = 1$ if and only if $\xi \geq 0$.

In the other direction, let $\widetilde{\psi} : X \to S^n$ be given by:

$$\widetilde{\psi}(i,x) = ((-1)^i\sqrt{1-\|x\|^2}, x).$$

The function $\widetilde{\psi}$ is clearly continuous, so to obtain a continuous function $X/\sim \to S^n$ all we need to show is that $\widetilde{\psi}$ is compatible with $\sim$. So suppose $(i,x) \sim (i',x')$. It is of course enough to consider the case where $x = x'$ and $x \in \partial D^n$. Then $\|x\| = 1$, hence $\widetilde{\psi}(i,x) = (0,x) = (0,x') = \widetilde{\psi}(i',x')$, as required.

To verify that $\psi \circ \phi = \mathsf{id}$, it is enough to show that $\psi(\phi_+(\xi,u)) = (\xi,u)$ for all $(\xi,u) \in S^n_+$, and similarly for ϕ_-. $\psi(\phi_+(\xi,u)) = \psi(\pi(1,u)) = \widetilde{\psi}(1,u) = (\sqrt{1-\|u\|^2}, u)$, and the required equation follows from the fact that $\xi^2 + \|u\|^2 = 1$ and that $\xi \geq 0$. The verification for ϕ_- is entirely analogous.

Finally, let us verify that $\phi \circ \psi = \mathsf{id}$. First note that $\phi(\widetilde{\psi}(x)) = \phi((-1)^i\sqrt{1-\|x\|^2}, x) = \pi(i,x)$, hence $\phi \circ \psi \circ \pi = \pi$. But of course, also $\mathsf{id} \circ \pi = \pi$, hence by the uniqueness part of Proposition 5.4.3, we get that $\phi \circ \psi = \mathsf{id}$, which completes the proof. □

Note how the proof of Proposition 5.4.4 constructs both directions of a homeomorphisms. We will see later (in Corollary 8.1.8) that in many cases, including this one, it is enough to construct one of the two maps and show that it is a bijection. This allows to produce relatively simple proofs of homeomorphisms in certain cases where an inverse function cannot defined explicitly, hence its continuity is hard to prove.

We will often refer to Corollary 8.1.8 when constructing homeomorphisms. Since Corollary 8.1.8 refers to notions that will not have been introduced yet, we ask the reader to temporarily take it on faith that the statement can indeed by applied, and encourage them to review these proofs after Corollary 8.1.8 itself has been proved.

5.4.2 Collapsing subspaces

If X is a topological space, and A a subset of X, define an equivalence relation $\sim$ on X by declaring that $x \sim x'$ if and only if $x = x'$ or both x and x' belong to A. In other words, $\sim$ is the minimal equivalence relation such that $a \sim a'$ for all $a, a' \in A$.

The quotient $X/\sim$ is denoted X/A, and we say that X/A is obtained from X by *collapsing* A. The idea is of course that A is reduced to a single point in the quotient, and everything else is left unchanged.

Exercise 49. Let X be the space obtain from $\mathbb{R}$ by collapsing the subset $\{-1, 1\}$. Show that X is homeomorphic to the set C of points (x, y) in $\mathbb{R}^2$ satisfying the following polynomial equation:

$$y^2 = x(x-1)^2.$$

[*Hint: consider the map* $\tilde{\phi} : \mathbb{R} \to C$ *defined by* $\phi(t) = (t^2, t^2(t^2-1)^2)$, *show that it defines a map on* X, *then find an inverse for that*]

An important example of space that is constructed by collapsing is the *cone*. In the following, $[a, b]$ will denote the closed interval of the real line with endpoints a and b.

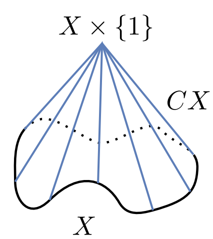

Figure 3: Cone on X

Definition 5.4.5. Let X a topological space. The *cone* on X, denoted CX, is the space obtained from $X \times [0, 1]$ by collapsing $X \times \{1\}$.

Exercise 50. If X is a subspace of $\mathbb{R}^n$, prove that CX is homeomorphic to the union of all those segments in $\mathbb{R}^{n+1} = \mathbb{R}^n \times \mathbb{R}$ obtained by joining the point $(0, 1)$ to all the points $(x, 0)$, where x ranges over the points of X.

For any space X, there is a canonical injection $j : X \to CX$ that maps x to $\pi(0, x)$.

Proposition 5.4.6. *The map* $j : X \to CX$ *is a homeomorphism with its image.*

Proof. Clearly j is continuous and bijective. To show that its inverse is also continuous, we prove that j is closed. So, let $C \subseteq X$ be closed. Since the only collapsed points in $X \times [0, 1]$ have second coordinate equal to 1, we have that $\pi^{-1}(j(C)) = C \times \{0\}$, which is closed. Therefore $j(C)$ is closed by Corollary 4.5.8. $\square$

Another basic construction on spaces is the *suspension*:

Definition 5.4.7. Let X be a topological space. The *suspension* of X, denoted SX, is the quotient space of $X \times [-1, 1]$ by the relation $\sim$ defined as follows: $(x, t) \sim (x', t')$ if $t = t'$ and at least one of the following conditions hold: $x = x'$ or $|t| = 1$.

Exercise 51. Show that SX is homeomorphic to the space obtained from $X \times [-1, 1]$ by first collapsing $X \times \{1\}$, and then collapsing $\pi(X \times \{-1\})$, where π denotes the projection from $X \times [-1, 1]$ into the quotient.

The name *suspension* suggests a picture where X is "suspended" mid-air by wires connecting it to two points on the two sides of X. This can be made precise (at least when X is a subspace of $\mathbb{R}^n$).

Exercise 52. If X is a subspace of $\mathbb{R}^n$, prove that CX is homeomorphic to the union of all those segments in $\mathbb{R}^{n+1} = \mathbb{R}^n \times \mathbb{R}$ obtained by joining the points $(0, \pm 1)$ to the points $(x, 0)$, where x ranges over the points of X.

The following proposition establishes a link between cone and suspension.

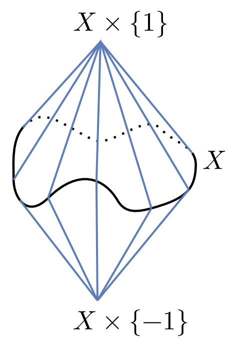

Figure 4: Suspension of X

Proposition 5.4.8. *Let X be a topological space. The suspension SX is homeomorphic to the space Y obtained by gluing two copies of CX along X. More precisely, Y is the quotient of $CX + CX$ by the minimal equivalence relation $\approx$ such that $\iota_0(j(x)) \approx \iota_1(j(x))$ for all $x \in X$.*

Proof. Thanks to Proposition 5.4.3, we can define a continuous function $f_0 : CX \to SX$ such that $f_0(\pi(x,t)) = \pi(x,t)$. Here we are using π to denote both the projection $X \times [0,1] \to CX$ and the projection $X \times [-1,1] \to X$. This should not cause any confusion.

Similarly, we define $f_1 : CX \to SX$ via $f_1(\pi(x,t)) = \pi(x,-t)$. Together, they define a function $\tilde{\phi} : CX + CX \to SX$, and it is immediate to verify that $\tilde{\phi}$ is compatibile with $\approx$, hence it determines a continuous map $\phi : Y \to SX$.

As for the the other direction, let $S_+X = \pi(X \times [0,1]) \subseteq SX$. Since $\pi^{-1}(S_+X) = X \times [0,1]$, which is closed in $X \times [-1,1]$, it follows that $\pi^{-1}(S_+X)$ is closed. Similarly $S\_X = \pi(X \times [-1,0])$ is a closed subset of SX. Define $\psi_+ : S_+X \to Y$ by $\psi_+(\pi(x,t)) = \pi(\iota_0(x,t))$ and $\psi\_ : S\_X \to Y$ by $\psi\_(\pi(x,t)) = \pi(\iota_1(x,t))$. It is easily checked that ψ_+ and $\psi\_$ coincide on $S_+X \cap S\_X$, hence they define a continuous function $\psi : SX \to Y$.

Verifying that ψ is the inverse of ϕ is now a simple matter of expanding the definitions and using the uniqueness part of the universal properties of the various quotients. $\square$

Exercise 53. Show that $D^{n+1} \cong CS^n$. Deduce from Proposition 5.4.4 and Proposition 5.4.8 that $SS^n \cong S^{n+1}$.

Exercise 54. Show that $CX/X \cong SX$. Deduce that $S^n \cong D^n/\partial D^n$.

5.5 Example: the torus

There are often many ways to define a topological space: we can regard it as a subspace of $\mathbb{R}^n$, or construct it indirectly using previously defined spaces. An example of space which can be obtained in several equivalent ways is the torus. The simplest definition, and often the most convenient to work with, is the following.

Definition 5.5.1. The *torus* is the product $S^1 \times S^1$ of two copies of the circle.

In fig. 5, the two circles labelled a and b denote the "axes" of the product, say $S^1 \times \{N\}$ and $\{N\} \times S^1$, where we have arbitrarily chosen N as a *base point* for both copies of the circle. The point (N, N) is the intersection of these two circles.

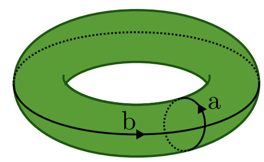

Figure 5: The *torus*

Another way to look at the torus is as the quotient of a square. More precisely, consider the equivalence relation $\sim$ on $[0,1] \times [0,1]$ generated by the following clauses:

$$(s, 0) \sim (s, 1),$$
$$(0, t) \sim (1, t)$$

and define $T := ([0,1] \times [0,1])/\sim$ to be the corresponding quotient space. We will show that this is an equivalent definition of the torus. In the following, we will often regard S^1 as a subspace of $\mathbb{C}$ through the usual identification $\mathbb{R}^2 \cong \mathbb{C}$. Under that identification, the circle can be regarded as the set of complex numbers of absolute value 1.

The equivalence relation $\sim$ is illustrated in fig. 6: the two arrows labelled a mark segments that end up being identified in the quotient. The direction of the arrows is the same, meaning that the equivalence relation identifies points on the line having the same y coordinate. Similarly for the arrows labelled b. Taking the quotient can be visualised as "folding" the square so as to make the top side of the square coincide with the bottom one, and then "bending" the resulting cylinder so as to join one of its boundary circles with the other one.

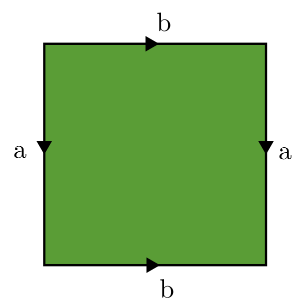

Figure 6: The torus as a quotient of $[0,1] \times [0,1]$

Proposition 5.5.2. *The quotient space T defined above is homeomorphic to the torus.*

Proof. Define a function $\tilde{\phi} : [0,1] \times [0,1] \to S^1 \times S^1$ by $\tilde{\phi}(s,t) = (e^{2\pi is}, e^{2\pi it})$. Since the exponential function $\mathbb{C} \to \mathbb{C}$ is continuous, it follows that $\tilde{\phi}$ is continuous. Surjectivity of $\tilde{\phi}$ follows from well-known properties of the exponential function.

Next, we can show that $\tilde{\phi}$ is compatible with the equivalence relation $\sim$, simply by applying $\tilde{\phi}$ to pairs of elements identified by $\sim$, and showing that they are mapped to the same element. Indeed, $\tilde{\phi}(0,t) = (1, e^{2\pi it}) = (e^{2\pi i}, t) = \tilde{\phi}(1,t)$, and similarly for the other defining clause of $\sim$. Therefore, by Proposition 5.4.3, we get a continuous surjective function $\phi : T \to S^1 \times S^1$

To show that ϕ is injective, first observe that we can the equivalence relation $\sim$ directly as follows: two points $x, x' \in [0,1] \times [0,1]$ satisfy $x \sim x'$ if and only if

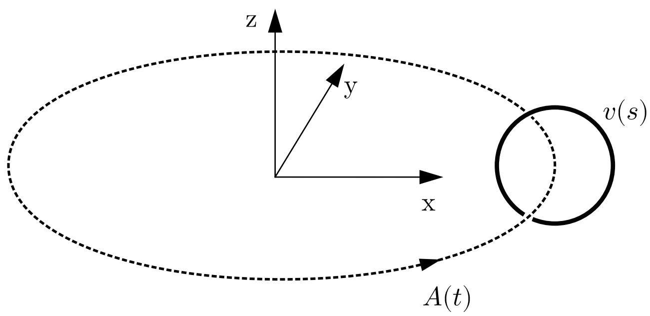

Figure 7: Parameterisation of a torus in $\mathbb{R}^3$

their coordinates differ by integers, i.e. $x - x' \in \mathbb{Z}^2$. Now, if $\phi(\pi(x)) = \phi(\pi(x'))$ for som, then it follows from the periodicity of the exponential function that $x - x' \in \mathbb{Z}^2$, hence $\pi(x) = \pi(x')$ by the above observation.

Finally, it will follow from Corollary 8.1.8 that ϕ is a homeomorphism, concluding the proof. $\square$

Note that fig. 5 is a 3-dimensional representation of the torus, while by Definition 5.5.1 the torus is naturally a subspace of $\mathbb{R}^4$. One can make the correspondence between these two representations precise, by exhibiting a homeomorphism between the torus and the subset of $\mathbb{R}^3$ depicted in fig. 5.

Exercise 55. Construct a map $\psi : S^1 \times S^1 \to \mathbb{R}^3$ using the following steps (cf. fig. 7):

- consider a circle of some radius $r > 0$ in the (x, z) plane (i.e. the plane defined by $y = 0$), and let $v_0 : S^1 \to \mathbb{R}^3$ a corresponding parameterisation;
- offset v_0 by a constant $R > r$ along the x axis, obtaining a new parameterisation $v : S^1 \to \mathbb{R}^3$;
- for $t \in S^1$, consider the matrix $A(t)$ which rotates space around the z axis by an angle $2\pi t$;
- let $\psi(s, t)$ be defined as $A(t)v(s)$.

Show that ψ is continuous and bijective. It will follow from Corollary 8.1.8 that it is a homeomorphism with its image.

Exercise 56. Let $p : \mathbb{R} \to S^1$ be the map defined by $p(t) = e^{2\pi i t}$. Show that the restriction of p to the half-open interval $[0, 1[$ is continuous and bijective, but not a homeomorphism.

5.6 Example: real projective spaces

Projective spaces are of fundamental importance in algebraic topology and algebraic geometry. Here we give their definition in the real case, and prove a few equivalent characterisations. The complex case will be examined later (section 8.6).

Definition 5.6.1. Let $\sim$ be the equivalence relation on $\mathbb{R}^{n+1} \setminus \{0\}$ where $x \sim y$ if and only if there exists $\lambda \in \mathbb{R}$ with $x = \lambda y$. The quotient space $\mathbb{R}^{n+1}/{\sim}$ is called the *n-dimensional real projective space*, and is denoted $\mathbb{R}P^n$.

Note that the equivalence classes for the relation $\sim$ on $\mathbb{R}^{n+1} \setminus \{0\}$ are simply lines through the origin (minus the origin itself, of course). It follows that there is a bijection between $\mathbb{R}P^n$ and the set of one-dimensional linear subspaces of $\mathbb{R}^{n+1}$. This observation can be used to equip the latter set with a topology, giving a first example of a *Grassmannian*, i.e. a space whose points are linear subspaces of some vector space.

The projection function $\pi : \mathbb{R}^{n+1} \setminus \{0\} \to \mathbb{R}P^n$ allows us to write points in $\mathbb{R}P^n$ in terms of $n+1$ coordinates, usually indexed starting from 0, which are called *homogeneous coordinates* of a point. By the definition of the equivalence relation, homogeneous coordinates are determined up to a scalar multiple, and they cannot all be zero. The point $\pi(x_0, ..., x_n)$ will also be written as $[x_0, ..., x_n]$.

Another way to think about $\mathbb{R}P^n$ is as as an n-dimensional Euclidean space with some extra points, which we picture to be "at infinity". This is made precise by the following.

Proposition 5.6.2. *Assume $n > 0$. Let U_0 be the set of points $[x_0, ..., x_n] \in \mathbb{R}P^n$ where $x_0 \neq 0$. Then U_0 is a dense open subspace of $\mathbb{R}P^n$ homeomorphic to $\mathbb{R}^n$, and the complement of U_0 is homeomorphic to $\mathbb{R}P^{n-1}$.*

Proof. Let $\widetilde{U}_0$ be the set of points $x = (x_0, ..., x_n) \in \mathbb{R}^{n+1}$ where $x_0 \neq 0$. Clearly, $\widetilde{U}_0$ is open, and it is easy to check that $\widetilde{U}_0 = \pi^{-1}(U_0)$, hence U_0 is open in $\mathbb{R}P^n$. Since $\widetilde{U}_0$ is dense in $\mathbb{R}^{n+1} \setminus \{0\}$, it follows (Exercise 14) that U_0 is dense in $\mathbb{R}P^n$.

We now construct a homeomorphism $\mathbb{R}^n \cong U_0$. Let $\phi : \mathbb{R}^n \to U_0$ be defined by by $\phi(u_1, ..., u_n) = [1, u_1, ..., u_n]$. To construct an inverse for ϕ, let $\widetilde{\psi} : \widetilde{U}_0 \to \mathbb{R}^n$ be given by $\widetilde{\psi}(x_0, x_1, ..., x_n) = (x_1/x_0, ..., x_n/x_0)$. It is clear that $\widetilde{\psi}$ is compatible with the equivalence relation $\sim$ on $\mathbb{R}^{n+1} \setminus \{0\}$, hence it determines a continuous function $\psi : U_0 \to \mathbb{R}^n$. It is now easy to verify that ϕ and ψ are inverses to each other.

Finally, let H_0 be the complement of U_0 in $\mathbb{R}P^n$, and denote by $\widetilde{H}_0$ its inverse image along π. Consider the homeomorphism $\alpha : \mathbb{R}^n \setminus \{0\} \to \widetilde{H}_0$ given by $\alpha(x_1, ..., x_n) = (0, x_1, ..., x_n)$. It is clear that both α and its inverse are compatible with the equivalence relations on both spaces, hence they determine a homeomorphism $\mathbb{R}P^{n-1} \cong H_0$ between the quotients. $\square$

Since points in a projective space are obtained by identifying whole lines in euclidean space, it makes sense to focus our attention only on those points that lie on a sphere, since the other points are clearly redundant, as they are always multiples of normalised points.

Unfortunately, this does not give a way to canonically identify a point in $\mathbb{R}P^n$ using coordinates in $\mathbb{R}^{n+1}$, because there are always *two* (antipodal) points on the sphere for any given line through the origin. Nevertheless, using only normalised points drastically reduces the "redundancy" of this representation from an infinite number of possible sets of coordinates to only two!

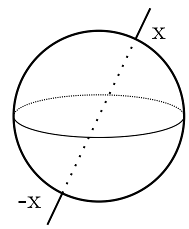

Proposition 5.6.3. *Let $\sim$ be the equivalence relation on the sphere S^n defined by $x \sim y$ if and only if $x = \pm y$. Then the quotient space $S^n/\sim$ is homeomorphic to $\mathbb{R}P^n$.*

Proof. Consider the inclusion function $\tilde{i} : S^n \to \mathbb{R}^{n+1} \setminus \{0\}$, and the function $\tilde{r} : \mathbb{R}^{n+1} \setminus \{0\} \to S^n$ defined by $\tilde{r}(x) = x/\|x\|$. Clearly $\tilde{r} \circ \tilde{i} = \mathsf{id}$, and both $\tilde{r}$ and $\tilde{i}$ are compatible with the equivalence relations on their respective domains. It follows that they determine functions $i : S^n/\sim \to \mathbb{R}P^n$ and $r : \mathbb{R}P^n \to S^n/\sim$, with $r \circ i = \mathsf{id}$. On the other hand, if $x = (x_0, ..., x_n)$ is a representative for a point in $\mathbb{R}P^n$, then $i(r[x_0, ..., x_n]) = [x_0/\|x\|, ..., x_n/\|x\|] = [x_0, ..., x_n]$, hence also $i \circ r = \mathsf{id}$, which concludes the proof. $\square$

The 1-dimensional projective space, called the *projective line*, is very special case, and it is actually homeomorphic to a space we have already encountered.

Proposition 5.6.4. $\mathbb{R}P^1 \cong S^1$.

Proof. Regard the circle as a subspace of $\mathbb{C}$, and consider the map $\tilde{f} : S^1 \to S^1$ given by $f(z) = z^2$. From the fact that $(-z)^2 = z^2$, we get that $\tilde{f}$ is compatible with the equivalence relation on S^1 defined in Proposition 5.6.3, hence it determines a continuous function $f : \mathbb{R}P^1 \to S^1$. Since every non-zero complex number has exactly two square roots, f is a bijection. The conclusion follows from Corollary 8.1.8. $\square$

What can we say about $\mathbb{R}P^2$? The argument in Proposition 5.6.4 does not generalise, and in fact we will be able to show (see Exercise 89) that $\mathbb{R}P^2$ is not homeomorphic to the 2-sphere, or any other space that we have seen so far.

6 Connectedness

6.1 Connected spaces

When we defined coproducts of spaces, we observed that a binary coproduct $X + Y$ can be thought of as the space obtained by regarding X and Y together as a single space. In particular, spaces can in general be made out of various "components", in a sense that will be made precise below.

To single out those spaces that are in some sense composed of a single "piece", we introduce the following definition.

Definition 6.1.1. A topological space X is said to be *connected* if it is non-empty, and every function $f : X \to 2$ is constant. A space that is not connected will be called *disconnected*.

Recall that $2 = \{0, 1\}$ is a discrete space with two elements. Clearly, a non-empty space is disconnected if and only if there exists a surjective function $X \to 2$, since every non-constant function to 2 is surjective.

Some authors consider the empty space to be connected. However, our convention makes certain important (for example Proposition 6.1.12 and Proposition 6.3.2) facts have cleaner statements.

Exercise 57. Show that a non-empty space X is connected if and only if for all spaces Y, Z, if $X \cong Y + Z$, then either $Y \cong \emptyset$ or $Z \cong \emptyset$.

Another equivalent formulation of connectedness is given by the following proposition.

Proposition 6.1.2. *A non-empty topological space X is connected if and only if for all disjoint open sets $U, V \subseteq X$ with $U \cup V = X$, either $U = \emptyset$ or $V = \emptyset$.*

Proof. Observe that functions $f : X \to 2$ are in bijective correspondence with pairs of disjoint open sets U, V with $U \cup V = X$. In fact, any such f determines U, V as the inverse images of the two points 0 and 1. Conversely, given U, V as above, we can define a corresponding function f that maps all the points in U to 0, all the points of V to 1, and continuity follows from Proposition 5.1.3.

Under this bijection, constant functions correspond to pairs where either U or V is empty. The statement then follows immediately. ☐

Corollary 6.1.3. *A non-empty topological space is connected if and only its only subsets that are both open and closed are the empty set and the whole space.*

A subset of a space which is both open and closed is sometimes called a *clopen* set. So a space is connected if and only if it has exactly two distinct clopen sets.

We now study how continuous functions behave with respect to connectedness.

Proposition 6.1.4. *Let X be a connected topological space, and $f : X \to Y$ a surjective function. Then Y is connected.*

Proof. First, since X is non-empty, Y must also be non-empty. If $g : Y \to 2$ is a continuous function, then $g \circ f$ is continuous as well, hence constant by connectedness of X. Therefore there exists $i \in 2$ such that $(g \circ f)^{-1}(i) = X$. But then $g^{-1}(i) = f(f^{-1}(g^{-1}(i))) = f(X) = Y$, hence g is constant. ☐

Another way to state Proposition 6.1.4 is to say that the image of a connected space through a continuous function is connected.

Corollary 6.1.5. *Let X be a connected topological space, and $\sim$ an equivalence relation on X. Then $X/{\sim}$ is connected.*

Corollary 6.1.6. *Connectedness is invariant under homeomorphism.*

Alternatively, Corollary 6.1.6 can be expressed by saying that connectedness is a *topological property*. We can sometimes use this fact to distinguish spaces, i.e. to show that specific pairs of spaces are *not* homeomorphism.

This is a general principle in topology, and more generally in geometry and algebra: an effective technique to prove that two objects cannot be possibly identified by any isomorphisms is to find some "quantity" (usually a number, or some sort of algebraic structure) that is relatively easy to calculate, and that is invariant under the appropriate notion of isomorphism. Then if the two objects happen to have different values for this quantity, we know that they are not isomorphic.

Connectedness is one of those invariant quantities in the context of topological spaces up to homeomorphism. It is a very simple quantity, namely just a truth value. Nonetheless, it can be useful to distinguish spaces, as we will see later.

Proposition 6.1.7. *The closed interval* $[0, 1]$ *is connected.*

Proof. Write $[0, 1] = U \cup V$, with U, V non-empty disjoint open sets. Let m be the supremum of U. Since U is closed, $m \in U$. If $m < 1$, then because U is open in $[0, 1]$, there exists $\varepsilon > 0$ such that $m + \varepsilon \in U$, which is impossible since m is the supremum of U. Therefore $1 \in U$, and similarly $1 \in V$, hence U and V are not disjoint, contradiction. $\square$

Proposition 6.1.8. *Let S be a connected subspace of a topological space. Then its closure $\overline{S}$ is connected.*

Proof. First, $\overline{S}$ is clearly non-empty, since it contains S. Let $f : \overline{S} \to 2$ be a continuous function. Since S is connected, f is constant on S, with value $u \in 2$. If $x \in \overline{S}$, by continuity of f there is a neighbourhood N of S such that f is constantly equal to $f(x)$ on N. Now if $a \in N \cap S$, we have $f(x) = f(a) = u$, hence f is constantly equal to u on all of $\overline{S}$. $\square$

Proposition 6.1.9. *Let $(A_i)_{i \in I}$ be a non-empty collection of connected subspaces of a topological space, and let X be their union. Assume that for any two indices i, j, the intersection $A_i \cap A_j$ is non-empty. Then X is connected.*

Proof. First, X is non-empty, since every A_i is non-empty, and there is at least one of them. Let $f : X \to 2$ be a continuous function. For all $i \in I$, $f|A_i$ is constant, hence there is an element $u_i \in 2$ such that $f(x) = u_i$ for all $x \in A_i$. Now, for any two indices i, j, let $x \in A_i \cap A_j$. Then $u_i = f(x) = u_j$, hence the function $u : I \to 2$ is constant. It follows that f itself is constant. $\square$

Lemma 6.1.10. *Let X, Y be connected topological spaces. Then $X \times Y$ is connected.*

Proof. Since X and Y are non-empty, $X \times Y$ is also non-empty, hence it contains a point $b := (x_0, y_0)$.

For $y \in Y$, let $A_y = \{x_0\} \times Y \cup X \times \{y\}$. Now, $\{x_0\} \times Y$ is connected, because it is homeomorphic to Y, and similarly $X \times \{y\}$ is connected. Their intersection is $\{b\}$, hence A_y is connected by Proposition 6.1.9.

Clearly, $X \times Y$ is the union of the A_y, and since every A_y contains the point b, we can apply Proposition 6.1.9 once more and deduce that $X \times Y$ is connected, as claimed. $\square$

Corollary 6.1.11. *Let $(X_i)_{i \in I}$ be a family of topological spaces, where I is finite. Then $\prod_{i \in I} X_i$ is connected.*

Proof. By induction, it is enough to show the statement just for the empty and binary products. The empty product is the terminal space, which is clearly connected. Connectedness of binary products is Lemma 6.1.10. $\square$

Proposition 6.1.12. *Let $(X_i)_{i \in I}$ be a family of topological spaces. Then $\prod_{i \in I} X_i$ is connected if and only if every X_i is connected.*

Proof. Let X denote the product of the X_i. If X is connected, then $\pi_i(X)$ is connected by Proposition 6.1.4, and since X is in particular non-empty, $\pi_i(X) = X_i$.

For the converse, let F be the collection of finite subsets of I. By the non-emptiness of the X_i, there exists a point $b \in X$. For $J \in F$, let Y_J be the subset of X consisting of those points x such that $x_i = b_i$ for all $i \notin J$. Clearly, $Y_J \cong \prod_{j \in J} X_j$, hence Y_J is connected by Corollary 6.1.11, and clearly $b \in Y_J$.

Therefore, $Y = \bigcup_{J \in F} Y_J$ is connected by Proposition 6.1.9. If we now show that Y is dense in X, connectedness of X will follow from Proposition 6.1.8.

So let U be a non-empty open set in X, and let us show that U meets Y. By Proposition 4.5.2, we can assume that U is a finite intersection of open sets of the form $\pi_i^{-1}(V_i)$ for some V_i non-empty open in X_i, i.e. there exists a finite $J \subseteq I$ such that

$$U = \{x \in X \mid x_j \in V_j \text{ for all } j \in J\}.$$

Now construct $y \in X$ so that $y_j \in V_j$ for $j \in J$, and $y_i = b_i$ for $i \notin J$. Then clearly $y \in U \cap Y$, so U meets Y. $\square$

6.2 Path connected spaces

In many cases, connected spaces satisfy a slightly stronger property, which we now define.

Definition 6.2.1. A topological space X is said to be path connected if it is non-empty, and for all $x, y \in X$ there exists a continuous function $\gamma : [0, 1] \to X$ such that $\gamma(0) = x$ and $\gamma(1) = y$.

A continuous function $\gamma : [0,1] \to X$ is called a *path* in X from $\gamma(0)$ to $\gamma(1)$. Definition 6.2.1 then captures the idea that if we start at any point x in X, we can always reach any other point by following a continuous path.

Proposition 6.2.2. *A path connected space is connected.*

Proof. Let X be path connected, and let $f : X \to 2$ be a surjective function. There exist $x, y \in X$ such that $f(x) = 0$ and $f(y) = 1$. By path connectedness, there exists a path γ from x to y. Now $f \circ \gamma$ is a continuous surjective function $[0,1] \to 2$, but $[0,1]$ is connected by Proposition 6.1.7, so we have reached a contradiction. □

Exercise 58. Let G be the set of points of the form $(x, \sin(1/x)) \in \mathbb{R}^2$ with $x > 0$. Show that $\overline{G}$ is connected, but not path connected. [*Hint: to show that $\overline{G}$ is not path connected, let $\gamma = \langle x, y\rangle$ be a path from say $(0,0)$ to some other point $(x_0, \sin(1/x_0)) \in G$. Let t_0 be the infimum of all $t \in [0,1]$ such that $x(t) > 0$. Continuity of γ implies that there exists an interval $I = [t_0, t_1]$ such that γ restricted to I is a path in $G \cap \mathbb{R} \times [y(t) - 1/4, y(t) + 1/4]$, which is a disconnected space. Obtain a contradiction.*]

Proposition 6.2.3. *Every non-empty convex subset S of $\mathbb{R}^n$ is path connected.*

Proof. For any two points $x, y \in S$, define a path γ by $\gamma(t) = (1-t)x + ty$. Then clearly γ is a path from x to y in $\mathbb{R}^n$, and convexity of S guarantees that $\gamma(t) \in S$. □

Corollary 6.2.4. *The spaces $\mathbb{R}^n$ and $\mathbb{C}^n$ are path connected.*

Path connected spaces enjoy many of the closure properties of connected spaces. As the following propositions show.

Proposition 6.2.5. *Let $f : X \to Y$ be a surjective continuous function between topological spaces. If X is path connected, so is Y.*

Proof. Let y, y' be points in Y. By surjectivity, there exist points $x, x' \in X$ such that $f(x) = y$ and $f(x') = y'$. Let γ be a path in X from x to x'. Then $f \circ \gamma$ is a path in Y between y and y'. □

Corollary 6.2.6. *A quotient of a path connected space is path connected.*

Proposition 6.2.7. *Let $(X_i)_{i\in I}$ be a family of path connected spaces. The product $\prod_{i\in I} X_i$ is path connected if and only if all the X_i are path connected.*

Proof. Let X denote the product of the X_i. If X is path connected, then it is non-empty, hence $X_i = \pi_i(X)$ is path connected by Proposition 6.2.5.

Conversely, suppose that every X_i is path connected, and let a, b be points in X. For all $i \in I$, let γ_i be a path in X_i from a_i to b_i. Then $\gamma = \langle \gamma_i \rangle_{i\in I}$ is a path in X from a to b. □

Paths in a topological space are a very fundamental notion in topology, and especially in algebraic topology. One of the reasons for their importance is that they can be *concatenated*, and concatenation of paths enjoys several important properties, as we will see later.

Definition 6.2.8. Let X be a topological space, let γ be a path in X from x to y and let σ be a path from y to z. Define the *concatenation* of γ and σ to be the path $\gamma * \sigma$ defined as follows:

$$(\gamma * \sigma)(t) = \begin{cases} \gamma(2t) & \text{for } t \leq 1/2 \\ \sigma(2t-1) & \text{for } t \geq 1/2. \end{cases}$$

It follows from Proposition 5.1.4 that $\gamma * \sigma$ is continuous, hence it defines a path from x to z. Intuitively, $\gamma * \sigma$ is a path that follows γ at double speed, and then switches to following σ, also at double speed.

Concatenation of paths allows us to prove the analogous result to Proposition 6.1.9 for path connectedness.

Proposition 6.2.9. *Let $(A_i)_{i \in I}$ be a family of path connected subspaces of a topological space, and let X be their union. Suppose that for all pairs of indices i, j, the intersection of A_i and A_j is non-empty. Then X is path connected.*

Proof. Let X, Y be two points in X, with $x \in A_i$ and $y \in A_j$ for some $i, j \in I$. Let $a \in A_i \cap A_j$. Now choose paths γ from x to a, and σ from a to y. The concatenation $\gamma * \sigma$ is then a path from x to y, as required. $\square$

6.3 Connected components

Geometric intuition suggests that a space can always be decomposed as a union of connected "pieces". This idea can in fact be made precise.

Definition 6.3.1. Let X be a topological space. A *connected component* of X is a maximal connected subset of X.

More explicitly, a subset C of X is a connected component if it is connected, and the only connected subset containing C is C itself. For example, a space X is connected if and only if X is a connected component of itself.

Components can be used to decompose a space into connected spaces.

Proposition 6.3.2. *The connected components of a topological space X are pairwise disjoint, and their union is X.*

Proof. Let C, C' be distinct connected components of X. If $a \in C \cap C'$, then by Proposition 6.1.9 $C \cup C'$ is connected, which violates maximality of C (since $C' \neq \emptyset$).

Now let $x \in X$ be any point. And let $\mathcal{A}$ be the collection of connected subsets of X containing x. Let $C = \bigcup \mathcal{A}$. By Proposition 6.1.9, C is connected. If

$C \subseteq C'$, with C' connected, then C' contains x, hence $C' \in \mathcal{A}$, and therefore $C' = C'$, proving that C is a connected component. So we have that any point of x belongs to some connected component, as required. $\square$

Proposition 6.3.3. *A space is connected if and only if it has exactly one connected component.*

Proof. If X is connected, then X is a connected component, as observed above. If C is any component of X, then $C \subseteq X$, hence $C = X$ by maximality of C, hence X is the only component.

Conversely, if C be the unique component of X, then $X = C$ by Proposition 6.3.2, so X is connected. $\square$

Proposition 6.3.4. *Connected components of a topological space X are closed subsets of X.*

Proof. Let C be a connected component of X. We know from Proposition 6.1.8 that $\overline{C}$ is also connected. But $C \subseteq \overline{C}$, hence by maximality $C = \overline{C}$, which shows that C is closed. $\square$

Corollary 6.3.5. *If a topological space X has a finite number of connected components, then its connected components are open.*

Proof. If C is a connected component of X, then it follows from Proposition 6.3.2 that the complement of C is a finite union of connected components. Since connected components are closed, the complement of C is also closed, hence C is open. $\square$

Exercise 59. Describe the connected components of $\mathbb{Q}$.

There is a completely analogous notion of path connected component, and similar results, the proofs of which are left to reader.

Definition 6.3.6. Let X be a topological space. A path connected component of X is a maximal path connected subset of X.

Proposition 6.3.7. *The path connected components of a topological space X are pairwise disjoint, and their union is X.*

Proposition 6.3.8. *A space is path connected if and only if it has exactly one connected component.*

The set of path connected components of a topological space X is denoted $\pi_0(X)$.

6.4 Examples

Proposition 6.4.1. *The sphere S^n is path connected for all $n \geq 1$.*

Proof. By Proposition 5.4.4, S^n can be written as a union of subsets homeomorphic to D^n. Since D^n is a convex subset of $\mathbb{R}^n$, it is path connected by Proposition 6.2.3. If $n \geq 1$, then the two disks have a non-empty intersection, hence S^n is path connected by Proposition 6.2.9. $\square$

Exercise 60. Generalise Proposition 6.4.1 to the suspension of a non-empty space X.

Note that $S^0 \cong 1 + 1$, hence it is disconnected.

Corollary 6.4.2. *Real projective spaces are connected.*

Corollary 6.4.3. *The torus is connected.*

Proposition 6.4.4. $\mathbb{R}^n$ *is not homeomorphic to* $\mathbb{R}$ *for* $n > 1$.

Proof. Suppose $\phi : \mathbb{R}^n \to \mathbb{R}$ is a homeomorphism. Then ϕ induces a homeomorphism between $\mathbb{R}^n \backslash \{0\}$ and $\mathbb{R} \backslash \{\phi(0)\}$. But the latter is disconnected, while the former is path connected.

One way to verify the last assertion is the following: write $\mathbb{R}^n \backslash \{0\}$ as the union $U_+ \cup U_-$, where U_+ (resp. U_-) is the set of non-zero vectors where the first coordinate is non-negative (resp. non-positive). Both U_+ and U_- are path connected, since every point in U_+ (resp. U_-) can be connected with a path to $(1, 0, ..., 0)$ (resp. $(-1, 0, ..., 0)$), and their intersection is homeomorphic to $\mathbb{R}^{n-1} \setminus \{0\}$, hence it is non-empty for $n > 1$. $\square$

7 Separation and countability axioms

Even though most of our examples so far have been of subspaces of Euclidean spaces and some of their quotients, we have seen some instances of topological spaces that exhibit unfamiliar characteristics, especially if one draws one's intuition from metric spaces. For example, convergence of sequences can behave very poorly on general topological spaces, and it does not necessarily capture the topology of the space.

In this section, we are going to isolate some "well-behavedness" properties of topological spaces that make them behave more like metric spaces in some regards. They generally fall into two categories:

- *separation axioms* express whether it is possible to "separate" points or closed sets from one another;
- *countability axioms* express whether the topology on a space can be recovered from countable collections of open sets.

There are numerous such axioms that have been considered by topologists. Most of them are only of interest in general topology, and do not often figure in related areas like algebraic topology or differential geometry. For this reason, we are going to focus our attention on those properties that have wide applicability.

7.1 Hausdorff spaces

Definition 7.1.1. A topological space X is said to be *Hausdorff* if for all pairs of distinct points $x, y \in X$, there exists disjoint (open) neighbourhoods U and V of x and y respectively.

We say that a space is Hausdorff if we can separate points by neighbourhoods, or simply by open sets. The Hausdorff property is by far the most important and widely used separation axiom, but it exists on a spectrum of increasingly stronger axioms, with names going from T_0 to T_4.

Definition 7.1.2. Let X be a topological space.

X is T_0 if for all pairs of distinct points $x, y \in X$, there exists a neighbourhood of one of the two that does not contain the other.

X is T_1 if for all pairs of distinct points $x, y \in X$, there exists a neighbourhood of x that does not contain y.

X is T_2 if it is Hausdorff, i.e. all pairs of distinct points have disjoint neighbourhoods.

X is T_3 if it is Hausdorff, and for all closed sets $C \subseteq X$, and all points $x \notin C$, there exist disjoint open sets U, V, with $C \subseteq U$ and $x \in V$.

X is T_4 if it is Hausdorff, and for all pairs of disjoint closed sets $C, D \subseteq X$, there exist disjoint open sets U, V, with $C \subseteq U$ and $D \subseteq V$.

It is clear that T_i implies T_j whenever $i > j$. A space is said to be *regular* if closed sets and points can be separated by open sets, so that a T_3 space is equivalently a Hausdorff regular space. Similarly, a *normal* space is one where disjoint closed sets can be separated by open sets, so that a T_4 space is the same as a Hausdorff normal space.

The following exercise shows a useful equivalent formulation of the T_1 axiom.

Exercise 61. Show that a topological space X is T_1 if and only if its points (or more precisely those subsets of X containing a single point) are closed.

None of the implications $T_{i+1} \Rightarrow T_i$ are reversible in general. Some of the counterexamples are easy to find, and they are left as an exercise.

Exercise 62. Find examples of:

- a topological space that is not T_0,
- a T_0 space that is not T_1,
- a T_1 space that is not Hausdorff

Counterexamples for the other implications are harder to construct. We refer the interested reader to [1] and [2].

In the following, we will mostly focus on Hausdorff spaces, although the other separation axioms listed above are sometimes useful to keep in mind.

It is clear that being Hausdorff is invariant under homeomorphism. However, unlike connectedness, it is *not* true that the image of a Hausdorff space through a continuous function is Hausdorff. In fact, any space X is the image of the identity function from X itself equipped with the discrete topology, and any discrete space is clearly Hausdorff.

Proposition 7.1.3. *A subspace of a Hausdorff topological space is Hausdorff.*

Proof. Let X be Hausdorff, and $S \subseteq X$. If $x, y \in S$ are distinct points, there exist disjoint subsets U, V, open in X, such that $x \in U$ and $y \in V$. Now $S \cap U$ and $S \cap V$ are disjoint, contain x and y respectively, and are open in S. □

Proposition 7.1.4. *If $(X_i)_{i \in I}$ is a family of Hausdorff spaces, the product $X := \prod_{i \in I} X_i$ is Hausdorff.*

Proof. If $x, y \in X$ are distinct points, there exists an index $i \in I$ such that $x_i \neq y_i$. Let U, V be disjoint open neighbourhoods of respectively x_i and y_i in X_i. Then $\pi_i^{-1}(U)$ and $\pi_i^{-1}(V)$ are disjoint open neighbourhoods of respectively x and y. □

Unfortunately, the quotient of a Hausdorff topological space is not in general Hausdorff. For example, let $X = 2 \times \mathbb{R}$ be the coproduct of two copies of the real line. By Proposition 7.1.4, X is Hausdorff. Let $\sim$ be the equivalence relation on X generated by $(0, x) \sim (1, x)$ for all $x \neq 0$. The resulting quotient space $X/{\sim}$ is called the *line with double origin*, since it consists of a single copy of all non-zero points of $\mathbb{R}$, plus two copies of 0.

Proposition 7.1.5. *The space $L = X/{\sim}$ defined above is not Hausdorff.*

Proof. Consider the point $a = \pi(0, 0)$. If U is an open set containing a, then $\pi^{-1}(U)$ is open in $2 \times \mathbb{R}$ and it contains $(0, 0)$, hence it contains an interval of the form $\{0\} \times]-\varepsilon, \varepsilon[$. A similar argument applies to $b = \pi(0, 0)$. It follows that every neighbourhood of a meets every neighbourhood of b. Therefore L is not Hausdorff. □

However, we will see that we can impose some conditions on an equivalence relation on a Hausdorff space to make sure that the corresponding quotient is Hausdorff.

Proposition 7.1.6. *A topological space X is Hausdorff if and only if the* diagonal *$\Delta_X := \{(x, x) \mid x \in X\}$ is a closed subset of $X \times X$.*

Proof. By simply reformulating the Hausdorff property on X in terms of $X \times X$, we get that X is Hausdorff if and only if for all $p \in X \times X$, if $p \notin \Delta_X$, then there exists open sets $U, V \subseteq X$ such that $(U \times V) \cap \Delta_X = \emptyset$. Since open sets of the form $U \times V$ are a base of the topology on X by Proposition 5.2.4, the latter condition is equivalent to Δ_X being closed, as required. $\square$

Proposition 7.1.7. *Let X be a Hausdorff topological space, and $\sim$ an equivalence relation on X. Assume the projection function $\pi : X \to X/{\sim}$ is open. Then $X/{\sim}$ is Hausdorff if and only if $R = \{(x, y) \mid x \sim y\}$ is a closed subset of $X \times X$.*

Proof. Let $\pi_2 : X \times X \to X/{\sim} \times X/{\sim}$ be the function defined by $\pi_2(x, y) = (\pi(x), \pi(y))$. Since π is continuous and open, it follows that π_2 is also continuous and open. By Exercise 26, a subset C of $X/{\sim} \times X/{\sim}$ is closed if and only if $\pi_2^{-1}(C)$ is closed in $X \times X$. Now, clearly $R = \pi_2^{-1}(\Delta_{X/\sim})$, hence the diagonal is closed in $X/{\sim}$ if and only if R is closed, and the conclusion follows from Proposition 7.1.6. $\square$

Exercise 63. The real projective spaces are Hausdorff.

The property of a space of being Hausdorff can be directly characterised in terms of convergence of filters.

Proposition 7.1.8. *A space X is Hausdorff if and only if every filter has at most one limit.*

Proof. Let $\mathcal{F}$ be a filter on a Hausdorff space X, and assume that $\mathcal{F}$ converges to both ℓ and ℓ', with $\ell \neq \ell'$. Let U, V be disjoint open neighbourhoods of ℓ and ℓ' respectively. Then $U \in \mathcal{F}$ and $V \in \mathcal{F}$, hence $\emptyset = U \cap V \in \mathcal{F}$, which is a contradiction.

Conversely, assume that X is not Hausdorff. Then there exist distinct points $\ell, \ell' \in X$ such that every neighbourhood of ℓ intersects every neighbourhood of ℓ'. It follows that $\mathcal{N}(\ell) \cup \mathcal{N}(\ell')$ has the finite intersection property, hence it is contained in a filter $\mathcal{F}$, which then converges to both ℓ and ℓ'. $\square$

7.2 First and second countability

If X is a metric space, and $x \in X$, then every neighbourhood of x contains a ball with a radius $r \in \mathbb{Q}$. Since the set of rationals is countable, this simply observation usually allows us to express many topological properties of metric spaces using only countable collections.

We can abstract this feature of metric spaces into a general definition.

Definition 7.2.1. Let X be a topological space, and $x \in X$. A *system of neighbourhoods* of x is a collection $\mathcal{U}$ of neighbourhoods of x such that every neighbourhood of x contains an element of $\mathcal{U}$.

So balls centred in x with rational radius (or even with radius of the form say 2^{-n} for $n \in \mathbb{N}$) form a system of neighbourhoods of x in any metric space.

Definition 7.2.2. A topological space X is said to be *first countable* if every point of X admits a countable system of neighbourhoods.

Proposition 7.2.3. *A metric space is first countable.*

Proof. Immediate consequence of the above observation. □

Exercise 64. Let X be an uncountable set equipped with the cofinite topology. Show that X is not first countable.

First countability is a property which can be checked one point at a time, and it does not say much about the topology of a space as a whole. Therefore, we introduce a new, stronger, notion.

Definition 7.2.4. A topological space X is said to be *second countable* if X has a countable base of open sets.

Proposition 7.2.5. *A second countable topological space X is first countable.*

Proof. Let $\mathcal{B}$ be a countable base for the topology of X. For any point $x \in X$, let $\mathcal{B}_x$ be the set of elements of $\mathcal{B}$ that contain x. Clearly, $\mathcal{B}_x$ is a system of neighbourhoods of x, and it is also countable, since it is a subset of $\mathcal{B}$. □

Exercise 65. Prove that a discrete space is second countable if and only if its underlying set is countable.

Note that metric spaces are not necessarily second countable. For example, let X be any uncountable set equipped with the metric d defined by $d(x, y) = 1$ for $x \neq y$. Then the topology on X induced by this metric is discrete (Exercise 5), hence X cannot be second countable by Exercise 65. On the other hand, metric spaces are always first countable by Proposition 7.2.3.

8 Compactness

8.1 Definition and basic properties

Recall that an *open cover* of a topological space X is a family of open sets $\mathcal{U} = (U_i)_{i \in I}$ such that the union of the U_i is X. A subset $J \subseteq I$ determines a new cover $(U_j)_{j \in J}$, which will be referred to as a *subcover* of $\mathcal{U}$.

The reader is likely familiar with the notion of *compactness* for subsets of $\mathbb{R}^n$, where a subset C is called *compact* if it is closed and bounded. We will now give a general and intrinsic definition of compactness. By general we mean that it makes sense for arbitrary spaces, and by intrinsic we mean that, unlike the definition above, it does not refer to some larger "ambient" space.

Definition 8.1.1. A topological space X is said to compact if all open covers of X have a finite subcover.

It is not immediately obvious that this definition specialises to the one for subsets of $\mathbb{R}^n$. Proving this fact will be one of the goals of this section.

Proposition 8.1.2. *A closed subspace of a compact topological space is compact.*

Proof. Let C be a closed subset of a topological space X. To show that C is compact, let $\mathcal{U} = (U_i)_{i \in I}$ be an open cover of C. By definition of subspace topology, there is a family $(V_i)_{i \in I}$ of open sets in X such that $U_i = V_i \cap C$. Therefore, the V_i together with $X \setminus C$ form a cover of X, so by compactness there exists a finite $J \subseteq I$ such that $(V_j)_{j \in J} \cup (X \setminus C)$ is a cover of X. It follows that $(U_j)_{j \in J}$ is a finite subcover of $\mathcal{U}$, as required. □

Proposition 8.1.3. *A compact subspace of a Hausdorff space is closed.*

Proof. Let K be a compact subspace of a topological space X, and assume that X is Hausdorff. To show that $X \setminus K$ is open, we fix a point $x \notin K$, and show that there is a neighbourhood of x which is disjoint from K. For $y \in K$, let U_y, V_y be disjoint neighbourhoods of x and y respectively. Clearly,

$$K \subseteq \bigcup_{y \in K} V_y,$$

hence there exists a finite $J \subseteq K$ such that the V_y, for $y \in K$, cover K. Let $W = \bigcap_{y \in J} U_y$. Since J is finite, W is a neighbourhood of x, and

$$W \cap K \subseteq W \cap \bigcup_{y \in J} V_y \subseteq \bigcup_{y \in J} W \cap V_y \subseteq \bigcup_{y \in J} U_y \cap V_y = \emptyset.$$

□

Proposition 8.1.4. *Let $f : X \to Y$ be a surjective continuous function, where X is a compact topological space. Then Y is compact.*

Proof. Let $\mathcal{U} = (U_i)_{i \in I}$ be an open cover of Y. By continuity of f, the family $(f^{-1}(U_i))_{i \in I}$ is an open cover of X, from which we can extract a finite subcover indexed by say $J \subseteq I$. Now

$$\bigcup_{j \in J} U_j = f\left(f^{-1}\left(\bigcup_{j \in J} U_j\right)\right) = f\left(\bigcup_{j \in J} f^{-1}(U_j)\right) = f(X) = Y,$$

so J determines a finite subcover of $\mathcal{U}$. □

Corollary 8.1.5. *A quotient of a compact topological space is compact.*

Corollary 8.1.6. *Compactness is invariant under homeomorphism.*

Putting together the basic facts about compact spaces proved above, we can finally prove the result about bijective continuous functions which we have been using numerous times in previous examples.

Proposition 8.1.7. *Let X be a compact space, Y a Hausdorff space. If $f : X \to Y$ is a continuous function, then f is closed.*

Proof. Let C be closed in X. By Proposition 8.1.2, C is compact. Therefore, $f(C)$ is compact by Proposition 8.1.4. Finally, $f(C)$ is closed in Y by Proposition 8.1.3. $\square$

Corollary 8.1.8. *Let X be a compact space, Y a Hausdorff space. If $f : X \to Y$ is a continuous bijective function, then f is a homeomorphism.*

Proof. By Proposition 8.1.7, f is closed, therefore f is a homeomorphism by Proposition 3.3.7. $\square$

Proposition 8.1.9. *Let $(X_i)_{i\in I}$ be a* finite *family of compact spaces. Then the coproduct $X := \coprod_{i\in I} X_i$ is compact.*

Proof. Let $(U_j)_{j\in J}$ be an open cover of X, and let $\iota_i : X_i \to X$ denote the coproduct injections. Then for all $i \in I$, the family $(\iota_i^{-1}(U_j))_{j\in J}$ is an open cover of X_i, hence by compactness we can find a finite subset F_i of J such that the $i_i^{-1}(U_j)$ cover X_i for $j \in F_i$, hence in particular

$$\bigcup_{j\in F_i} U_j \supseteq \iota_i(X_i).$$

It follows that $F := \bigcup_{i\in I} F_i$ determines a subcover of X, and F is finite because I and the F_i are. $\square$

Corollary 8.1.10. *Let $(K_i)_{i\in I}$ be a finite collection of compact subspaces of a topological space X. Then the union $K := \bigcup_{i\in I} K$ is compact.*

Proof. The space K is the image of the map $\coprod_{i\in I} K_i \to X$ induced by the inclusion functions $K_i \to X$. Since $\coprod_{i\in I} K_i$ is compact by Proposition 8.1.9, it follows that K is compact by Proposition 8.1.4. $\square$

To prove that a space is compact, it is sometimes useful to be able to shrink the open sets of a cover when constructing a finite subcover. In order to make this precise, we introduce the notion of refinement.

Definition 8.1.11. Let $\mathcal{A} = (A_i)_{i\in I}$ be a cover of a set X. A *refinement* of $\mathcal{A}$ is a cover $\mathcal{B} = (B_j)_{j\in J}$ such that for all $j \in J$ there exists $i \in I$ with $B_j \subseteq A_i$.

In other words, a refinement of a cover is a cover obtained by replacing some of the elements of the original cover with smaller ones, and possibly removing some altogether. In particular, a subcover is a refinement.

Lemma 8.1.12. *A topological space X is compact if and only if every open cover of X admits a finite refinement.*

Proof. Since, as observed above, every subcover is a refinement, one direction is obvious. For the converse, let $\mathcal{U} = (U_i)_{i\in I}$ be an open cover of X, and suppose that it has a finite refinement $(V_j)_{j\in J}$. For all $j \in J$, let $k(j)$ be such that $V_j \subseteq U_{k(j)}$, and observe that the family $(U_k(j))_{j\in J}$ is a subcover of $\mathcal{U}$, as required. $\square$

Lemma 8.1.13. *Let X be a topological space, and $\mathcal{B}$ be a base of the topology of X. Then X is compact if and only if all covers of X consisting of elements of $\mathcal{B}$ admit a finite subcover.*

Proof. If X is compact, then it clearly also satisfies the compactness condition on covers consisting of elements of $\mathcal{B}$. Conversely, let $\mathcal{U} = (U_i)_{i\in I}$ be an arbitrary cover. For all $i \in I$, write U_i as

$$\bigcup_{j\in J_i} V_j,$$

for some set J_i, and some $V_j \in \mathcal{B}$. Then $\mathcal{V} = (V_j)_{\substack{i\in I\\ j\in J_i}}$ is a refinement of $\mathcal{U}$ consisting of elements of $\mathcal{B}$. By the assumption on X, the cover $\mathcal{V}$ has a finite subcover, which is therefore a finite refinement of $\mathcal{U}$. It then follows from Lemma 8.1.12 that X is compact. $\square$

Proposition 8.1.14. *Let X and Y be compact topological spaces. Then $X \times Y$ is compact.*

Proof. Fix a cover $\mathcal{U}$ of $X \times Y$. By Lemma 8.1.13, we can assume that $\mathcal{U}$ is of the form $(U_i \times V_i)_{i\in I}$, where the U_i (resp. V_i) are open in X (resp. Y).

For $x \in X$, let J_x be a finite subset of I such that

$$\{x\} \times Y \subseteq \bigcup_{j\in J_x} U_j \times V_j;$$

this exists because $\{x\}\times Y$ is homeomorphic to Y, hence compact. In particular,

$$Y = \bigcup_{j\in J_x} V_j.$$

Let $W_x = \bigcap_{j\in J_x} U_j$. Clearly $x \in W_x$, hence $(W_x)_{x\in X}$ is an open cover of X. By compactness, there exists a finite $K \subseteq X$ such that

$$X = \bigcup_{x\in K} W_x.$$

Now, if $J' = \bigcup_{x\in K} J_x$, we have:

$$\bigcup_{j\in J'} U_j \times V_j = \bigcup_{x\in K} \bigcup_{j\in J_x} U_j \times V_j \supseteq \bigcup_{x\in K} W_x \times \left(\bigcup_{j\in J_x} V_j\right) = \bigcup_{x\in K} W_x \times Y = X \times Y,$$

and we have found a finite subcover of $\mathcal{U}$. $\square$

Corollary 8.1.15. *A finite product of compact topological spaces is compact.*

Proof. Immediate by induction from Proposition 8.1.14 and the fact that the terminal space is compact. ☐

Proposition 8.1.16. *The closed interval* $[0, 1]$ *is compact.*

Proof. Let $(U_i)_{i\in I}$ be an open cover of $[0, 1]$, and let S be the set of points $x \in [0, 1]$ such that $[0, x]$ is covered by a finite subset of the U_i.

Clearly S is non-empty, because certainly $0 \in S$, hence S has a supremum m. We want to show that $m = 1$.

If that is not the case, let $i \in I$ such that $m \in U_i$, and let $\varepsilon > 0$ be such that $m + \varepsilon \in U_i$. Since clearly $m \in \overline{S}$, there must be a point $m' \in U_i \cap S$, so let $J \subseteq I$ be a finite set such that $(U_j)_{j\in J}$ is a cover of $[0, m']$. It then follows that $J \cup \{i\}$ determines a finite cover of $[0, m + \varepsilon]$, hence $m + \varepsilon \in S$, contradicting the fact that m is the supremum of S. ☐

Proposition 8.1.17. *Let* K *be a compact subset of* $\mathbb{R}$*. Then* K *is bounded.*

Proof. Clearly,

$$(]x - 1, x + 1[\cap K)_{x\in K}$$

is an open cover of K, so by compactness of K, there exists a finite $J \subseteq K$ such that

$$K \subseteq \bigcup_{x\in J}]x - 1, x + 1[.$$

If a, b are respectively the minimum and maximum element of J, it follows that $K \subseteq [a - 1, b + 1]$, hence it is bounded. ☐

Corollary 8.1.18. *Let* X *be a compact topological space, and* $f : X \to \mathbb{R}$ *a continuous function. Then* f *is bounded.*

Proof. Immediate consequence of Proposition 8.1.17 and Proposition 8.1.4. ☐

Proposition 8.1.19. *A subset of* $\mathbb{R}^n$ *is compact if and only if it is closed and bounded.*

Proof. If C is a closed and bounded subset of $\mathbb{R}^n$, then it is homeomorphic to a closed subspace of a cube $[0, 1]^n$, which is compact by Corollary 8.1.15. Therefore, C is compact by Proposition 8.1.2.

Conversely, let C be a compact subspace of $\mathbb{R}^n$. Then C is closed by Proposition 8.1.3. Finally, the norm function $\| - \| : C \to \mathbb{R}^n$ is continuous, hence bounded by Corollary 8.1.18, which means that C itself is bounded. ☐

8.2 Examples

Proposition 8.2.1. *The spheres S^n and the disks D^n are compact spaces.*

Proof. Immediate consequence of Proposition 8.1.19. $\square$

Corollary 8.2.2. *The torus is compact.*

Proof. The torus is a product of compact spaces, hence compact by Proposition 8.1.14. $\square$

Corollary 8.2.3. *The real projective spaces are compact.*

Proof. Regard $\mathbb{R}P^n$ as a quotient of S^n, then apply Corollary 8.1.5. $\square$

Using compactness, we can already show that many pairs of spaces are not homeomorphic. For example:

Proposition 8.2.4. *No sphere is homeomorphic to an open subset of $\mathbb{R}^n$.*

Proof. Spheres are compact, but it follows from Proposition 8.1.19 that the only compact open subset of $\mathbb{R}^n$ is the empty set. $\square$

8.3 Compactness and sequences

The reader may be familiar with a different definition of compactness, given in terms of sequences. As we will see, this definition is equivalent to the general notion of compactness of Definition 8.1.1 in the case of metric spaces, but the two notions disagree for general topological spaces.

Recall that, given a sequence $x : \mathbb{N} \to X$, and a strictly increasing function $k : \mathbb{N} \to \mathbb{N}$, the composition $x \circ k : \mathbb{N} \to X$ is new sequence, which is said to be a *subsequence* of x.

Definition 8.3.1. A space is said to be *sequentially compact* if every sequence in X admits a convergent subsequence.

It is not too surprising that, without some countability assumption, sequential compactness and compactness cannot be directly compared, and we will look at examples showing that neither of the two notions implies the other in general. However, for first countable spaces, we can show that compactness is stronger than sequential compactness.

Lemma 8.3.2. *Let ℓ be a point with a countable system of neighbourhoods in a topological space X. A sequence x has a subsequence which converges to ℓ if and only if for all neighbourhoods N of ℓ, the sequence x is frequently in N.*

Proof. Suppose first that $x \circ k \to \ell$ for some strictly increasing function $k : \mathbb{N} \to \mathbb{N}$. Then for all neighbourhoods N of ℓ, $x \circ k$ is eventually in N, which immediately implies that x is frequently in N.

Conversely, let $(V_n)_n$ be a countable system of neighbourhoods of ℓ. By replacing V_n with $V_1 \cap \cdots \cap V_n$, we can assume that $V_n \supseteq V_{n+1}$ for all n. Our goal is to construct a subsequence y of x such that for all neighbourhoods N of ℓ, y is eventually in N. To that end, we have to define a strictly monotone function $k : \mathbb{N} \to \mathbb{N}$, and we do so by induction.

As induction hypothesis, assume we have defined $k(i)$ for $i < n$. By assumption, x is frequently in V_n, hence there exists $m \geq k(n-1)+1$ such that $x_m \in V_n$. Therefore, we set $k(n) := m$. This produces a function k which is strictly monotone by construction. Now let $y = x \circ k$. For any fixed $n \in \mathbb{N}$, we have $y_i = x_{k(i)} \in V_i \subseteq V_n$, hence in particular y is eventually in V_n. Since the V_n form a system of neighbourhoods of ℓ, it follows that y satisfies the required property. □

Proposition 8.3.3. *A first countable compact topological space is sequentially compact.*

Proof. Let X be a first countable compact topological space, and suppose by contradiction that there exists a sequence u in X with no convergent subsequence. By Lemma 8.3.2, every point $x \in X$ has a neighbourhood U_x such that it is not the case that u is frequently in U_x, or in other words, u is eventually in the complement of U_x.

By compactness, there exists a finite $J \subseteq X$ such that

$$X = \bigcup_{x \in J} U_x.$$

Let n_x be such that $u_m \notin U_x$ for $m \geq n_x$. Then for $m \geq \max_{x \in J} n_x$ we have $x_m \notin \bigcup_{x \in J} U_x = X$, which is clearly a contradiction. □

To see why the first countability assumption is necessary, let us consider the space

$$X := \prod_{t \in 2^{\mathbb{N}}} 2,$$

where 2 denotes a discrete space with 2 elements, and $2^{\mathbb{N}}$ is the set of functions $\mathbb{N} \to 2$. There is a “canonical” sequence $x = (x^n)_{n \in \mathbb{N}}$ in X, defined as follows:

$$x^n_t = t(n).$$

It is easy to see that x does not have any convergent subsequence. For if there exists a strictly increasing $k : \mathbb{N} \to \mathbb{N}$ such that $x \circ k$ converges to some $\ell \in X$, then in particular the sequence $\left(x^{k(n)}_t\right)_{n \in \mathbb{N}}$ should converge for any $t : \mathbb{N} \to 2$. But if we choose t such that $t(n) = 1$ if n is of the form $k(2m)$ for some m, and $t(n) = 0$ otherwise, then the resulting sequence alternates between 0 and 1, hence it does not converge.

We have therefore proved that the space X is not sequentially compact. However, it will follow from Tychonoff's theorem (Theorem 8.4.10) that X is compact.

The following exercise shows that sequential compactness does not imply compactness even for first countable spaces.

Exercise 66. Show that the first uncountable ordinal ω_1, with the order topology, is first countable and sequentially compact, but not compact.

To get compactness from sequential compactness, we need to strengthen the countability assumption.

Proposition 8.3.4. *Let X be a second countable sequentially compact topological space. Then X is compact.*

Proof. By Lemma 8.1.13, it is enough to show that every countable open cover has a finite subcover. Suppose by contradiction that there exists a countable open cover $(U_n)_{n\in\mathbb{N}}$ with no finite subcover, and let x_n be a point in the complement of $\bigcup_{i<n} U_i$. Let $x = (x_n)_{n\in\mathbb{N}}$ be the resulting sequence, and $k : \mathbb{N} \to \mathbb{N}$ a strictly increasing function such that $x \circ k$ converges to a point $\ell \in X$.

If $m \in \mathbb{N}$ is such that $\ell \in U_m$, then $x \circ k$ is eventually in U_m, which implies that x is frequently in U_m. In particular, there exists $n \geq m$ such that $x_n \in U_m$, which contradicts the defining property of x_n. □

Thanks to Proposition 8.3.3 and Proposition 8.3.4, we can show that sequential compactness and compactness are equivalent for metric spaces.

Lemma 8.3.5. *Let X be a sequentially compact metric space. Then for all $\varepsilon > 0$, there exists a finite subset S of X such that*

$$X = \bigcup_{x\in S} B_\varepsilon(x).$$

Proof. Fix $\varepsilon > 0$, and assume by contradiction that X cannot be covered by finitely many balls of radius ε. By induction, we can construct a sequence $(x_n)_{n\in\mathbb{N}}$ such that

$$x_n \notin \bigcup_{i<n} B_\varepsilon(x_i).$$

Now let $k : \mathbb{N} \to \mathbb{N}$ be a strictly increasing function such that $x \circ k$ converges to a point $\ell \in X$. Let $n_0 \in \mathbb{N}$ be such that $d(x_{k(n)}, \ell) < \varepsilon/2$ for $n \geq n_0$. Then

$$d(x_{k(n)}, x_{k(n+1)}) \leq d(x_{k(n)}, \ell) + d(x_{k(n+1)}, \ell) < \varepsilon/2 + \varepsilon/2 = \varepsilon,$$

which means that $x_{k(n+1)} \in B_\varepsilon(x_{k(n)})$, contradiction. □

Proposition 8.3.6. *Let X be a metric space. Then X is compact if and only if it is sequentially compact.*

Proof. Since metric spaces are first countable, it follows from Proposition 8.3.3 that compactness implies sequential compactness for metric spaces.

For the converse, thanks to Proposition 8.3.4, it is enough to show that if a metric space X is sequentially compact, then it is second countable. For all $n \in \mathbb{N}$, let S_n be a finite subset of X such that X is covered by the family of balls of radius 2^{-n} with centres in S_n, as given by Lemma 8.3.5. Define a collection of open sets $\mathcal{B}$ as follows:

$$\mathcal{B} = \{B_{2^{-n}}(y) \mid n \in \mathbb{N}, y \in S_n\}.$$

Clearly $\mathcal{B}$ is countable, so it remains to show that it generates the topology on X. If $B_r(x)$ is an arbitrary ball, it is enough to find an element of $\mathcal{B}$ that contains x and is contained in $B_r(x)$. To that end, choose n such that $2^{-n} < \varepsilon/2$. Now, if $y \in S_n$ is such that $d(x, y) < 2^{-n}$, then clearly $x \in B_{2^{-n}}(y) \subseteq B_r(x)$, as required. □

8.4 Compactness and filters

We have seen that compactness can be characterised in terms of convergence of sequences, but only for metric spaces. Using filters, we can obtain a characterisation of compactness based on convergence that works for arbitrary topological spaces. First, we reformulate compactness using closed sets, simply by "dualising" the definition.

Proposition 8.4.1. *A topological space X is compact if and only if every collection of closed sets in X with the finite intersection property has non-empty intersection.*

Proof. Observe that open covers on a topological space X are in bijective correspondence with collections of closed sets with empty intersection. The correspondence works by assigning to a cover $\mathcal{U}$ the collection of the complements of the open sets of $\mathcal{U}$. Using this bijection, one can see that X is compact if and only if every collection of closed sets with empty intersection admits a finite subcollection with empty intersection.

By taking the contra-positive, we get that a space is compact if and only if for every collection of open sets such that every finite subcollection has non-empty intersection, the whole collection has non-empty intersection, which is exactly the statement that we want to prove. □

Proposition 8.4.2. *A topological space X is compact if and only if every filter on X is contained in a convergent filter.*

Proof. Let $\mathcal{F}$ be a filter on a compact space X, and let $\mathcal{A}$ be the collection of closed subsets of X that belong to $\mathcal{F}$. Since $\mathcal{F}$ has the finite intersection property, so does $\mathcal{A}$, hence by compactness of X there exists $\ell \in X$ which is contained in all the elements of $\mathcal{A}$. To show that $\mathcal{F}$ is contained in a filter that converges, by Proposition 4.6.5 it is enough to show that for all neighbourhoods

N of ℓ, and all $A \in \mathcal{F}$, we have $N \cap A \neq \emptyset$. By contradiction, suppose $N \cap A = \emptyset$. Then in particular $U \cap A = \emptyset$ for some open neighbourhood U of ℓ, from which it follows that $U \cap \overline{A} = \emptyset$. But $\overline{A} \in \mathcal{A}$, hence $\ell \in \overline{A}$, which is a contradiction.

Conversely, let $\mathcal{A}$ be a collection of closed sets with the finite intersection property, and let $\mathcal{F}$ be a filter containing $\mathcal{A}$. By the assumption on X, $\mathcal{F}$ (hence $\mathcal{A}$) is contained in filter that converges to a point $\ell \in X$, therefore in particular every neighbourhood of ℓ meets every element of $\mathcal{A}$. It follows that for all $C \in \mathcal{A}$, we have $\ell \in \overline{C} = C$, which means that the intersection of $\mathcal{A}$ contains ℓ, hence it is non-empty. $\square$

Because the formulation of compactness in terms of filters is based on enlarging a filter to a convergent one, it is sometimes useful to work with filters that are "maximal". This idea is captured by the following definition.

Definition 8.4.3. Let X be a set. An *ultrafilter* on X is a maximal filter on X, i.e. a filter $\mathcal{F}$ such that for all filters $\mathcal{G}$ on X, if $\mathcal{F} \subseteq \mathcal{G}$ then $\mathcal{F} = \mathcal{G}$.

Exercise 67. Let $x \in X$. Prove that the collection of subsets of X that contain x is an ultrafilter on X.

The ultrafilter defined in Exercise 67 is called a *principal* ultrafilter. It turns out that principal ultrafilters are the only examples of ultrafilters that can in some sense be "explicitly" defined. However, existence of non-principal ultrafilters can indeed be derived in a sufficiently powerful set theory, as the following proposition shows.

Proposition 8.4.4. *Every filter is contained in an ultrafilter.*

Proof. Let $\mathcal{F}$ be a filter on a set X. The set of filters containing $\mathcal{F}$ is clearly non-empty, and it is clear that the union of a totally ordered family of filters is a filter. Therefore, Zorn's lemma implies that there is a maximal filter $\mathcal{G}$ containing $\mathcal{F}$, and $\mathcal{G}$ is a ultrafilter by construction. $\square$

Corollary 8.4.5. *Let X be an infinite set. Then there exist non-principal ultrafilters on X.*

Proof. Let $\mathcal{F}$ be the filter of cofinite subsets of X, and let $\mathcal{G}$ be an ultrafilter containing $\mathcal{F}$. If $\mathcal{G}$ is principal, then it contains $\{x\}$ for some $x \in X$, hence its complement cannot be in $\mathcal{G}$ even though it is cofinite. $\square$

Exercise 68. Prove that every ultrafilter on a finite set is principal.

Exercise 69. Show that a filter $\mathcal{F}$ on a set X is an ultrafilter if and only if, for all $A \subseteq X$, either $A \in \mathcal{F}$ or $X \setminus A \in \mathcal{F}$.

Proposition 8.4.6. *Let $\mathcal{F}$ an ultrafilter on a set X, and let $f : X \to Y$ be any function. Then the filter $f_*(\mathcal{F})$ is an ultrafilter on Y.*

Proof. Let $B \subseteq Y$, and assume $B \notin f_*(\mathcal{F})$. Then $f^{-1}(B) \notin \mathcal{F}$, hence $X \setminus f^{-1}(B) \in \mathcal{F}$, by Exercise 69. But $X \setminus f^{-1}(B) = f^{-1}(Y \setminus B)$, hence $Y \setminus B \in f_*(\mathcal{F})$. The conclusion then follows from Exercise 69 again. □

Corollary 8.4.7. *A topological space X is compact if and only if every ultrafilter in X converges.*

Proof. Let X be compact, and $\mathcal{F}$ an ultrafilter on X. Then $\mathcal{F} \subseteq \mathcal{G}$, where $\mathcal{G}$ is a convergent filter. But $\mathcal{F}$ is maximal, hence $\mathcal{F} = \mathcal{G}$, so $\mathcal{F}$ is convergent.

Conversely, let $\mathcal{F}$ be a filter on X. There is an ultrafilter $\mathcal{G}$ containing $\mathcal{F}$ by Proposition 8.4.4, and $\mathcal{G}$ is a convergent filter by the assumption on X. □

Proposition 8.4.8. *A topological space X is Hausdorff if and only if every ultrafilter in X has at most one limit.*

Proof. If X is Hausdorff, then every filter on X has at most one limit by Proposition 7.1.8, hence in particular every ultrafilter.

Conversely, suppose every ultrafilter has at most one limit. If $\mathcal{F}$ is any filter on X, then $\mathcal{F} \subseteq \mathcal{G}$ for some ultrafilter $\mathcal{G}$ on X. If $\mathcal{F}$ converges to $\ell, \ell' \in X$, then also $\mathcal{G}$ converges to ℓ and ℓ', hence $\ell = \ell'$ by the assumption on X. □

Corollary 8.4.9. *A topological space X is compact and Hausdorff if and only if every ultrafilter in X has a unique limit.*

Using the characterisation of compactness in terms of ultrafilters, we can now prove a general statement about compactness of products topological spaces. To appreciate why the formulation in terms of ultrafilters is useful, the reader is encouraged to try to prove this result directly using the definition of compactness in terms of open covers.

Theorem 8.4.10. *Let $(X_i)_{i \in I}$ be a family of compact topological spaces. Then the product $X := \prod_{i \in I} X_i$ is compact.*

Proof. Let $\mathcal{F}$ be an ultrafilter on X. Then $(\pi_i)_*(\mathcal{F})$ is an ultrafilter on X_i by Proposition 8.4.6, hence it converges to some $x_i \in X_i$. Let $x = (x_i)_{i \in I}$ be the resulting point of X. To show that $\mathcal{F}$ converges to x, it is enough to show that $\mathcal{F}$ contains neighbourhoods of x of the form $\pi_i^{-1}(U_i)$, where U_i is a neighbourhood of x_i in X_i. But we know that $U_i \in (\pi_i)_*(\mathcal{F})$, which means that $\pi_i^{-1}(U_i) \in \mathcal{F}$, as required. □

8.5 One-point compactification

It is possible to turn any topological space X into a compact one by adding a single point. The resulting space is called the *one-point* (or *Alexandroff*) *compactification* of X. To study its properties, we first introduce an auxiliary definition which generalises the notion of compactness.

Definition 8.5.1. We say that a space X is *locally compact* if every point of X has a compact neighbourhood.

A compact space is trivially locally compact, since we can take the whole space as a compact neighbourhood of every point.

Exercise 70. Prove that there are subspaces of a Euclidean space that are not locally compact.

Lemma 8.5.2. *Let X be a compact Hausdorff space. Then X is regular, i.e. closed sets and points can be separated by open sets.*

Proof. Let C be a closed set in X, and $x \notin C$. Since X is compact, C is compact by Proposition 8.1.2. For any $y \in C$, let V_y and U_y be disjoint open neighbourhoods of y and x respectively. By compactness of C, there exists a finite subset J of C such that $V := \bigcup_{y \in J} V_y \supseteq C$. Now set $U := \bigcap_{y \in J} U_y$. Then U and V are disjoint open sets containing x and C respectively. $\square$

Exercise 71. Show that a compact Hausdorff space is normal, i.e. pairs of disjoint closed sets can be separated by open sets.

Proposition 8.5.3. *Let X be a compact Hausdorff space. Then every point of x has a system of compact neighbourhoods.*

Proof. Let U be an open subset of X, and $x \in X$. We have to construct a compact neighbourhood of x contained in U. Let W and V disjoint open sets containing $X \setminus U$ and x respectively, which exist by Lemma 8.5.2. Then $K := X \setminus W \subseteq U$, so it is a neighbourhood of x. Finally, K is closed, hence compact by Proposition 8.1.2. $\square$

Corollary 8.5.4. *Let X be a locally compact Hausdorff space. Then every point of x has a system of compact neighbourhoods.*

Proof. Let $x \in X$, and K a compact neighbourhood of x. Let $U \subseteq K$ be an open set containing x. Then U is also open in K, hence it contains a compact neighbourhood K' of x by Proposition 8.5.3. It is then clear that K' is also a neighbourhood of x in the topology of X. $\square$

Corollary 8.5.5. *An open subspace of a compact Hausdorff space X is locally compact.*

Proof. If $U \subseteq X$ is open, and $x \in U$, then U contains a compact neighbourhood K of x by Proposition 8.5.3, and K is also a neighbourhood of x in the topology of U. $\square$

One-point compactification allows us to prove a converse for Corollary 8.5.5. If X is a topological space, we define the one-point compactification of X^∞ to be the *set* $X + 1$, equipped with the topology we are going to describe below. We will identify X with the corresponding subset of X^∞ along the coproduct injection $X \to X^\infty$. The remaining point of X^∞, determined by the other coproduct injection $1 \to X^\infty$, will be denoted by ∞.

We declare a subset $U \subseteq X^\infty$ open if one of the following conditions holds:

- $\infty \notin U$, and U is open in X;
- $\infty \in U$, and $X \setminus U$ is closed and compact.

In particular, X itself is open as a subset of X^∞.

Proposition 8.5.6. *The collection of open subsets of X^∞ defined above forms a topology.*

Proof. Let $(U_i)_{i\in I}$ be a family of open sets, and let V be their union. If none of them contains ∞, then neither does V, hence V is open in X^∞ since it is open in X. If at least one of them, say U_{i_0}, contains ∞, then let $K = X \setminus U_{i_0}$. Now,

$$X \setminus V = \bigcap_{i\in I}(X \setminus U_i) \subseteq K,$$

so $X \setminus V$ is closed and contained in K. Since K is also closed, $X \setminus V$ is closed in K as well, hence $X \setminus V$ is compact by Proposition 8.1.2.

Now suppose that I is finite, and let W be the intersection of the U_i. If all of the U_i contain ∞, then so does their union. Since $X \setminus W$ is a finite union of compact spaces, it is compact by Corollary 8.1.10, so W is open. If instead one of the U_i does not contain ∞, then $\infty \notin W$, and we can write W as a finite intersection of open sets of X. □

Note that the topology on X when regarded as an open subspace of X^∞ does coincide with the original topology on X, essentially by construction.

Proposition 8.5.7. *For any topological space X, the one-point compactification X^∞ of X is compact.*

Proof. Let $(U_i)_{i\in I}$ be an open cover of X. Let $i_0 \in I$ be such that $\infty \in U_{i_0}$, and let $K = X \setminus U_{i_0}$. Then $(U_i \cap K)_{i\in I}$ is an open cover of K, hence there exists a finite subset J of I such that

$$\bigcup_{j\in J} U_j \supseteq K.$$

In particular, $J \cup \{i_0\}$ determines a finite subcover of X^∞. □

Proposition 8.5.8. *Let X be a topological space. Then X is locally compact and Hausdorff if and only if X^∞ is compact and Hausdorff.*

Proof. If X^∞ is Hausdorff, then so is X by Proposition 7.1.3, and if X^∞ is compact, then X is locally compact by Corollary 8.5.5.

Conversely, suppose that X is locally compact and Hausdorff. We already know that X^∞ is compact, so we need to prove that it is Hausdorff. Since we already know that we can separate points in X, it remains to show that we can separate a point $x \in X$ from ∞. So let K be a compact neighbourhood of x, and U an open neighbourhood of x contained in K. Since X is Hausdorff, K is closed by Proposition 8.1.3. Therefore, $X^\infty \setminus K$ is an open neighbourhood of ∞, and it is disjoint from U by construction. $\square$

Corollary 8.5.9. *A Hausdorff space is locally compact if and only if it is an open set of a compact Hausdorff space.*

Exercise 72. Show that X is dense in X^∞ if and only if X is not compact and not empty.

The following proposition shows that the one-point compactification is in some sense "the only way" to turn a space into a compact Hausdorff space by adding a point, at least in the case of locally compact Hausdorff spaces.

Proposition 8.5.10. *Let Y be a compact Hausdorff space, and $p \in Y$ any point. Let X be the complement of p. Then the map $f : X^\infty \to Y$ defined by $f(\infty) = p$ and $f(x) = x$ for $x \in X$ is a homeomorphism.*

Proof. It is clear that f is bijective. By Corollary 8.1.8, it is enough to show that f is continuous. If U is an open subset of Y that does not contain p, then $U \subseteq X$, so $f^{-1}(U) = U$ is open in X, since X itself is open thanks to the fact that Y is Hausdorff.

If, on the other hand, $p \in U$, then $\infty \in f^{-1}(U)$, so we need to check that $X^\infty \setminus f^{-1}(U)$ is compact (since then Proposition 8.1.3 will imply that it is closed in X). But $X^\infty \setminus f^{-1}(U) = Y \setminus U$, so it is closed in Y, hence compact by Proposition 8.1.2. $\square$

Exercise 73. Show that S^n is the one-point compactification of $\mathbb{R}^n$. [*Hint: remember the stereographic projection $S^n \setminus \{N\} \to \mathbb{R}^n$*]

8.6 Example: complex projective spaces

The definition of complex projective spaces is entirely analogous to that of their real counterparts.

Definition 8.6.1. Let $\sim$ be the equivalence relation on $\mathbb{C}^{n+1} \setminus \{0\}$ where $x \sim y$ if and only if there exists $\lambda \in \mathbb{C}$ with $x = \lambda y$. The quotient space $\mathbb{C}^{n+1}/\sim$ is called the *n-dimensional complex projective space*, and is denoted $\mathbb{C}P^n$.

Note that, despite the similar definition, $\mathbb{C}P^n$ is not homeomorphic to any of the $\mathbb{R}P^m$, except when $n = m = 0$, in which case they are both the terminal space. Unfortunately, we will not be able to prove this general statement in these notes.

Nevertheless, similar results to those about real projective spaces hold in the complex case. Their proofs are completely analogous and are left to the reader.

Proposition 8.6.2. *Assume $n > 0$. Let U_0 be the set of points $[x_0, ..., x_n] \in \mathbb{C}P^n$ where $x_0 \neq 0$. Then U_0 is a dense open subspace of $\mathbb{C}P^n$ homeomorphic to $\mathbb{C}^n$, and the complement of U_0 is homeomorphic to $\mathbb{C}P^{n-1}$.*

Proposition 8.6.3. *Let $\sim$ be the equivalence relation on $S^{2n+1} \subseteq \mathbb{C}^{n+1}$ where $x \sim y$ if and only if there exists $\lambda \in S^1$ such that $x = \lambda y$. Then $\mathbb{C}P^n \cong S^{2n+1}/\sim$.*

Corollary 8.6.4. *The complex projective spaces are connected and compact Hausdorff spaces.*

Interestingly, the analogous of Proposition 5.6.4 requires a slightly different approach. The strategy that we used in the real case was to construct a map $S^1 \to S^1$ and show that it induces a homeomorphism $\mathbb{R}P^1 \cong S^1$.

In the complex case, note that puncturing $\mathbb{C}P^1$ produces $\mathbb{C}$, therefore we should expect $\mathbb{C}P^1$ to be homeomorphic to S^2. This is indeed the case, but trying to execute the above strategy would require us to find a non-trivial map $S^3 \to S^2$ that is compatible with the equivalence relation on S^3, and it is not obvious how to construct such a map.

In fact, it is easier to prove that $\mathbb{C}P^1 \cong S^2$ directly, then apply the above strategy in reverse to obtain the map $S^3 \to S^2$ above. This would also work in the real case, and we leave it to the reader to fill in the details and produce another proof of Proposition 5.6.4.

Proposition 8.6.5. $\mathbb{C}P^1 \cong S^2$.

Proof. Both spaces are homeomorphic to the one-point compactification of $\mathbb{C}$ by Proposition 8.5.10. □

The projection into the quotient $S^3 \to \mathbb{C}P^1$ can then be regarded as a map $h : S^3 \to S^2$. This map is called the *Hopf map* (or *Hopf fibration*), and it is a tool of fundamental importance in the study of algebraic invariants of spheres. Note that for any point of S^2, its fibre (i.e. inverse image) is homeomorphic to S^1.

Exercise 74. Make the homeomorphisms involved in the proof of Proposition 8.6.5 explicit, and obtain a formula for the Hopf fibration.

8.7 Example: matrix groups and quaternions

The set of $n \times m$ matrices with real coefficients is obviously in bijection with $\mathbb{R}^{nm}$, for example by regarding a matrix as the vector obtained by concatenating all its columns. There are other ways to set up such a bijection, but they all differ by a permutation of the coordinates of $\mathbb{R}^{nm}$. Since permuting coordinates is a homeomorphism $\mathbb{R}^k \to \mathbb{R}^k$, this means that there is a well-defined topology on the set of $n \times m$ matrices that makes all these bijections into homeomorphisms. Therefore, from now on we will regard the set of matrices as a topological space.

We will mostly focus on square $n \times n$ matrices, and denote them by $\mathsf{Mat}(n, \mathbb{R})$. The whole space $\mathsf{Mat}(n, \mathbb{R})$ is, by construction, just a Euclidean space, so there is nothing new there. However, by exploiting the rest of the algebraic structure on matrices we can define new interesting subspaces.

Definition 8.7.1. For any $n \geq 1$, the *real general linear group* in dimension n is the space of invertible $n \times n$ matrices with real coefficients, and is denoted $GL(n, \mathbb{R})$.

A matrix with coefficients in a field is invertible if and only if its determinant is non-zero, therefore $GL(n, \mathbb{R})$ can be regarded as the inverse image of $\mathbb{R} \setminus \{0\}$ along the function $\det : \mathsf{Mat}(n, \mathbb{R}) \to \mathbb{R}$ that computes the determinant. By writing out an explicit formula for the determinant of a matrix, one can observe that det is a polynomial function of the coefficients, therefore continuous. It follows that $GL(n, \mathbb{R})$ is an open subset of $\mathsf{Mat}(n, \mathbb{R})$.

Furthermore, since $\mathbb{R}\setminus\{0\}$ is disconnected, it follows that $GL(n, \mathbb{R})$ is also disconnected. Let $GL_+(n, \mathbb{R})$ be the open subspace of $GL(n, \mathbb{R})$ consisting of matrices with positive determinant, and $GL_-(n, \mathbb{R})$ be its complement (so, matrices with negative determinant).

Exercise 75. Prove that $GL_-(n, \mathbb{R})$ is homeomorphic to $GL_+(n, \mathbb{R})$.

Proposition 8.7.2. *$GL_+(n, \mathbb{R})$ is connected. Therefore, $GL_+(n, \mathbb{R})$ and $GL_-(n, \mathbb{R})$ are precisely the connected components of $GL(n, \mathbb{R})$.*

We will prove Proposition 8.7.2 later.

Exercise 76. Show that $GL(n, \mathbb{R})$ is *not* compact.

Definition 8.7.3. For any $n \geq 1$, the *orthogonal group* in dimension n is the space of orthogonal $n \times n$ matrices, and is denoted $O(n)$.

Recall that a matrix A is said to be *orthogonal* if $A^t A = I$, where A^t denotes the transpose of A. In other words, a matrix is orthogonal if and only if all its columns are pairwise orthogonal, and have length 1, i.e. they form an orthonormal basis of $\mathbb{R}^n$ with respect to its standard scalar product.

In particular, the homeomorphism $\mathsf{Mat}(n, \mathbb{R}) \to (\mathbb{R}^n)^n$ that sends a matrix to the sequence of its columns maps $O(n)$ into the subspace $(S^{n-1})^n$.

Proposition 8.7.4. *Through the identification above, $O(n)$ is a closed subspace of $(S^{n-1})^n$.*

Proof. The equation defining $O(n)$ can be regarded as a system of polynomial (more precisely, quadratic) equations in the coefficients of a matrix. Therefore, their solution space is the inverse image of a point along a continuous function, which is then closed in $\mathsf{Mat}(n, \mathbb{R})$, hence in $(S^{n-1})^n$. □

Corollary 8.7.5. *The orthogonal group $O(n)$ is compact.*

Proof. By Corollary 8.1.15, $(S^{n-1})^n$ is compact, and by Proposition 8.7.4, $O(n)$ is homeomorphic to a closed subspace of it, hence it is compact by Proposition 8.1.2. □

Exercise 77. Prove that $O(2) \cong S^1 \times O(1)$.

Just like $GL(n,\mathbb{R})$, $O(n)$ is not connected, since $\det(O(n)) = O(1) = \{-1, 1\}$. This motivates the following definition:

Definition 8.7.6. The *special orthogonal group* $SO(n)$ is the subspace of $O(n)$ consisting of matrices of determinant 1.

Let us regard $SO(n)$ as the subspace of $SO(n+1)$ consisting of all those matrices of the form

$$\left[\begin{array}{c|c} 1 & 0^t \\ \hline 0 & A \end{array}\right],$$

where A is an arbitrary matrix in $SO(n)$. Note that a matrix $X \in SO(n+1)$ belongs to this subspace if and only if its first column is e_1, the first vector of the standard basis, i.e. if and only if e_1 is an eigenvector of X with eigenvalue 1.

Lemma 8.7.7. *Let $n > 1$, $v \in S^{n-1}$, and H be the subspace of $SO(n)$ consisting of matrices whose first column is v. Then H is homeomorphic to $SO(n-1)$.*

Proof. First of all, H is non empty. In fact, one can extend v to a positive orthonormal basis of $\mathbb{R}^n$, and the matrix A having those vectors as columns will be in H.

Now, if B is any element of H, then $A^{-1}Be_1 = A^{-1}v = e_1$, hence $A^{-1}B \in SO(n-1)$. Conversely, if $C \in SO(n-1)$, then $AC \in H$. Therefore, multiplication by A establishes a homeomorphism between $SO(n-1)$ and H, as required. □

Proposition 8.7.8. *$SO(n)$ is connected.*

Proof. By induction on n, the base case $n = 1$ being trivial, since $SO(1) = 1$, the terminal space. Therefore, let $n > 1$, and assume that $SO(k)$ is connected for $k < n$.

Let $f : SO(n) \to S^{n-1}$ be the function that returns the first column of a matrix. The function f is clearly continuous and surjective, and Proposition 8.1.7 implies that it is closed. By Exercise 26, S^{n-1} has the final topology induced by f.

We now show that f is an open map. Let U be an open subset of $SO(n)$. Observe that

$$f^{-1}(f(U)) = \bigcup_{A \in SO(n-1)} AU,$$

where, as above, we are identifying $SO(n-1)$ with the subgroup of $SO(n)$ consisting of those matrices that fix the first vector of the standard basis.

Since multiplying by a matrix in $SO(n)$ induces a homeomorphism $SO(n) \to SO(n)$, it follows that $f^{-1}(f(U))$ can be written as a union of open sets, hence it is open, and therefore $f(U)$ is open.

Now suppose $SO(n) = U \cup V$, where U and V are non-empty open sets. We have to show that U and V meet. Since S^{n-1} is connected for $n > 1$, it follows that $f(U) \cap f(V)$ cannot be empty, hence it contains an element v.

Let X be the inverse image of v along f, i.e. the space of all matrices in $SO(n)$ such that their first column is v. By Lemma 8.7.7, $X \cong SO(n-1)$, hence X is connected by induction hypothesis. But $U \cap X$ is non-empty, since $v \in f(U \cap X)$, and similarly $V \cap X$ is non empty. It follows that $U \cap V \neq \emptyset$, as required. $\square$

As promised, we will now prove that $GL_+(n, \mathbb{R})$ is connected.

Lemma 8.7.9. *Let A be a matrix in $GL(n, \mathbb{R})$. Then there exists a matrix $B \in O(n)$ and a path in $GL(n, \mathbb{R})$ connecting A to B.*

Proof. Let $v_1, ..., v_n$ be the columns of A. The idea of the proof is that we execute Gram-Schmidt's orthogonalisation, with the addition of a parameter t that, when set to 0, will leave the matrix unchanged, and when set to 1, will produce an orthogonal matrix.

Let
$$w_i(t) = v_i - t \sum_{j<i} v_j \frac{v_j \cdot v_i}{\|v_j\|}.$$

If $Q(t)$ is the matrix whose columns are the $w_i(t)$, then Q defines a path in $\mathsf{Mat}(n, \mathbb{R})$ from $Q(0) = A$ to a matrix $Q(1)$. Now $Q(1)$ is the Gram-Schmidt orthogonalisation of A, hence $Q(1)$ can be connected to an orthogonal matrix simply by scaling every column.

It remains to show that every $Q(t)$ on the path actually belongs to $GL(n, \mathbb{R})$, i.e. that $Q(t)$ is invertible for all $0 \leq t \leq 1$. Observe that the definition of Q can be rewritten as a matrix multiplication:
$$Q(t) = AR(t)$$
where $R(t)$ is an upper triangular matrix with ones on the diagonal, hence invertible. Since $A \in GL(n, \mathbb{R})$, it follows that $Q(t)$ is also invertible. $\square$

Corollary 8.7.10. *$GL_+(n, \mathbb{R})$ is connected.*

Proof. For all $A \in GL_+(n, \mathbb{R})$, Lemma 8.7.9 implies that there exists a path connecting A to a matrix in $O(n)$. Since A has positive determinant, it follows that this path connects A to $SO(n)$. Let C_A be the union of $SO(n)$ and the image of this path.

We clearly have that $GL_+(n, \mathbb{R})$ is the union of the C_A, since $A \in C_A$, and each of the C_A is connected by Proposition 6.1.9, hence $GL_+(n, \mathbb{R})$ is connected by Proposition 6.1.9 again. $\square$

Exercise 78. Show that $GL_+(n, \mathbb{R})$ is also path connected. [*Hint: use that it is an open subset of Euclidean space*]. Deduce that $SO(n)$ is also path connected.

We can give similar definitions of matrix groups in the complex case. Of course, the starting point is $\mathsf{Mat}(n, \mathbb{C})$, which gets its topology from $\mathbb{C}^{n^2}$, or alternatively, $\mathbb{R}^{2n^2}$.

Definition 8.7.11. The *complex general linear group* is the subspace of $\mathsf{Mat}(n, \mathbb{C})$ consisting of invertible matrices, and is denoted $GL(n, \mathbb{C})$.

Definition 8.7.12. The *unitary group* is the subspace of $GL(n, \mathbb{C})$ consisting of unitary matrices, i.e. those matrices A such that $A^*A = I$, where A^* is the transpose conjugate of A. It is denoted $U(n)$.

Definition 8.7.13. The *special unitary group* is the subspace of $U(n)$ consisting of those matrices that have determinant 1, and is denoted $SU(n)$.

We can immediately observe a difference with the linear case, in that the determinant of an invertible matrix with complex coefficient is an element of $GL(1, \mathbb{C}) = \mathbb{C} \setminus \{0\}$, which is connected. Therefore, we cannot conclude that $GL(n, \mathbb{C})$ is disconnected as we could for $GL(n, \mathbb{R})$. In fact, we have the following:

Proposition 8.7.14. *$GL(n, \mathbb{C})$, $U(n)$ and $SU(n)$ are path connected for all $n \geq 1$.*

The proof uses the same ideas as the one of Proposition 8.7.8, so it is left as an exercise for the reader. We have similar results about compactness.

Proposition 8.7.15. *$U(n)$ and $SU(n)$ are compact. $GL(n, \mathbb{C})$ is not compact.*

Not all of the matrix groups are "new" spaces. Among the trivial examples, $O(1)$ is a discrete space with two points, and $SO(1)$ and $SU(1)$ are the terminal space. We have already observed that $O(2) \cong S^1 \times O(1)$, from which it follows that $SO(2) \cong S^1$. Finally, it is immediate from the definition that $U(1)$ is also homeomorphic to S^1.

Let us now turn our attention to one of the higher dimensional examples.

Recall that the *quaternion algebra* $\mathbb{H}$ is the 4-dimensional real algebra generated by elements i, j, k, subject to the relations $ij = k$, $jk = i$, $ki = j$, $i^2 = j^2 = k^2 = -1$. Note that $\mathbb{H}$ is not commutative. We regard $\mathbb{H}$ as a topological space by identifying it with $\mathbb{R}^4$.

Quaternions (i.e. elements of $\mathbb{H}$) have a conjugation operation $* : \mathbb{H} \to \mathbb{H}$, defined by $(w+ix+jy+kz)^* = w-ix-iy-iz$. If $q = w+ix+jy+kz$ is any quaternion, then its *norm* q^*q is a positive real number equal to $w^2 + x^2 + y^2 + z^2$, hence it is the Euclidean norm of q when regarded as a vector in $\mathbb{R}^4$. In particular, we can identify S^3 with the set of *unit* quaternions, i.e. quaternions of norm 1.

It is easy to see that $\mathbb{H}$ is isomorphic to the subalgebra of 2×2 matrices with complex coefficients of the form $\begin{bmatrix} a & -b^* \\ b & a^* \end{bmatrix}$. The norm of the corresponding quaternion is $|a|^2 + |b|^2$, hence such a matrix is in $SU(2)$ if and only if it corresponds to a unit quaternion. So we have proved the following:

Proposition 8.7.16. *$SU(2) \cong S^3$.*

There is also an important relation between unit quaternions and $SO(3)$, i.e. rotations of 3-dimensional Euclidean space, which the next exercises explore.

Exercise 79. Let V be the 3-dimensional real subspace of $\mathbb{H}$ generated by i, j, k. Show that for any unit quaternion q, the function $x \mapsto qxq^*$ restricts to an *orthogonal* linear automorphism of V, when regarded as a subspace of $\mathbb{R}^4$ with the Euclidean metric. Use this fact to construct a group homomorphism $SU(2) \to SO(3)$ and prove that it is continuous.

Exercise 80. Show that a unit quaternion q determines the identity rotation of V if and only if $q = 1$ or $q = -1$. Deduce that two quaternions determine the same rotation if and only if they correspond to antipodal points on S^3.

Exercise 81. Let A be a matrix in $SO(3)$. Define:

$$\begin{aligned} w &= \tfrac{1}{2}\sqrt{1 + \mathsf{tr}(A)} \\ x &= \tfrac{1}{4w}(A_{32} - A_{23}) \\ y &= \tfrac{1}{4w}(A_{13} - A_{31}) \\ z &= \tfrac{1}{4w}(A_{21} - A_{12}). \end{aligned}$$

Show that the quaternion $q = w + ix + jy + kz$ determines the rotation corresponding to A.

Exercise 82. Use the previous exercises to construct a homeomorphism $SO(3) \cong \mathbb{R}P^3$.

Quaternions can be used as a base to build projective spaces, just like the reals and the complex numbers. We will not pursue that here, but suggest the following exercises for the interested readers.

Exercise 83. Give a definition of the *quaternionic projective space* $\mathbb{H}P^n$. [*Hint: just like for the real and complex case, define an equivalence relation on* $\mathbb{H}^{n+1}$*, and then quotient by it*]

Exercise 84. Prove that $\mathbb{H}P^n$ is homeomorphic to a quotient of S^{4n+3}. Deduce that $\mathbb{H}P^n$ is compact.

Exercise 85. Prove that $\mathbb{H}P^1 \cong S^4$. Use this homeomorphism to construct a map $S^7 \to S^4$ (the *quaternionic Hopf map*) with fibre homeomorphic to S^3.

9 Elements of algebraic topology

With the topological invariants introduced so far, we have been able to show that certain pairs of spaces are not homeomorphic. However, those invariants alone are not sufficient to distinguish many of the common examples we have seen, including Euclidean spaces of different dimensions, the various spheres, and two-dimensional surfaces.

For example, we have seen that it is easy to prove that $\mathbb{R}$ is not homeomorphic to $\mathbb{R}^2$ by removing one point from each (we say *puncturing them*) and observing that they cannot be homeomorphic, because the latter is (path) connected and the former is not. However, the argument does not generalise to higher Euclidean spaces, because a punctured $\mathbb{R}^n$ and a punctured $\mathbb{R}^{n-1}$ are both path connected when $n > 1$.

One way to understand why what we have so far is not enough to solve this seemingly very simple problem is to think back about the proof technique that we used. We can conceptualise the above argument as follows: to show that two spaces are not homeomorphic, assign a certain "object" to every space in such a way that homeomorphic spaces get isomorphic objects, then compute that object for the two spaces of interest, and show that they are not isomorphic. Such an object will be called an "invariant".

In our case, the invariant was the set π_0 of path connected components. Clearly, homeomorphic spaces have isomorphic sets of path connected components, therefore the argument goes through. However, π_0 alone is not enough to distinguish punctured Euclidean spaces of higher dimensions.

This suggests that we need more sophisticated invariants for higher dimensions. As it turns out, most of those invariants have the additional benefit of carrying an algebraic structure. Therefore, the study of these invariants and the techniques needed to compute them is called *algebraic topology*, since it uses ideas from both topology and algebra.

In the rest of these notes, we will introduce the most elementary such invariant, the *fundamental group* π_1, and we will use it to make some progress in the problem introduced above. Namely, we will show that $\mathbb{R}^2$ is not homeomorphic to $\mathbb{R}^n$ if $n \neq 2$. To completely solve the problem of distinguishing Euclidean spaces, we would need higher analogues to the fundamental group, called *higher homotopy groups* and denoted by π_n, but they are unfortunately out of the scope of these notes.

We will, however, apply the fundamental group to other problems, and obtain important results such as Brouwer's fixed point theorem as a simple consequence of the "functoriality" of π_1.

9.1 Homotopies and the fundamental group

In section 6.2 we have introduced the notion of *path* in a topological space, and constructed path concatenation as a map that takes two "consecutive" paths on a space and produces a new one. Since paths can only be concatenated if their endpoints are compatible, meaning that the second endpoint of the first

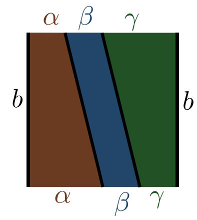

Figure 8: Associativity of path concatenation

path coincides with the first endpoint of the second path, the concatenation operation does not quite give us an algebraic structure on paths.

One way around this difficulty is restricting our attention to *loops*, i.e. paths that begin and end at the same point. Since all loops around a given *base point* b can be concatenated, we immediately get the beginning of an algebraic structure on paths. For this reason, we will often consider spaces equipped with a choice of base point. The base point is often arbitrary, and we will see that many constructions do not depend on the specific choice of the base point (only on its connected component), but having one always available simplifies certain statements.

Definition 9.1.1. A *pointed (topological) space* is a topological space X equipped with a distinguished point $b \in X$, called the *base point*.

There is also a second difficulty to take care of. Namely, the operation of concatenation of loops is not associative, and it does not possess a neutral element, making it difficult to deal with algebraically. To see why associativity fails, let us fix three loops α, β, γ on a pointed space X, and consider the two possible way to concatenate them in that order: $\alpha * (\beta * \gamma)$ and $(\alpha * \beta) * \gamma$.

To compute the former path, we first concatenate β and γ, giving rise to a path that traces β at double speed on the interval $[0, \frac{1}{2}]$, followed by γ, also at double speed, on the interval $[\frac{1}{2}, 1]$.

Then, we add α at the beginning, which means that now α will be traced at double the speed on the interval $[0, \frac{1}{2}]$, while β and γ will both be traced at four times the original speed, on the intervals $[\frac{1}{2}, \frac{3}{4}]$ and $[\frac{3}{4}, 1]$ respectively.

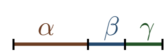

On the other hand, the other way of composing the three paths yields a different parameterisation of the result, where α and β are now traced at four times the speed on the intervals $[0, \frac{1}{4}]$ and $[\frac{1}{4}, \frac{1}{2}]$ respectively, while γ is on $[\frac{1}{2}, 1]$ going at twice the speed.

Therefore, the two resulting paths are clearly not equal. On the other hand, it is pretty clear that it is possible to convert one into the other in a continuous way. This can be realised by defining a continuous function from the square $[0, 1] \times [0, 1]$ into X such that two opposite sides are mapped to the two composed loops, while the other two opposite sides stay constant on the base point b.

This function can be described explicitly as follows:

$$h(s,t) = \begin{cases} \alpha((2t+2)s) & \text{for } 0 \le s \le \frac{1}{2t+2} \\ \beta(4s+t-2) & \text{for } \frac{1}{2t+2} \le s \le \frac{t+3}{4t+4} \\ \gamma((4-2t)s+2t-3) & \text{for } \frac{t+3}{4t+4} \le s \le 1. \end{cases}$$

This explicit formula can be obtained by computing the expressions for the two loops $\alpha * (\beta * \gamma)$ and $(\alpha * \beta) * \gamma$ given by repeated applications of Definition 6.2.8, then linearly interpolating the two formulas with the additional variable t. Continuity follows immediately from Proposition 5.1.4.

The above observation motivates the following definition.

Definition 9.1.2. Let $f, g : X \to Y$ two continuous functions between topological spaces. A *homotopy* between f and g is a continuous function

$$h : X \times [0,1] \to Y,$$

such that for all $x \in X$, $h(x,0) = f(x)$ and $h(x,1) = g(x)$. The two functions f and g are said to be *homotopic*, and we write $f \sim g$, if there exists a homotopy between them.

The function h defined above is then a homotopy between the loops $\alpha * (\beta * \gamma)$ and $(\alpha * \beta) * \gamma$. Furthermore, $h(0,s) = h(1,s) = b$ for all $s \in [0,1]$. We abstract this property with a new definition.

Definition 9.1.3. Let $f, g : X \to Y$ be continuous functions, h a homotopy between them, and A a subspace of X. We say that h is a homotopy *relative to* A if $h(a,t)$ is constant in t for all points $a \in A$. When such a homotopy exists, we say that f and g are *homotopic relative to* A, and write $f \sim g$ (rel A).

Note that if f and g are homotopic relative to A, then in particular they restrict to the same function on A. Also note that the absolute version of homotopy given in Definition 9.1.2 can be thought of as homotopy relative to the empty subspace of X, hence we only need to concern ourselves with the relative case when proving properties of homotopies.

Given a homotopy h between f and g, and a fixed $t \in [0,1]$, we will sometimes write h_t to denote the function $X \to Y$ given by $h_t(x) = h(x,t)$. Clearly, h_t is a continuous function, $h_0 = f$ and $h_1 = g$. Dually, if we fix a point $x \in X$, we can define a path h^x in Y connecting $f(x)$ and $g(x)$.

Since we have proved that the two possible ways to compose three loops are homotopic relative to $\{0,1\}$, this suggests a way to resolve the problem with associativity of path concatenation.

Proposition 9.1.4. *Let X, Y be topological spaces, and A a subspace of X. The relation of homotopy relative to A is an equivalence relation on the set of continuous functions $X \to Y$.*

Proof. First, it is clear that the homotopy relation is reflexive, since given any continuous function $u : X \to Y$, a homotopy h relative to A between u and itself is simply given by $h(x,t) = u(x)$.

As for symmetry, if h is a relative homotopy between u and v, we define a homotopy k between v and u as $k(x,t) = h(x, 1-t)$.

Finally, let h be a homotopy between u and v, and k a homotopy between v and w. Define a new homotopy j by setting $j(x,t) = (h^x * k^x)(t)$. By expanding the definition of path concatenation, it is immediate to check that j is continuous, hence it defines a homotopy between u and v. □

Exercise 86. Let $f, g : X \to Y$ be continuous functions, A a subspace of X, h a homotopy between f and g relative to A, and $u : Y \to Z$ any continuous function. Show that $u \circ h$ is a homotopy between $u \circ f$ and $u \circ g$ relative to A.

Definition 9.1.5. Let X be a pointed space. Define $\pi_1(X)$ to be the quotient of the set of loops on X by the equivalence relation of homotopy relative to $\{0,1\}$.

Note that X in Definition 9.1.5 is a *pointed* space, i.e. it comes with a choice of a specified base point. When we want to be more explicit, we can denote such a pointed space as a pair (X, b), where X is a topological space and $b \in X$, and correspondingly write $\pi_1(X, b)$.

Exercise 87. Prove that two paths in any topological space X are homotopic relative to $\{1\}$ if and only if they have the same endpoint. Prove that if X is path connected, then any two paths are homotopic (relative to empty set).

Exercise 87 shows that homotopies between paths that are not relative to $\{0,1\}$ do not carry much information on the space. For this reason, we shall simply say that two paths are homotopic to mean that they are homotopic relative to $\{0,1\}$, unless otherwise specified.

Proposition 9.1.6. *For any pointed space X, path concatenation induces an associative binary operation on $\pi_1(X)$.*

Proof. First, we show that path concatenation is well-defined up to homotopy. That is, given loops α, α', β, if α and α' are homotopic, then so are $\alpha * \beta$ and $\alpha' * \beta$. The analogous result for the right side can of course be proved similarly. So let h be a homotopy between α and α', and define a function k as $k(s,t) = (h_t * \beta)(s)$. Clearly k is continuous, and $k_0 = h_0 * \beta = \alpha * \beta$, while $k_1 = h_1 * \beta = \alpha' * \beta$, so k is the desired homotopy.

Therefore, we can regard path concatenation as a function on homotopy classes of loops, i.e. as a binary operation:

$$* : \pi_1(X) \times \pi_1(X) \to \pi_1(X),$$

and the above explicit construction of the homotopy between $\alpha * (\beta * \gamma)$ and $(\alpha * \beta) * \gamma$ shows that it is associative. □

The set $\pi_1(X)$, equipped with the operation induced from path concatenation, can actually be shown to be a group, called the *fundamental group* of the pointed space X. The following two exercises spell out why $\pi_1(X)$ is indeed a group.

Exercise 88. If (X, b) is a pointed space, let e be the *constant (or trivial) loop* at the base point, i.e. the path defined by $e(s) = b$. Show that the homotopy class of e is the neutral element of the operation on $\pi_1(X, b)$.

Exercise 89. Let α be a loop on a pointed space X. Define the *inverse* of α to be the path α^{-1} given by $\alpha^{-1}(s) = \alpha(1-s)$. Show that $\alpha * \alpha^{-1}$ is homotopic to the constant loop e, and that $\left(\alpha^{-1}\right)^{-1} = \alpha$. Deduce that the homotopy class of α^{-1} in $\pi_1(X)$ is an inverse for the homotopy class of α.

9.2 Homotopy equivalence

We will now attempt to compute fundamental groups in some very simple cases.

Proposition 9.2.1. *Let X be a convex subset of $\mathbb{R}^n$, and $b \in X$. Then $\pi_1(X, b)$ is the trivial group.*

Proof. If α is any loop in X, we can show that α is homotopic to the trivial loop. In fact, let h be the function defined by $h(s,t) = (1-t)\alpha(s) + tb$. The function h is clearly continuous, and it lies in X thanks to convexity. Furthermore, $h_0 = \alpha$ and h_1 is the trivial loop.

This shows that there is only one homotopy class of loops in X, hence $\pi(X, b)$ must be the trivial group. ☐

We can generalise the example of a convex space by abstracting out the proof of Proposition 9.2.1.

Definition 9.2.2. A pointed space (X, b) is said to be *contractible* if the identity function $X \to X$ is homotopic to the constant function with value b.

More explicitly, a (X, b) is contractible if there exists a continuous function $h : X \times [0,1] \to X$ such that $h(x, 0) = x$ and $h(x, 1) = b$. Note that a pointed convex subset of $\mathbb{R}^n$ is contractible, since in that case we can define $h(x,t) = (1-t)x + tb$.

Exercise 90. Prove that contractibility of a pointed space does not depend on the choice of the base point.

Because of Exercise 90, we will often speak of a contractible topological space, without reference to a base point. Note, however, that for a space to be contractible there must exists at least *one* choice of base point, which means that contractible spaces are in particular non-empty.

Exercise 91. Prove that a contractible space is path connected.

Definition 9.2.3. A pointed space X is said to be *simply connected* if it is path connected, and $\pi_1(X)$ is the trivial group.

Exercise 92. Show that X is simply connected if and only if every two paths in X with the same endpoints are homotopic (relative to $\{0,1\}$).

Exercise 93. Let X be a path connected topological base, and b, b' points in X. Prove that (X, b) is simply connected if and only if (X, b') is.

Because of Exercise 93, we will, just like for contractible spaces, speak of a simply connected space without necessarily a reference to a base point. Explicitly, a topological space X is said to be simply connected if it is path connected, and for all (or any) of its points b, the pointed spaced (X, b) are simply connected.

Lemma 9.2.4. *Let h be a homotopy between two paths γ and σ, not necessarily relative to the endpoints. Let $\alpha = h^0$ and $\beta = h^1$. Then σ and $\alpha^{-1} * \gamma * \beta$ are homotopic relative to the endpoints.*

We give two proofs of Lemma 9.2.4. The first one constructs the homotopy explicitly:

Proof. If ρ is any path, let $\hat{\rho}$ be defined by $\hat{\rho}(s,t) = \rho(\mathsf{max}\{s,t\})$. Then $\hat{\rho}$ is a homotopy between ρ and the trivial loop at $\rho(1)$, relative to $\{1\}$. Define a homotopy k such that

$$k_t := (\hat{\alpha}_t)^{-1} * h_t * \hat{\beta}_t.$$

It is easy to see that k is relative to $\{0,1\}$. Furthermore $k_0 = \alpha^{-1} * \gamma * \beta$, and $k_1 = \sigma$. □

The second proof uses the contractibility of convex subsets of $\mathbb{R}^n$:

Proof. Let l, r, t, b be the left, right, top and bottom sides of the square $Q := [0,1] \times [0,1]$ respectively. Explicitly: $l(t) = (0,t)$, $r(t) = (1,t)$, $t(s) = (s,1)$, $b(s) = (s,0)$. Then clearly $\gamma = h \circ b$, $\sigma = h \circ t$, $\alpha = h \circ l$ and $\beta = h \circ l$. Note that $l^{-1} * b * r * t^{-1}$ is a loop in Q, hence it is homotopic to a trivial loop. It follows that $t \sim l^{-1} * b * r$ (rel $\{0,1\}$), and therefore $\sigma \sim \alpha^{-1} * \gamma * \beta$ (rel $\{0,1\}$) by Exercise 86. □

Proposition 9.2.5. *If X is a contractible space, then X is simply-connected.*

Proof. Let h be a homotopy between the identity of X and a constant. If γ is a loop in X, define k by $k(s,t) = h(\gamma(s), t)$. Then k is a homotopy between γ and the trivial loop, but not necessarily relative to the endpoints. If $\alpha = k^0$ and $\beta = k^1$, we get by Lemma 9.2.4 that γ is homotopic to $\alpha^{-1} * \beta$. Since γ is a loop, $\alpha = \beta$, therefore γ is homotopic to a trivial loop. □

Contractibility captures the idea that a space X is essentially the same as a single point, at least when reasoning up to homotopy, like when calculating π_1. We will now generalise this idea to other pairs of spaces.

Definition 9.2.6. Let $f : X \to Y$ be a continuous function between topological spaces. We say that f is a *homotopy equivalence* if there exists a continuous function $g : Y \to X$ such that $g \circ f$ is homotopic to the identity of X, and $f \circ g$ is homotopic to the identity of Y. We say that X and Y are *homotopy equivalent* if there exists a homotopy equivalence between them.

The notion of homotopy equivalence is a straightforward generalisation of the notion of homeomorphism, where we have replaced equality by homotopy. In particular, homeomorphic spaces are homotopy equivalent, since homotopy is a reflexive relation.

The advantage of introducing a more relaxed notion of equivalence between spaces is that it turns out that in many situations homotopy equivalence is enough to distinguish spaces. If we know that two spaces are not homotopy equivalent, then this is sufficient to conclude that they are not homeomorphic. The reason why this is useful in practice is that topological invariants such as the fundamental group turn out to be also invariant under homotopy equivalence.

To see why this is the case, let us first examine the interaction between arbitrary continuous functions and the fundamental group.

Definition 9.2.7. Let X and Y be pointed topological spaces, with base points x and y respectively. A *pointed map* $X \to Y$ is a continuous function $f : X \to Y$ such that $f(x) = y$.

It is clear that the identity of any topological space X can be regarded as a pointed map for any choice of a base point. Also, pointed maps can be composed in the obvious way.

Proposition 9.2.8. *Let $f : X \to Y$ be a pointed map between pointed spaces. The map on loops $\gamma \mapsto f \circ \gamma$ determined by f induces a group homomorphism $f_* : \pi_1(X) \to \pi_1(Y)$.*

Proof. If γ and ρ are homotopic loops, then $f \circ \gamma$ and $f \circ \rho$ are homotopic by Exercise 86, hence the assignment $\gamma \mapsto f \circ \gamma$ determines a function $f_* : \pi_1(X) \to \pi_1(Y)$, so we only need to verify that it is a homomorphism of groups.

In fact, if e denotes the constant loop (on any base point), we have $f \circ e = e$ and $f \circ (\gamma * \rho) = (f \circ \gamma) * (f \circ \rho)$ directly from the definitions. Then, in particular, if we write $[\gamma]$ for the homotopy class of the loop γ in π_1, we get $f_*([e]) = [e]$ and $f_*([\gamma][\rho]) = f_*([\gamma]) f_*([\rho])$, as required. □

Exercise 94. Show that the assignment $f \mapsto f_*$ defined in Proposition 9.2.8 satisfies the following *functoriality* properties:

$$\begin{aligned} \mathsf{id}_* &= \mathsf{id}, \\ (g \circ f)_* &= g_* \circ f_*. \end{aligned}$$

One can summarise Exercise 94, together with the various constructions that lead to it, by saying that π_1 is a *functor* from the category of pointed topological spaces to the category of groups.

We have constructed the induced homomorphism f_* for pointed maps f. However, if $f : X \to Y$ is just a continuous function, we can still consider f_* once we fix a base point x_0 for X, by regarding f as a pointed map $(X, x_0) \to (Y, f(x_0))$, so that f_* is a homomorphism $\pi_1(X, x_0) \to \pi_1(Y, f(x_0))$.

Proposition 9.2.9. *Let α be a path in X connecting x_1 with x_2. Then the assignment $\gamma \mapsto \alpha^{-1} * \gamma * \alpha$ induces an isomorphism $\phi_\alpha : \pi_1(X, x_1) \to \pi_1(X, x_2)$, called the* change of base point *isomorphism.*

Proof. If γ and γ' are homotopic loops on x_1, then it is easy to see that $\alpha^{-1} * \gamma * \alpha$ and $\alpha^{-1} * \gamma * \alpha$ are homotopic as well. Therefore, the assignment $\gamma \mapsto \alpha^{-1} * \gamma * \alpha$ uniquely determines a function $\phi_\alpha : \pi_1(X, x_1) \to \pi_1(X, x_2)$. To show that it is a group homomorphism, we simply calculate:

$$\begin{aligned}
&\phi_\alpha([e]) = [\alpha^{-1} * e * \alpha] = [\alpha^{-1} * \alpha] = [e] \\
&\phi_\alpha([\gamma][\rho]) = [\alpha^{-1} * \gamma * \rho * \alpha] = [\alpha^{-1} * \gamma * \alpha * \alpha^{-1} * \rho * \alpha] = \phi_\alpha([\gamma])\phi_\alpha([\rho]).
\end{aligned}$$

□

Corollary 9.2.10. *If X is a path connected topological space, and x_1, x_2 two points in X, then $\pi_1(X, x_1) \cong \pi_1(X, x_2)$.*

Note, however, that the isomorphism between the two fundamental groups in Corollary 9.2.10 is not uniquely determined, since it can depend on the choice of a path between x_1 and x_2. For this reason, even when dealing with path connected spaces, it is often best to make the choice of a base point explicit.

Lemma 9.2.11. *Let $f, g : X \to Y$ be homotopic maps between topological spaces, and let x_0 be a base point for X. Then there exists an isomorphism $\phi : \pi_1(Y, f(x_0)) \to \pi_1(Y, g(x_0))$ such that $g_* = \phi \circ f_*$, i.e. the following diagram is commutative:*

$$\begin{tikzcd}
& \pi_1(X, x_0) \arrow[dl, "f_*"'] \arrow[dr, "g_*"] & \\
\pi_1(Y, f(x_0)) \arrow[rr, "\phi"'] & & \pi_1(Y, g(x_0)).
\end{tikzcd}$$

Proof. Let h be a homotopy between f and g. Then in particular h determines a path $\alpha = h^{x_0}$ between $f(x_0)$ and $g(x_0)$, hence we get a change of base isomorphism $\phi_\alpha : \pi_1(Y, f(x_0)) \to \pi_1(Y, g(x_0))$.

Now, if γ is any loop at x_0, define $k(s, t) = h(\gamma(s), t)$. Then k is a homotopy between $f \circ \gamma$ and $g \circ \gamma$, but not necessarily relative to the endpoints. Note, however, that $\alpha = k^0 = k^1$. Lemma 9.2.4 then implies that $g \circ \gamma = \alpha^{-1} * (f \circ \gamma) * \alpha$, and therefore $g_*([\gamma]) = \phi_\alpha(f_*([\gamma]))$, as required. □

Lemma 9.2.12. *Let $f : A \to B$, $g : B \to C$, $h : C \to D$ be functions. If $g \circ f$ is invertible and $h \circ g$ is invertible, then all of f, g and h are also invertible.*

Proof. Since $g \circ f$ is surjective, g is also surjective. Since $h \circ g$ is injective, g is also injective. Therefore, g is invertible. But then $f = g^{-1} \circ (g \circ f)$ must also be invertible, and similarly for h. □

Corollary 9.2.13. *If $f : X \to Y$ be a homotopy equivalence between spaces, and x_0 any point in X. Then $f_* : \pi_1(X, x_0) \to \pi_1(Y, f(x_0))$ is an isomorphism of groups.*

Proof. Let $g : Y \to X$ such that $g \circ f \sim \mathsf{id}$ and $f \circ g \sim \mathsf{id}$. It then follows from Lemma 9.2.11 that $g_* \circ f_* = (g \circ f)_* = \phi \circ \mathsf{id} = \phi$, where $\phi : \pi_1(X, x_0) \to \pi_1(X, g(f(x_0)))$ is an isomorphism. In particular, $g_* \circ f_*$ is an isomorphism, and similarly $f_* \circ g_*$. It follows that f^* is an isomorphism by Lemma 9.2.12. □

9.3 Covering spaces

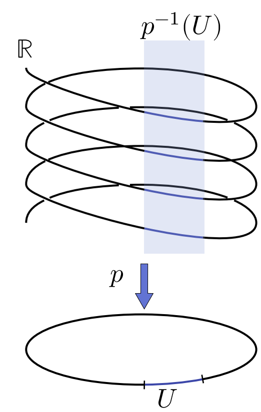

Let $p : \mathbb{R} \to S^1$ be the function defined by $p(t) = e^{2\pi i t}$. It follows from the property of the complex exponential that p is surjective and periodic of period 1, and that its restriction to any open interval $]a, b[$ with $b - a \leq 1$ is a homeomorphism with its image.

In fact, more is true: for any point ξ in the image S^1, there exists an open neighbourhood U of ξ whose inverse image along p is an open set of $\mathbb{R}$ homeomorphic to $U \times \mathbb{Z}$, with the function p acting as the first projection. Indeed, we can take U to be any open neighbourhood of ξ that is not the entire circle.

We will abstract this idea into the notion of a *covering map*. The underlying principle is that we can think of a map like $p : \mathbb{R} \to S^1$ as a "family" of spaces indexed by points of S^1. For each point $\xi \in S^1$, the corresponding set is the fibre of p over it, i.e. the inverse image along p of $\{\xi\}$, which is this case is a copy of the integers $\mathbb{Z}$.

Definition 9.3.1. Let $p : E \to X$ be a continuous function between topological spaces. We say that p is a *covering map* if every point of X is contained in an open set U such that the inverse image of U along p is a union:

$$p^{-1}(U) = \bigcup_{\alpha \in F} V_\alpha,$$

where F is some non-empty set, the V_α are disjoint open subsets of E, and p induces a homeomorphism of V_α with U.

Such an open set U is said to be *evenly covered* by p.

Exercise 95. Show that a covering map is surjective and open.

Exercise 96. Show that a homeomorphism is a covering map.

If $p : E \to X$ is a covering map, E is said to be a *covering space* of X. We can think of such a p as representing a "locally constant" family of discrete spaces over X, whose union is E. On every evenly covered open set U in X, the family is constantly equal to some set F.

In fact, if $x \in U$, then $p^{-1}(x)$ consists of a point in V_α for every element $\alpha \in F$, and each of these points is the intersection of $p^{-1}(x)$ with one of the V_α, hence it is open in $p^{-1}(x)$. So $p^{-1}(x)$ is discrete and in bijection with F, hence it is homeomorphic to F regarded as a discrete space. The following proposition gives an alternative characterisation of the condition for an open set U to be evenly covered, which is sometimes easier to work with.

Proposition 9.3.2. *Let $p : E \to X$ be a continuous function between topological spaces. An open set $U \subseteq X$ is evenly covered by p if and only if there exists a discrete space F, and a homeomorphism $\phi : U \times F \to p^{-1}(U)$ such that the following diagram is commutative:*

$$\begin{tikzcd} U \times F \arrow[rr, "\phi"] \arrow[dr, "\pi"'] & & p^{-1}(U) \arrow[dl, "p"] \\ & U, & \end{tikzcd}$$

where $\pi : U \times F \to U$ is the first projection.

Proof. First of all, it is clear that U is evenly covered by the first projection map $U \times F \to U$ whenever F is discrete. Therefore, if U satisfies the condition above, it follows from the commutativity of the diagram that U is evenly covered by p as well.

Conversely, suppose U is an evenly covered open subset of X, so that $p^{-1}(U) = \bigcup_{x \in F} V_x$ for a non-empty family of disjoint open subsets V_x of E. Now define a map $\phi : p^{-1}(U) \to U \times F$ as follows: on every V_α, simply take $p|V_\alpha : V_\alpha \to U$, and glue all of those into a single map using Proposition 5.1.3. Conversely, define $\psi : U \times F \to p^{-1}(U)$ as $\psi(u, \alpha) = (p|V_\alpha)^{-1}(u)$, which is well-defined and continuous because $p|V_\alpha$ is required to be a homeomorphism. It is clear that ϕ and ψ are inverses, hence they define the required homeomorphism. The commutativity of the diagram above is immediate. $\square$

Corollary 9.3.3. *The first projection $X \times F \to X$ is a covering map, for all topological spaces X and discrete spaces F.*

The following proposition shows that the fibre of a covering map to a connected space does not depend on the point.

Proposition 9.3.4. *Let X be a connected space, and $p : E \to X$ a covering map. Then there exists a discrete space F such that $p^{-1}(x) \cong F$ for all $x \in X$.*

Proof. Let $x_0 \in X$ be an arbitrary point, and F be the corresponding fibre. Let U be the set of points x in X such that the fibre over x is homeomorphic to F. If $x \in U$, there exists an open neighbourhood V of x that is evenly covered by p. It follows that the fibre of p on V is always homeomorphic to $p^{-1}(x)$, hence $V \subseteq U$. Therefore, U is open. On the other hand, if $x \notin U$, then by the same argument there exists an open neighbourhood V of x such that the fibre over V is *not* homeomorphic to F, therefore U is closed. Since U contains x_0, it cannot be empty, hence by connectedness of X it must be equal to X itself, as claimed. □

It is very important to note that, although Proposition 9.3.4 says that the fibre of a covering map $p : E \to X$ are in a sense always constant if X is connected, it does *not* imply that the covering space can be thought of as a projection map from a product space $X \times F$. Intuitively, we can always find a homeomorphism between the fibre over any point $x \in X$ and F, but this homeomorphism cannot necessarily be chosen "in a continuous way" as x varies in X, only for x in an evenly covered open set.

The importance of covering spaces lies in their connection with the fundamental group. For example, the covering map $p : \mathbb{R} \to S^1$ defined above makes it easy to calculate the fundamental group of S^1. In the following, we will regard the circle as a pointed space by fixing $1 \in \mathbb{C}$ as its base point. We first prove that covering maps have important *lifting properties* of paths and homotopies.

Proposition 9.3.5. *Let $p : E \to X$ be a covering map, U an evenly covered open set in X and γ a path in U starting at x_0. For any $\tilde{x}_0 \in E$ such that $p(\tilde{x}_0) = x_0$, there exists a unique path $\tilde{\gamma}$ in E starting at $\tilde{x}_0$ and such that $p \circ \tilde{\gamma} = \gamma$. Such a path $\tilde{\gamma}$ is called a* lift *of γ.*

Proof. Since U is evenly covered, it is enough by Proposition 9.3.2 to find a lift for the first projection $U \times F \to U$ starting from some point $\tilde{x}_0 \in U \times F$ which is mapped to x_0. Such a point must then be of the form (x_0, α) for some $\alpha \in F$. Now let $\pi' : U \times F \to F$ be the second projection. If $\tilde{\gamma}$ is a lift of γ, then $\pi' \circ \tilde{\gamma}$ must be constant, because $[0, 1]$ is connected. Therefore, $\pi'(\tilde{\gamma}(t)) = \alpha$ for all $t \in [0, 1]$. It follows that $\tilde{\gamma}(t) = (\gamma(t), \alpha)$, proving uniqueness. As for existence, the equation above defines a continuous path $\tilde{\gamma}$ on $U \times F$ which is clearly mapped onto γ. □

Exercise 97. The previous result can be generalised to the situation where we have a continuous function $f : B \to X$ for some connected space B, and a "partial lift" of f on a connected subspace A of B, i.e. a continuous function $g : A \to E$ such that $p \circ g = f|A$. Use a similar argument as in the proof of Proposition 9.3.5 to construct a function $h : B \to E$ that extends g and such that $p \circ h = f$.

Proposition 9.3.6. *Let $p : E \to X$ be a covering map and γ a path in X starting at x_0. For any $\tilde{x}_0 \in E$ such that $p(\tilde{x}_0) = x_0$, there exists a unique lift $\tilde{\gamma}$ of γ starting at $\tilde{x}_0$.*

Proof. Let $(U_j)_{j \in J}$ be a cover of X consisting of evenly covered open sets. Since $[0, 1]$ is compact, there exists a finite sequence $0 = t_0 < t_1 < \cdots < t_n = 1$ such that every interval $[t_i, t_{i+1}]$ is mapped into one of the U_j by γ. We prove by induction on i that the restriction of γ to $[0, t_i]$ has a unique lift $\tilde{\gamma}$ such that $\tilde{\gamma}(0) = \tilde{x}_0$. This is clearly true for $i = 0$. As for the inductive step, assume we know that the restriction of γ to $[0, t_i]$ has a unique lift $\tilde{\gamma}$, for $i < n$. Note that the restriction of γ to $[t_i, t_{i+1}]$ lies entirely into an evenly covered open set, therefore we can find a unique lift ρ of $\gamma|[t_i, t_{i+1}]$ such that $\rho(t_i) = \tilde{\gamma}(t_i)$ by Proposition 9.3.5. This implies that we can extend $\tilde{\gamma}$ to $[0, t_{i+1}]$, and the uniqueness of ρ implies that the extension is unique, concluding the proof. $\square$

What makes covering spaces useful in the calculation of homotopy groups is that we can use them to lift not just paths, but also homotopies.

Proposition 9.3.7. *Let $p : E \to X$ be a covering map, $h : B \times [0,1] \to X$ a homotopy between functions $f, g : B \to X$, and $\tilde{f}$ a lift of f, i.e. a function $B \to E$ such that $p \circ \tilde{f} = f$. Then there exists a unique lift of h to a homotopy $\tilde{h}$ in E such that $\tilde{h}_0 = \tilde{f}$.*

Proof. Let $b \in B$ be any point. One can find a finite sequence $0 = t_0 < t_1 < \cdots < t_n = 1$, and open neighbourhoods U_i of b, such that $U_i \times [t_i, t_{i+1}]$ is mapped into an evenly covered open set of X. Let U be the intersection of the U_i. Arguing by induction, similarly to the proof of Proposition 9.3.6, we can then construct a lift $\tilde{h}$ for the restriction of h on $U \times [0,1]$, with $\tilde{h}_0 = \tilde{f}$ on U.

Varying b in B, we get an open cover of B, and for each open set U of the cover, a lift of the restriction of h to $U \times [0,1]$. If U and V are two open sets in the cover, and $b \in U \cap V$, then the two partial lifts are in particular lifts of the path obtained by restricting h to $\{b\} \times [0,1]$, hence they must coincide on points of the form (b,t) for all t, thanks to the uniqueness of Proposition 9.3.6. It follows from Proposition 5.1.3 that all the partial lifts of h can be glued into a lift of h.

Uniqueness follows immediately, again from the uniqueness of Proposition 9.3.6 applied the paths obtained by restricting h to $\{b\} \times [0,1]$, for all $b \in B$. $\square$

Proposition 9.3.8. $\pi_1(S^1) \cong \mathbb{Z}$.

Proof. For all $n \in \mathbb{Z}$, let γ_n be a loop on 1 that winds around the circle n times, with positive winding being interpreted as travelling counter-clockwise, and negative winding clockwise. Explicitly, we define the path $\tilde{\gamma}_n$ on $\mathbb{R}$ as $\tilde{\gamma}_n(t) = nt$, and $\gamma_n = p \circ \tilde{\gamma}_n$.

This determines a map $\phi : \mathbb{Z} \to \pi_1(S_1)$, defined as $\phi(n) = [\gamma_n]$. To see that ϕ is a homomorphism, observe that $\tilde{\gamma}_n * (\tilde{\gamma}_m + n)$ and $\tilde{\gamma}_{n+m}$ are both paths connecting 0 with $n + m$ in $\mathbb{R}$, hence they are homotopic relative to $\{0, 1\}$ by Exercise 92. Since $p(x + n) = p(x)$, it follows that $p \circ (\tilde{\gamma}_m + n) = \gamma_m$, and therefore $\gamma_n * \gamma_m$ is homotopic to γ_{n+m}.

Now let ρ be any loop on 1 in S^1. By Proposition 9.3.6, there exists a lift $\tilde{\rho}$ of ρ such that $\tilde{\rho}(0) = 0$. Then $p(\tilde{\rho}(1)) = \rho(1) = 1$, hence $\tilde{\rho}(1)$ must be an integer n. Since $\tilde{\gamma}_n$ and $\tilde{\rho}$ are homotopic (again because they have the same endpoints and $\mathbb{R}$ is simply connected), it follows that $\phi(n) = [\gamma_n] = [\rho]$, showing that ϕ is surjective.

As for injectivity, suppose $\phi(n) = e$, where e is the class of the trivial loop in $\pi_1(S^1)$. That implies that there exists a homotopy h between the trivial loop and γ_n. By Proposition 9.3.7, h can be lifted to a homotopy $\tilde{h}$ between the trivial path in $\mathbb{R}$ on 0 and some lifting $\tilde{\rho}$ of γ_n. Since the homotopy is relative to $\{0, 1\}$, we must have $\tilde{\rho}(1) = 0$. But $\tilde{\rho}$ must also be homotopic to $\tilde{\gamma}_n$, hence $\tilde{\rho}(1) = n$. It follows that $n = 0$, as required. □

9.4 A special case of Seifert-van Kampen theorem

The Seifert-van Kampen theorem can be used to calculate the fundamental group of a union of two open sets U and V in terms of the fundamental groups of U, V and $U \cap V$. In this subsection, we will prove the special case of the theorem where both U and V are simply connected, which is already very useful, and will help us prove that some of the examples of spaces we have seen so far are simply connected.

Theorem 9.4.1. *Let X be a topological space, suppose that X is covered by two simply connected open sets U and V, and that $U \cap V$ is path connected. Then X is simply connected.*

Proof. Let $b \in U \cap V$ be any point, and take b to be the base point of X. Let γ be a loop in X at b. The two open sets $\gamma^{-1}(U)$ and $\gamma^{-1}(V)$ cover $[0, 1]$. By compactness, there exist a finite number of points $0 = t_0 < t_1 < \cdots < t_n = 1$ such that for all $0 \leq i < n$, $\gamma|[t_i, t_{i+1}]$ is valued either in U or V. Let γ_i be the reparameterisation of $\gamma|[t_i, t_{i+1}]$ to $[0, 1]$. Then it is easy to see that γ is homotopic to the concatenation of the γ_i. By removing some of the t_i, we can assume that if γ_i has values in U, then γ_{i+1} has values in V, and vice versa. In particular, $\gamma(t_i) \in U \cap V$ for all i.

For all $0 < i < n$, let α_i be a path in $U \cap V$ from b to $\gamma(t_i)$. Set α_0 and α_n to be the trivial loop at b. Define $\gamma'_i := \alpha_i * \gamma_i * \alpha_i^{-1}$. Then γ'_i is a loop at b for all i. Furthermore, since $\alpha_i^{-1} * \alpha_{i+1}$ is a loop in $U \cap V$, we can in particular regard it as a loop in U, hence it is homotopic to a trivial loop. In particular, the concatenation of all the γ'_i defines a loop γ' in X which is homotopic to γ.

Since U and V are simply connected, all the loops γ'_i are homotopic to the trivial loop at b, hence so is γ', proving that X is simply connected. □

Lemma 9.4.2. *Let X be any pointed space. Then CX is contractible.*

Proof. We define an explicit homotopy between the identity on CX and a constant. Namely, $h([x, s], t) = [(x, 1 - (1 - s)t)]$. Note that h is well defined, since if $s = 1$, then $1 - (1 - s)t = 1$ for all t. Furthermore, $h([x, s], 0) = [(x, 1)]$, hence h_0 is a constant map, and $h([x, s], 1) = [x, s]$, hence h_1 is the identity. □

Proposition 9.4.3. *Let X be a path connected space. Then the suspension SX is simply connected.*

Proof. Write $SX = U_+ \cup U_-$, where U_+ and $U-$ are open sets containing the two copies C_+X and C_X of CX as in Proposition 5.4.8. Now, C_+X and C_X are both contractible, and it is easy to see that the inclusion of C_+X (resp. C_X) in U_+ (resp. $U\_$) is a homotopy equivalence. Their intersection is homeomorphic to $X \times \mathbb{R}$, which is path connected. We can therefore apply Theorem 9.4.1 and conclude that SX is simply connected. □

Corollary 9.4.4. *The sphere S^n is simply connected for $n > 1$.*

Proof. By Exercise 53, $S^n = SS^{n-1}$. If $n > 1$, then S^{n-1} is path connected, and therefore S^n is simply connected by Proposition 9.4.3. □

9.5 Applications

Knowledge of the fundamental group of spheres can be directly applied to the problem of distinguishing some Euclidean spaces.

Proposition 9.5.1. *Let $n > 2$, then $\mathbb{R}^n$ is* not *homeomorphic to $\mathbb{R}^2$.*

Proof. Let $\phi : \mathbb{R}^n \to \mathbb{R}^2$ a homeomorphism. Then ϕ induces a homeomorphism $\mathbb{R}^n \setminus \{0\} \to \mathbb{R}^2 \setminus \{0\}$. It is easy to see that $\mathbb{R}^k \setminus \{0\}$ is homotopy equivalent to S^{k-1} for all $k > 0$, hence in particular $\mathbb{R}^n \setminus \{0\}$ is simply-connected, whereas $\mathbb{R}^2$ has $\mathbb{Z}$ as fundamental group, so they cannot be homeomorphic. □

The idea of Proposition 9.5.1 could be extended to higher dimensional spaces, if only we had access to an invariant which is non-trivial for the m-sphere, and trivial for higher-dimensional spheres. Such invariants do indeed exist. The simplest example is the m-th homology group, which is quite easy to calculate (at least for spheres), and is sufficient for this purpose. Another is given by the higher homotopy group π_m, which is the natural generalisation of the fundamental group. Homotopy groups are much harder to calculate, even for spheres, but they sometimes provide more information on the underlying spaces.

Another simple application of homotopy groups is the famous *Brouwer's fixed point theorem*:

Theorem 9.5.2. *Let $f : D^2 \to D^2$ be a continuous map. Then f has a fixed point.*

Proof. By contradiction, assume that f has no fixed points. For $x \in D^2$, trace a half-line starting from $f(x)$ and going through x, and let $g(x)$ be the intersection of this half-line with the circle S^1. By writing out an explicit formula for $g(x)$, it is easy to see that g is a continuous function $D^2 \to S^1$, and if $x \in S^1$, then $g(x) = x$.

In other words, if $i : S^1 \to D^2$ is the inclusion function, we have that $g \circ i = \mathsf{id}$. By Exercise 94, it follows that $g_* \circ i_* = \mathsf{id}$, where g_* and i_* are the corresponding maps on the fundamental groups.

However, D^2 is contractible, which implies that i_* (and hence $g_* \circ i_*$) is the zero map. It follows that the identity $\mathbb{Z} \to \mathbb{Z}$ is the zero map, which is a contradiction. $\square$

Using covering spaces, and the fact that S^n is simply connected for all $n > 1$, we can now calculate the fundamental group of real projective spaces.

Lemma 9.5.3. *Let $p : E \to X$ be a covering space, with X path connected, and let $b \in X$ be a base point. Then $\pi_1(X, b)$ acts on the fibre F of p over b. The action is transitive if and only if E is path connected. The action is free if and only if $\pi_1(E, u)$ is trivial for all $u \in F$.*

Proof. If γ is a loop at b, and $u \in F$, let $\tilde{\gamma}_u$ be the unique lift of γ to E with $\tilde{\gamma}_u(0) = u$. It follows from Proposition 9.3.7 that $\tilde{\gamma}_u(1)$ only depends on the homotopy class of γ in $\pi_1(X, b)$, hence the assignment $(u, [\gamma]) \mapsto u \cdot [\gamma] := \tilde{\gamma}_u(1)$ defines a map $F \times \pi_1(X, b) \to F$, which we take to be the (right) action of $\pi_1(X, b)$ on F.

If γ is the trivial loop, then $\tilde{\gamma}_u$ can be chosen to be the trivial loop at u, hence in particular $\tilde{\gamma}_u(1) = 1$, which implies that $u \cdot e = u$, where e is the unit element of $\pi_1(X, b)$.

If γ and σ are loops at b, the path $\tilde{\gamma}_u * \tilde{\sigma}_{\tilde{\gamma}_u(1)}$ is a lift of $\gamma * \sigma$. Therefore, $u \cdot [\gamma * \sigma] = \tilde{\sigma}_{\tilde{\gamma}_u(1)}(1) = \tilde{\gamma}_u(1) \cdot [\sigma] = (u \cdot [\gamma]) \cdot [\sigma]$, which completes the proof that $\pi_1(X, b)$ acts on F.

Now assume that E is path connected, and let $u, v \in F$. Let $\tilde{\gamma}$ be a path connecting u and v in E, and $\gamma = p \circ \tilde{\gamma}$ its projection in X. Then $\tilde{\gamma}$ is clearly a lift of γ starting in u, hence $u \cdot [\gamma] = v$, which shows that the action is transitive. Conversely, assume that the action is transitive, and let $u, v \in E$. We want to find a path in E connecting u and v. Since X is path connected, we can assume that $u, v \in F$, thanks to Proposition 9.3.6. By transitivity, there exists $[\gamma] \in \pi_1(X, b)$ such that $u \cdot [\gamma] = v$, which means that $\tilde{\gamma}_u$ connects u and v, as desired.

As for the last statement, assume that $\pi_1(E, u)$ is trivial, and let $[\gamma] \in \pi_1(X, b)$ be such that $u \cdot [\gamma] = u$. Then $\tilde{\gamma}_u$ is a loop at u, which implies that it is homotopic to the trivial loop. Therefore γ is also homotopic to the trivial loop, hence $[\gamma] = e$. Conversely, suppose that the action is free, and let $\tilde{\gamma}$ be a loop at $u \in F$. If $\gamma = p \circ \tilde{\gamma}$, then $\tilde{\gamma}$ is a lift of γ, hence $u \cdot [\gamma] = \tilde{\gamma}(1) = u$, and hence γ is homotopic to the trivial loop by freeness. It then follows from Proposition 9.3.7 that $\tilde{\gamma}$ must also be homotopic to a trivial loop. $\square$

Corollary 9.5.4. *Let $p : E \to X$ be a covering space where E is simply connected. Then for any base point $b \in X$, the action of $\pi_1(X, b)$ on the fibre F over b is free and transitive, hence in particular there is a bijection $\pi_1(X, b) \cong F$.*

Corollary 9.5.4 generalises Proposition 9.3.8, and allows us to directly calculate the fundamental groups of real projective spaces.

Proposition 9.5.5. *For all $n > 1$, $\pi_1(\mathbb{R}P^n) = \mathbb{Z}/2\mathbb{Z}$.*

Proof. It is easy to see that the quotient projection $S^n \to \mathbb{R}P^n$ is a covering map. Since S^n is simply connected for $n > 1$, it follows from Corollary 9.5.4 that $\pi_1(\mathbb{R}P^n)$ has two elements, hence it must be isomorphic to the cyclic group $\mathbb{Z}/2\mathbb{Z}$. □

As an immediate consequence, $\mathbb{R}P^1$ is the only real projective space that is equivalent to a sphere.

As a final application of the results we have proved about the fundamental group, we give a proof of the famous *Borsuk-Ulam* theorem, which is often stated as: *at any point in time, there exist two antipodal points on Earth that have the same temperature and pressure.*

Theorem 9.5.6. *If $f : S^2 \to \mathbb{R}^2$ is a continuous function, there exists $x \in S^2$ such that $f(x) = f(-x)$.*

Proof. By contradiction, assume that for all $x \in S^2$ we have $f(x) \neq f(-x)$. That allows us to define a continuous function $g : S^2 \to S^1$ as

$$g(x) = \frac{f(x) - f(-x)}{\|f(x) - f(-x)\|},$$

which satisfies $g(-x) = -g(x)$. Let $\overline{g} : \mathbb{R}P^2 \to \mathbb{R}P^1$ be the map induced by g on projective spaces, and let p denote the projection $S^n \to \mathbb{R}P^n$.

If b is any base point on S^2, let σ be a path from b to $-b$ such that $p_*([\sigma])$ generates $\pi_1(\mathbb{R}P^2)$. For example, one can take

$$\sigma(t) = (\cos(\pi t), \sin(\pi t), 0).$$

Since $\overline{g}_* : \mathbb{Z}/2\mathbb{Z} \to \mathbb{Z}$ must be the zero map, it follows that $p_*(g_*([\sigma]))$ is the unit of $\pi_1(\mathbb{R}P^1)$, hence $p \circ g \circ \sigma$ is homotopic to a trivial loop. Therefore, its lift $g \circ \sigma$ along the covering map p must also be a loop.

However, $g(\sigma(1)) = g(-b) = -g(b) = g(\sigma(0))$, which shows that $g \circ \sigma$ is not a loop. Contradiction. □

References

[1] Munkres, J.R. 2000. *Topology.* Prentice Hall, Incorporated.

[2] Steen, L.A. and Seebach, J.A. 2003. *Counterexamples in Topology.* Dover Publications.