Clark Barwick

LECTURE NOTES FOR TOPOLOGY I

Clark Barwick
Department of Mathematics, Harvard University, One Oxford Street, Cambridge, MA 02138, USA.
E-mail : `clarkbar@gmail.com`

LECTURE NOTES FOR TOPOLOGY I

Clark Barwick

CONTENTS

ACT I

GENERAL TOPOLOGY

CHAPTER 1

THE CATEGORY OF TOPOLOGICAL SPACES

1.1. Topological spaces

Notation 1.1.1. — Given a set S, we denote by $\mathscr{P}(S)$ the *power set* of S, i.e., the set of all subsets of S. Any map $f : T \longrightarrow S$ induces a map

$$f^{-1} : \mathscr{P}(S) \longrightarrow \mathscr{P}(T),$$

where for any subset $A \subset S$,

$$f^{-1}A := \{t \in T \mid f(t) \in A\} \subset T.$$

One may verify that f^{-1} respects arbitrary intersections and arbitrary unions: for any subset $\mathscr{A} \subset \mathscr{P}(S)$,

$$f^{-1}\left(\bigcap_{A \in \mathscr{A}} A\right) = \bigcap_{A \in \mathscr{A}} f^{-1}A \quad \text{and} \quad f^{-1}\left(\bigcup_{A \in \mathscr{A}} A\right) = \bigcup_{A \in \mathscr{A}} f^{-1}A;$$

moreover, f^{-1} respects complements: for any set $A \in \mathscr{P}(S)$, one may verify that

$$f^{-1}(\complement_S A) = \complement_T(f^{-1}A).$$

Definition 1.1.2. — A *topological space* (or more simply a *space*) is a pair $(X, \mathscr{O}p(X))$ comprised of a set X and a subset $\mathscr{O}p(X) \subset \mathscr{P}(X)$ satisfying the following properties.

(1.1.2.1) For any finite subset $\mathscr{U} \subset \mathscr{O}p(X)$, the intersection

$$\bigcap_{U \in \mathscr{U}} U \in \mathscr{O}p(X).$$

(1.1.2.2) For any subset $\mathscr{U}' \subset \mathscr{O}p(X)$, the union

$$\bigcup_{U' \in \mathscr{U}'} U' \in \mathscr{O}p(X).$$

In this case, the set $\mathscr{O}p(X)$ is said to be a *topology on* X. The elements $U \in \mathscr{O}p(X)$ will be called *open* in X. A subset $V \in \mathscr{P}(X)$ will be called *closed* if its complement $\complement V := X - V$ is an open set.

1.1.3. — In the definition above, note that when $\mathscr{U} = \varnothing$, condition (1.1.2.1) translates to the condition that $X \in \mathscr{O}p(X)$. Similarly, when $\mathscr{U}' = \varnothing$, condition (1.1.2.2) translates to the condition that $\varnothing \in \mathscr{O}p(X)$.

1.1.4. — We have described a topology on a set X by prescribing the set $\mathscr{O}p(X)$ of open sets; we may also describe a topology on a set X by prescribing the set $\mathscr{C}l(X)$ of *closed* sets.

A topology on a set X is, equivalently, a subset $\mathscr{C}l(X) \subset \mathscr{P}(X)$ satisfying the following axioms.

(1.1.4.1) For any subset $\mathscr{V} \subset \mathscr{C}l(X)$, the intersection

$$\bigcap_{V \in \mathscr{V}} V \in \mathscr{C}l(X).$$

(1.1.4.2) For any finite subset $\mathscr{V}' \subset \mathscr{C}l(X)$, the union

$$\bigcup_{V' \in \mathscr{V}'} V' \in \mathscr{C}l(X).$$

As a result, we sometimes refer to the pair $(X, \mathscr{C}l(X))$ as a topological space. In practice, it will always be clear whether we are describing the set of open subsets of X or the set of closed subsets of X.

Definition 1.1.5. — Suppose X a set, and suppose $\mathscr{O}p(X)$ and $\mathscr{O}p'(X)$ two topologies on X. If $\mathscr{O}p'(X) \subset \mathscr{O}p(X)$, then we say that $\mathscr{O}p(X)$ is *finer* than $\mathscr{O}p'(X)$ and that $\mathscr{O}p'(X)$ is *coarser* than $\mathscr{O}p(X)$.

Definition 1.1.6. — Suppose $(X, \mathscr{O}p(X))$ a topological space. Then there are many different kinds of subsets of X.

(1.1.6.1) For any point $x \in X$, an *open neighborhood* of x is an open set $U \in \mathscr{O}p(X)$ containing x. A *neighborhood* of x is a subset $V \subset X$ containing an open neighborhood of x.

(1.1.6.2) For any subset $A \subset X$, the *closure* of A is the set

$$\overline{A} := \{x \in X \mid \text{for any neighborhood } V \text{ of } x,\ V \cap A \neq \varnothing\}.$$

In other words, the points of $\overline{A}$ are those points $x \in X$ with the property that every neighborhood of x meets A.

(1.1.6.3) Consider two subsets $A \subset B \subset X$. We say that A *is dense in* B if $\overline{A} \cap B = B$.

(1.1.6.4) Dual to the closure is the *interior*. For any subset $A \subset X$, the *interior* of A is the set

$$A^\circ := \{x \in A \mid \text{there exists a neighborhood } V \text{ of } x \text{ such that } V \subset A\}.$$

In other words, the points of A° are those points $x \in A$ with the property that there is a neighborhood of x contained in A.

Exercise 1. — Suppose $(X, \mathscr{O}p(X))$ a topological space. Suppose that $A \subset X$ is a subset. Prove that the *closure* $\overline{A} \subset X$ can be characterized in the following two ways.

(1.1) The closure $\overline{A}$ is the set

$$\overline{A} := \{x \in X \mid \text{for any open neighborhood } V \text{ of } x,\ V \cap A \neq \varnothing\}.$$

(1.2) The closure $\overline{A}$ is the smallest closed subset containing A.

(1.3) The closure $\overline{A}$ is the intersection of all the closed subsets containing A:

$$\overline{A} = \bigcap_{V \in \mathscr{C}l(X), A \subset V} V.$$

Dually, prove that the *interior* $A^\circ \subset X$ can be characterized in the following two ways.

(1.4) The interior A° is the set

$$A^\circ := \{x \in A \mid \text{there exists an open neighborhood } V \text{ of } x \text{ such that } V \subset A\}.$$

(1.5) The interior A° is the largest open subset contained in A.

(1.6) The interior A° is the union of all the open subsets contained in A:

$$A^\circ = \bigcup_{U \in \mathscr{O}p(X), U \subset A} U.$$

Show that the closure and the interior operations are dual operations in the sense that, for any subset $A \subset X$,

$$\complement((\complement A)^\circ) = \overline{A} \quad \text{and} \quad \complement(\overline{\complement A}) = A^\circ.$$

Exercise* ***2.*** — Suppose X a set. A *Kuratowski closure operator on* X is a map

$$\varkappa : \mathscr{P}(X) \longrightarrow \mathscr{P}(X)$$

satisfying the following properties.

(2.1) For any subset $U \subset X$, one has $U \subset \varkappa(U)$.

(2.2) For any subset $U \subset X$, one has $\varkappa(\varkappa(U)) = \varkappa(U)$.

(2.3) For any finite subset $\mathscr{U} \subset \mathscr{P}(X)$, one has

$$\varkappa\left(\bigcup_{U\in\mathscr{U}} U\right) = \bigcup_{U\in\mathscr{U}} \varkappa(U).$$

Given a Kuratowski closure operator $\varkappa$, define subsets

$$\mathscr{C}l(X^{\varkappa}) := \{V \in \mathscr{P}(X) \mid \varkappa(V) = V\} \subset \mathscr{P}(X) \quad \text{and thus} \quad \mathscr{O}p(X^{\varkappa}) := \{U \in \mathscr{P}(X) \mid \complement U \in \mathscr{C}l(X^{\varkappa})\} \subset \mathscr{P}(X)$$

Show that the pair $X^{\varkappa} := (X, \mathscr{O}p(X^{\varkappa}))$ is a topological space.

Moreover, show that the map

$$\begin{aligned} \{\text{Kuratowski closure operators on } X\} &\longrightarrow \{\text{topologies on } X\} \\ \varkappa &\longmapsto \mathscr{O}p(X^{\varkappa}) \end{aligned}$$

is a bijection, where the inverse is the operation $A \longmapsto \overline{A}$ described in 1.1.6. That is, a topology on a set X can be specified simply by describing a Kuratowski closure operator.

1.2. First examples

Example 1.2.1. — For any set S, there are two topologies that exist automatically on S.

(1.2.1.1) The *indiscrete topology* S^{ι} is the pair $(S, \{\varnothing, S\})$.

(1.2.1.2) The *discrete topology* S^{δ} is the pair $(S, \mathscr{P}(S))$.

Example 1.2.2. — The *Sierpiński space* Σ_0 is the set $\{0,1\}$ equipped with the subset:

$$\mathscr{O}p(\Sigma_0) := \{\varnothing, \{1\}, \{0,1\}\} \subset \mathscr{P}(\{0,1\}).$$

One may verify directly that the Sierpiński space is in fact a topological space. The set $\{0,1\}$ can also be given a topology given by the subset

$$\mathscr{O}p(\Sigma_1) := \{\varnothing, \{0\}, \{0,1\}\} \subset \mathscr{P}(\{0,1\}).$$

This defines a space Σ_1, which is — as we will explain — *homeomorphic* to Σ_0. The topologies

$$\{0,1\}^{\iota},\ \Sigma_0,\ \Sigma_1, \text{ and } \{0,1\}^{\delta}$$

specify all the topologies on the set $\{0,1\}$.

Exercise 3. — Describe explicitly all the topologies on the set $\{0,1,2\}$.

Example 1.2.3. — Any set S can be given the *cofinite* topology, in which a subset $U \subset S$ is declared to be *open* if and only if either $S = \varnothing$, or else the complement $\complement_S U$ is finite. In this case, the closed subsets of S are the finite sets and S itself.

Similarly, any set S can be given the *cocountable* topology, in which a subset $U \subset S$ is declared to be *open* if and only if either $S = \varnothing$, or else the complement $\complement_S U$ is countable. In this case, the closed subsets of S are the countable sets and S itself.

Example 1.2.4. — Suppose $(X, \mathscr{O}p(X))$ a topological space, and suppose $A \subset X$ a subset of X. Then A may be given the *subspace topology*, defined by

$$\mathscr{O}p(A) := \{U \in \mathscr{P}(U) \mid \text{there exists } U' \in \mathscr{O}p(X) \text{ such that } U = U' \cap A\}.$$

A *subspace* of X is thus a subset $A \subset X$, equipped with the subspace topology.

Example 1.2.5. — Let us consider the vector space $\mathbf{R}^n$. For any point $x \in \mathbf{R}^n$ and any real number $\varepsilon > 0$, the *(open) ball of radius r centered at x* is the subset

$$B(x,\varepsilon) := \{y \in \mathbf{R}^n \mid \|y - x\| < \varepsilon\}.$$

Now the set $\mathbf{R}^n$ can be given a topology in the following manner: we will say that a subset $U \subset \mathbf{R}^n$ is *open* if for any point $x \in U$, there is a real number $\varepsilon > 0$ such that $B(x,\varepsilon) \subset U$. This is known as the *standard topology* on $\mathbf{R}^n$.

Example 1.2.6. — We can now combine the previous two examples to get a whole world of interesting topological spaces.

(1.2.6.1) The open ball of radius 1 at 0 is sometimes called the *n-cell*.

(1.2.6.2) The *closed n-disk* is the set
$$D^n := \{x \in \mathbf{R}^n \mid ||x|| \leq 1\}.$$
It is equal to the closure $\overline{B(0,1)}$ of the ball of radius 1 at 0.

(1.2.6.3) The $(n-1)$*-sphere* is the set
$$S^{n-1} := \{x \in \mathbf{R}^n \mid ||x|| \leq 1\}.$$
It is equal to the complement of the n-cell in this n-disk.

(1.2.6.4) The rational numbers $\mathbf{Q} \subset \mathbf{R}$ can be given the subspace topology.

(1.2.6.5) The irrational numbers $\mathbf{R} - \mathbf{Q} \subset \mathbf{R}$ can also be given the subspace topology.

Exercise 4. — Show that the subspace topology on $\mathbf{Q}$ is *not* the discrete topology on $\mathbf{Q}$. Show also that $\mathbf{Q}$ is *dense* in $\mathbf{R}$.

1.2.7. — One of the most difficult aspects of topology is that, despite the fact that the subset $\mathscr{O}p(X) \subset \mathscr{P}(X)$ is the information that is needed in order to specify a topology on a set X, it is frequently difficult to give an explicit description of the elements of $\mathscr{O}p(X)$.

This leads us to the notion of a *subbase* for a topology.

Proposition 1.2.8. — *Suppose X a set, and suppose $\mathscr{S} \subset \mathscr{P}(X)$ a collection of subsets of X. Then there exists a unique coarsest topology $\mathscr{O}p_{\mathscr{S}}(X)$ containing $\mathscr{S}$.*

Proof. — Define a topology in the following manner: let $\mathscr{B}$ be the set of all finite intersections of elements of $\mathscr{S}$ (including the empty intersection, which is all of X). Then let $\mathscr{O}p_{\mathscr{S}}(X)$ be the set of all unions of elements of $\mathscr{B}$.

This is indeed a topology. Unions of unions of elements of $\mathscr{B}$ are unions of elements of $\mathscr{B}$. Moreover, for any finite subset $\mathscr{U} \subset \mathscr{O}p_{\mathscr{S}}(X)$, one may write any element $U \in \mathscr{U}$ as a union of elements of $\mathscr{B}$:
$$U = \bigcup_{i \in I_U} B_i,$$
where $B_i \in \mathscr{B}$. Now
$$\bigcap_{U \in \mathscr{U}} U = \bigcap_{U \in \mathscr{U}} \bigcup_{i \in I_U} B_i = \bigcup_{U \in \mathscr{U},\ i_U \in I_U} \bigcap_{U \in \mathscr{U}} B_{i_U},$$
and since each set $\bigcap_{U \in \mathscr{U}} B_{i_U}$ is an element of $\mathscr{B}$ by construction, this intersection is a union of elements of $\mathscr{B}$.(1)

Any topology containing $\mathscr{S}$ must contain all unions of finite intersections of elements of $\mathscr{S}$, so any such topology must be finer than $\mathscr{O}p_{\mathscr{S}}(X)$. Thus $\mathscr{O}p_{\mathscr{S}}(X)$ is the unique coarsest topology containing $\mathscr{S}$. □

Definition 1.2.9. — In the situation of the previous proposition, the topology $\mathscr{O}p_{\mathscr{S}}(X)$ is said to be *generated by* $\mathscr{S}$, and the set $\mathscr{S}$ is said to be a *subbase* for the topological space $(X, \mathscr{O}p_{\mathscr{S}}(X))$.

Exercise 5. — Show that the standard topology on $\mathbf{R}$ [1.2.5] is generated by the set
$$\{(a,\infty) \mid a \in \mathbf{R}\} \cup \{(-\infty, b) \mid b \in \mathbf{R}\} \subset \mathscr{P}(\mathbf{R}).$$

Example 1.2.10. — More generally, on any totally ordered set X, one may define the *order topology*, which is generated by the set
$$\{(a,\infty) \mid a \in X\} \cup \{(-\infty, b) \mid b \in X\} \subset \mathscr{P}(X),$$
where (a,∞) and $(-\infty, b)$ are shorthand:
$$(a,\infty) := \{x \in X \mid x > a\} \quad \text{and} \quad (-\infty, b) := \{x \in X \mid x < b\}.$$

(1)If this formula seems confusing, contemplate the case when $\mathscr{U}$ has cardinality 2.

Example 1.2.11. — Suppose X and Y two spaces. Then their product

$$X \times Y := \{(x,y) \mid x \in X,\ y \in Y\}$$

has a topology, called the *product topology* generated by the set

$$\{U \times Y \mid U \in \mathscr{O}p(X)\} \cup \{X \times V \mid V \in \mathscr{O}p(Y)\}.$$

Exercise 6. — Show that the standard topology on $\mathbf{R}^n$ can be described inductively in the following way.
(6.1) If $n = 0$, then $\mathbf{R}^0 = \star$, and the standard topology is just the discrete topology.
(6.2) If $n = 1$, then $\mathbf{R}^1 = \mathbf{R}$, and the standard topology is just the order topology, as you showed above.
(6.3) If $n > 1$, then $\mathbf{R}^n = \mathbf{R}^{n-1} \times \mathbf{R}$, and the standard topology on $\mathbf{R}^n$ is just the product topology of the standard topology on $\mathbf{R}^{n-1}$ with the standard topology on $\mathbf{R}$.

1.3. Continuous maps and homeomorphisms

Definition 1.3.1. — Suppose X and Y topological spaces. Then a map $f : X \longrightarrow Y$ is *continuous* if the inverse image of any open set is open, that is, if for any $U \in \mathscr{O}p(Y)$, the set $f^{-1}U \in \mathscr{O}p(X)$.

Example 1.3.2. — Of course the identity map of any space is a homeomorphism. More generally, suppose X a space, and suppose $A \subset X$ is a subset, equipped with the subspace topology. Then the *inclusion map* $i : A \longrightarrow X$ (given by the rule $i(a) = a$) is continuous.

Exercise 7. — In fact, the subspace topology is rigged precisely to make this happen. Show that if X is a space, and $A \subset X$ is a subset, then the subspace topology on A is the coarsest topology on A such that the inclusion map $i : A \longrightarrow X$ is continuous.

Example 1.3.3. — The composition of two continuous functions is continuous.

Example 1.3.4. — Suppose X is a set with two topologies $\mathscr{O}p(X)$ and $\mathscr{O}p'(X)$. Then $\mathscr{O}p(X)$ is finer than $\mathscr{O}p'(X)$ if and only if the identity map on X is a continuous map

$$(X, \mathscr{O}p(X)) \longrightarrow (X, \mathscr{O}p'(X)).$$

Example 1.3.5. — Suppose X and Y two spaces, and consider the product $X \times Y$ with the product topology. Then we have two *projection maps*

$$\mathrm{pr}_1 : X \times Y \longrightarrow X \quad \text{and} \quad \mathrm{pr}_2 : X \times Y \longrightarrow Y,$$

given by the rules $\mathrm{pr}_1(x,y) = x$ and $\mathrm{pr}_2(x,y) = y$. Both of these maps are continuous.

Exercise 8. — Once again, the way we defined the product topology guaranteed that this property was as sharp as possible. If X and Y are spaces, then the product topology is the coarsest topology on $X \times Y$ such that the two projection maps

$$\mathrm{pr}_1 : X \times Y \longrightarrow X \quad \text{and} \quad \mathrm{pr}_2 : X \times Y \longrightarrow Y,$$

are continuous.

Definition 1.3.6. — A map $f : X \longrightarrow Y$ is *open* if the image of any open set is open, that is, if for any $U \in \mathscr{O}p(X)$, $fU \in \mathscr{O}p(Y)$.

Example 1.3.7. — Open maps need not be continuous. Suppose X a topological space, and suppose S any set. Let us endow S with the discrete topology S^{δ}. Then one sees that any map $p : X \longrightarrow S^{\delta}$ is open; however, in order for p to be continuous, for any $s \in S$ the *fiber* over s — i.e., the inverse image $f^{-1}\{s\}$ — would have to be an open set. Hence the *floor function* $\lfloor \cdot \rfloor : \mathbf{R} \longrightarrow \mathbf{Z}$ (which just rounds any real number down to the nearest integer) is open, but not continuous.

Example 1.3.8. — If X and Y are topological spaces, then the projection maps pr_1 and pr_2 are both continuous (as we have seen) and open.

Example 1.3.9. — Consider the function $f : \mathbf{R} \longrightarrow \mathbf{R}$ given by the rule $f(x) = x^2$. This is continuous, but not open.

Exercise 9. — Show that if X is a space, and $A \subset X$ is a subset with the subspace topology, then the inclusion map $i : A \longrightarrow X$ is open if and only if A is an open subset of X.

Definition 1.3.10. — A continuous map $f : X \longrightarrow Y$ is a *homeomorphism* if it is both a bijection and an open map. Two spaces X and Y are said to be *homeomorphic* if there is a homeomorphism $X \longrightarrow Y$.

1.3.11. — A homeomorphism between two topological spaces is a guarantee that those two spaces are *indistinguishable from the point of view of topology*. In group theory or set theory, an isomorphism or bijection between two objects (groups in the first case, sets in the second) guaranteed that the objects involved were the same, up to some "relabeling," so anything you knew about the first group or set was reflected in the second group of set.

As we will see, the same phenomenon happens here. In the case of homeomorphisms, not only do you have the same points up to some relabeling, but you have the same open sets up to that relabeling!

1.3.12. — Here's another way to think of homeomorphism. If $f : X \longrightarrow Y$ is a continuous map, then the inverse image can be viewed as a map

$$f^{-1} : \mathscr{O}p(Y) \longrightarrow \mathscr{O}p(X).$$

The statement that f is a homeomorphism is the statement that both f and f^{-1} are bijections.

Alternately, if $f : X \longrightarrow Y$ is a bijection, then it has an inverse $f^{-1} : Y \longrightarrow X$. The statement that f is a homeomorphism is the statement that both f and f^{-1} are continuous.

(We're being totally sloppy about notation here, but only because everyone else is too, and you might as well get used to it now. Mathematicians use the symbol f^{-1} in two different ways, often in close proximity to one another. (1) In the first description, we were assuming that f is continuous, and we were talking about the function $\mathscr{O}p(Y) \longrightarrow \mathscr{O}p(X)$ that takes an open set U to its inverse image $f^{-1}U$, which is again an open set. (2) In the second description, we were assuming that f is a bijection, and we were talking about the function $Y \longrightarrow X$ which is the inverse to f.)

Example 1.3.13. — Of course the identity map is a homeomorphism.

Example 1.3.14. — The two Sierpiński spaces Σ_0 and Σ_1 are homeomorphic.

Example 1.3.15. — Note that, in order for $f : X \longrightarrow Y$ to be a homeomorphism, we do not only require that f be a bijective continuous map. That is not enough: let us think of the set

$$U(1) := \{z \in \mathbf{C} \mid |z| = 1\}$$

with the subspace topology from $\mathbf{C} \cong \mathbf{R}^2$. Now contemplate the map

$$\rho : [0, 2\pi) \longrightarrow U(1)$$

given by the rule $\rho(x) = e^{ix}$. This wraps the half-open interval around the circle. It is a bijective continuous map, but it had better not be a homeomorphism!

Example 1.3.16. — Here's a very important example to contemplate. Consider the tangent function as a map

$$\tan : \left(-\tfrac{\pi}{2}, \tfrac{\pi}{2}\right) \longrightarrow \mathbf{R}.$$

It's a homeomorphism! This suggests something that may at first seem surprising: topology doesn't "see" the difference between an infinitely long line and an open interval of finite length.

This might seem a little odd, since we've been thinking of topological spaces as sets with a concept of "nearness" of points. But this leads to an important idea about what topology actually is: one might say that topology is what is left over from geometry when you *throw out* the notion of distance.

Exercise 10. — Continuing with that idea, find a homeomorphism between $\mathbf{R}$ and the ray $(0, \infty)$.

Exercise 11. — Suppose n is a nonnegative integer. Show that any two n-dimensional balls $B^n(x, \delta)$ and $B^n(y, \varepsilon)$ in $\mathbf{R}^n$ are homeomorphic, and so are their closures.

Exercise * ***12.*** — Suppose n is a nonnegative integer. Show that the following three spaces are all homeomorphic:

(12.1) the n-dimensional ball $B^n(x, \varepsilon)$ for any $x \in \mathbf{R}^n$ and any $\varepsilon > 0$,
(12.2) the euclidean space $\mathbf{R}^n$, and

(12.3) the product space $\prod_{i=1}^{m} B^{k_i}(x_i, \varepsilon_i)$ for any decomposition $n = \sum_{i=1}^{m} k_i$, and points $x_i \in \mathbf{R}^{k_i}$, and any real numbers $\varepsilon_i > 0$.

Example 1.3.17. — Here's a challenging example. Consider the spaces $\mathbf{R}^m$ and $\mathbf{R}^n$ with the standard topology. It is a fact that they are homeomorphic if and only if $m = n$. Certainly one direction here is obvious! But how on earth could you show that there is *no* homeomorphism between two spaces? What kind of contradiction could you get from the existence of a homeomorphism? These kinds of questions will lead us naturally to

Example 1.3.18. — Here's another example that merits some contemplation. Consider the function $f(x) = 1/x$. This can regarded as a map

$$f : \mathbf{R} - \{0\} \longrightarrow \mathbf{R} - \{0\}.$$

Since f is its own inverse, f here is a homeomorphism. (In general, a homeomorphism from a space to itself is called an *automorphism*.) Of course we have removed the point $0 \in \mathbf{R}$, because in elementary school we were told that $1/0$ is "undefined." But let's try to define it anyhow.

We note that, as x approaches 0 from the right, $1/x$ increases without bound; as x approaches 0 from the left, $1/x$ decreases without bound. If we wanted to add a point that would play the role of $1/0$, then this leads us to the following idea: consider the set $\mathbf{R} \sqcup \{\infty\}$, where ∞ here is just a formal symbol that we've tacked on. Let's topologize this set in the following way: we will contemplate the set

$$\{(a,b) \mid a,b \in \mathbf{R}\} \cup \{(-\infty, c) \cup \{\infty\} \cup (d, \infty) \mid c, d \in \mathbf{R}\} \subset \mathscr{P}(\mathbf{R} \cup \{\infty\}),$$

and we will contemplate the topology generated by this set. That is, we will look at all unions of all finite intersections of open intervals in $\mathbf{R}$ along with sets

$$(-\infty, c) \cup \{\infty\} \cup (d, \infty)$$

that "wrap around" our added point ∞.

Now we think of our map $f : \mathbf{R} - \{0\} \longrightarrow \mathbf{R} - \{0\}$, and we note that we can extend it to a map $F : \mathbf{R} \cup \{\infty\} \longrightarrow \mathbf{R} \cup \{\infty\}$, where we set $F(x) := 1/x$ for any $x \in \mathbf{R} - \{0\}$, we set $F(0) := \infty$, and we set $F(\infty) := 0$. With the topology we've given $\mathbf{R} \cup \{\infty\}$, this is continuous!

The space $\mathbf{R} \cup \{\infty\}$ we've constructed here is called a *compactification* of $\mathbf{R}$. (We'll explain more about that later.) The subspace $\mathbf{R} \subset \mathbf{R} \cup \{\infty\}$ is open, and ∞ is contained in the closure of any ray $(-\infty, c)$ or (d, ∞).

Exercise 13. — A little bit of thought should convince you that the space $\mathbf{R} \cup \{\infty\}$ we've described in the previous example is in fact a circle. Let's find a homeomorphism that makes this explicit. Translate our circle S^1 so that we are looking at the unit circle around the point $(0,1)$ in $\mathbf{R}^2$:

$$S^1(1) := \{\theta \in \mathbf{R}^2 \mid \|\theta - (0,1)\| = 1\}.$$

Now let us define a map $g : S^1(1) - \{(0,2)\} \longrightarrow \mathbf{R}$ in the following way. For any point $\theta \in S^1(1) - \{(0,2)\}$, let ℓ_θ be the unique line passing through $(0,2)$ and θ in $\mathbf{R}^2$. Now define $g(\theta)$ to be the x-coordinate of the point of intersection between ℓ_θ and the x-axis $\{(x,y) \in \mathbf{R}^2 \mid y = 0\}$.

(13.1) Write an explicit formula for g.
(13.2) Show that g is a homeomorphism.
(13.3) Show that g can be extended to a homeomorphism

$$G : S^1(1) \longrightarrow \mathbf{R} \cup \{\infty\}$$

by setting $G(\theta) := g(\theta)$ for any $\theta \in S^1(1) - \{(0,2)\}$ and $G((0,2)) := \infty$.

(13.4) In the previous example, we constructed a homeomorphism $F : \mathbf{R} \cup \{\infty\} \longrightarrow \mathbf{R} \cup \{\infty\}$ (which is not the identity). Using the homeomorphism G, we can rewrite this as a homeomorphism

$$G^{-1} \circ F \circ G : S^1(1) \longrightarrow S^1(1).$$

What homeomorphism is this? How might you draw its graph?

Exercise 14. — The space $\mathbf{R}\cup\{\infty\}$ is not the only compactification of $\mathbf{R}$; we were led to consider it by our example. If we thought of another example, we might be led to a different compactification. Think, for example, of the function $h(x)=x^3$, which we regard as a map

$$h:\mathbf{R}\longrightarrow\mathbf{R}.$$

(14.1) Show that h is a homeomorphism. (Be careful if you decide to study h^{-1}!)

(14.2) Show that $h(x)$ increases without bound as x does, and likewise $h(x)$ decreases without bound as x does.

(14.3) Describe a space $\mathbf{R}\cup\{-\infty,\infty\}$ (where $-\infty$ and ∞ are each just formal symbols) with the following properties. (1) $\mathbf{R}$ is an open subspace of $\mathbf{R}\cup\{-\infty,\infty\}$. (2) The closure of a ray $(-\infty,a)\subset\mathbf{R}$ in $\mathbf{R}\cup\{-\infty,\infty\}$ is the set

$$[-\infty,a]:=\{-\infty\}\cup(-\infty,a],$$

and the closure of a ray $(b,\infty)\subset\mathbf{R}$ in $\mathbf{R}\cup\{-\infty,\infty\}$ is the set

$$[b,\infty]:=[b,\infty)\cup\{\infty\}.$$

(14.4) Describe an extension of h to a homeomorphism

$$H:\mathbf{R}\cup\{-\infty,\infty\}\longrightarrow\mathbf{R}\cup\{-\infty,\infty\}.$$

(14.5) Show that the space $\mathbf{R}\cup\{-\infty,\infty\}$ is homeomorphic to the interval $[-1,1]$.

(14.6) Show that the map

$$p:\mathbf{R}\cup\{-\infty,\infty\}\longrightarrow\mathbf{R}\cup\{\infty\}$$

(with the topologies defined here and above) is continuous. Show, moreover, that for any topological space X, and any continuous map $f:\mathbf{R}\cup\{-\infty,\infty\}\longrightarrow X$ such that $f(-\infty)=f(\infty)$, there exists a unique map $f':\mathbf{R}\cup\{\infty\}\longrightarrow X$ such that $f=f'\circ p$.

CHAPTER 2

CONSTRUCTING TOPOLOGIES AND CONTINUOUS MAPS

2.1. Quotients

Definition 2.1.1. — Suppose X a space, and suppose $A \subset X$ a subspace. Then a *quotient of X by A* is a pointed space $(X/A, \alpha)$ and a continuous map $q : X \longrightarrow X/A$ such that $q(A) = \{\alpha\}$, and for any space Y and any continuous map $f : X \longrightarrow Y$ such that $f(A) = \{a\}$ for some element $a \in Y$, there exists a unique continuous map $f' : X/A \longrightarrow Y$ such that $f'(\alpha) = a$ and $f = f' \circ q$.

2.1.2. — Suppose X a space, and suppose $A \subset X$ a subspace. The definition above is nonconstructive, in the sense that it does not make it apparent that such a space X/A exists. In fact, it always does. Here's one way to construct it.

Let $X/A = (X - A) \sqcup \{\alpha\}$; consider the map $q : X \longrightarrow X/A$, given by the formula

$$q(x) := \begin{cases} x & \text{if } x \in X - A; \\ \alpha & \text{otherwise.} \end{cases}$$

We topologize X/A in the following manner. We say that a subset $U \subset X/A$ is *open* if and only if the inverse image $q^{-1}U$ is open in X. This forces q to be continuous. (In fact, this is the finest possible topology on X/A such that q is continuous.)

I claim that the space X/A and the continuous map $q : X \longrightarrow X/A$ satisfy the conditions of the definition of a quotient. To see this, suppose Y a topological space, and suppose $f : X \longrightarrow Y$ a continuous map such that $f(A) = \{a\}$ for some element $a \in Y$. Let us define $f' : X/A \longrightarrow Y$ in the following way

$$f'(x) := \begin{cases} x & \text{if } x \in X - A; \\ a & \text{otherwise.} \end{cases}$$

One sees that this is the only possible map such that $f = f' \circ q$. Moreover, f' is continuous: for any open set $V \subset Y$, $f'^{-1}V$ is open in X/A if and only if $f^{-1}V = q^{-1}f'^{-1}V$ is open in X, but this is automatic since f is continuous.

Exercise 15. — Show that the above description of X/A is unique up to unique homeomorphism. That is, show that if X is a space, $A \subset X$ is a subspace, and both $q : X \longrightarrow W$ and $q' : X \longrightarrow W'$ are each quotients of X by A, then there is a unique homeomorphism $h : W \longrightarrow W'$ such that $h \circ q = q'$.

Exercise* 16. — For any integer $n \geq 0$, show that the quotient D^n/S^{n-1} is homeomorphic to S^n.

2.1.3. — Contemplate the quotient $X/\varnothing$. The answer may surprise you!

2.2. Coverings and locality

Exercise 17. — Subbases for topologies permit the easy identification of continuous maps. Suppose X and Y topological spaces, and suppose $\mathscr{S}$ a subbase for the topology on Y. Then show that a map $f : X \longrightarrow Y$ is continuous if and only if the inverse image of any element $U \in \mathscr{S}$ under f is open in X.

2.2.1. — The previous result is helpful because it permits one to check only a small number of conditions on subsets of Y in order to guarantee the continuity of a function.

Definition 2.2.2. — Suppose X and Y two topological spaces, suppose $f : X \longrightarrow Y$ a map, and suppose $x \in X$. Then one says that f is *continuous at* x if for any neighborhood V of $f(x) \in Y$, there is a neighborhood U of $x \in X$ such that $f(U) \subset V$.

2.2.3. — Note that I have written "neighborhood" in the above definition. I could easily as well have written "open neighborhood": f is continuous at x if and only if for any open neighborhood V of $f(x) \in Y$, there is an open neighborhood U of $x \in X$ such that $f(U) \subset V$.

Lemma 2.2.4. — *Suppose X a topological space, and suppose $U \subset X$ a subset thereof. Then U is open if and only if, for any point $x \in U$, there exists an open neighborhood W of x with $W \subset U$.*

Proof. — If U is open, then of course U itself is an open neighborhood of any point $x \in U$.

Conversely, one may for the interior U° of U, which is the union of all the open sets $W \subset U$. One sees that $U^\circ = U$: indeed, $U^\circ \subset U$, and for any point $x \in U$, there exists an open neighborhood W of x with $W \subset U$, whence $x \in U^\circ$. □

Proposition 2.2.5. — *Suppose X and Y two topological spaces, and suppose $f : X \longrightarrow Y$ a map. Then f is continuous if and only if it is continuous for every point $x \in X$.*

Proof. — Suppose f continuous, and suppose $x \in X$ a point. Then for any neighborhood V of $f(x) \in Y$, one may contemplate the interior V°, which is open, and its inverse image $f^{-1}(V^\circ)$, which is open as well. Now the image $f(f^{-1}(V^\circ)) \subset V^\circ \subset V$. Hence f is continuous at every point $x \in X$.

Suppose, conversely, that f is continuous at every point $x \in X$. Suppose $V \subset Y$ an open set. The claim is that $f^{-1}(V)$ is open. For any point $x \in f^{-1}(V)$, there is an open neighborhood W of $x \in f^{-1}(V)$ with $f(W) \subset V$, or equivalently, with $W \subset f^{-1}(V)$. Now we apply the lemma above to see that $f^{-1}(V)$ is open. □

Definition 2.2.6. — Suppose X a space. A collection $\mathscr{V} \subset \mathscr{P}(X)$ of subspaces of X is said to *cover* X — or to be *a covering of* X — if $X = \bigcup_{V \in \mathscr{V}} V$. One says that $\mathscr{V}$ is an *open covering of* X if every $V \in \mathscr{V}$ is an open set; likewise, one says that $\mathscr{V}$ is an *closed covering of* X if every $V \in \mathscr{V}$ is an closed set.

For a subspace $A \subset X$, we say that a collection $\mathscr{V} \subset \mathscr{P}(X)$ of subspaces of X is said to *cover* A — or to be *a covering of* A — if $A \subset \bigcup_{V \in \mathscr{V}} V$, or, equivalently, of $A = \bigcup_{V \in \mathscr{V}} A \cap V$.

Exercise 18. — Suppose X a space, and suppose $\mathscr{V}$ an open covering of X. Now suppose, for any $V \in \mathscr{V}$, that $f_V : V \longrightarrow Y$ is a continuous function, and suppose that, for any elements $V, W \in \mathscr{V}$, one has

$$f_V|_{V \cap W} = f_W|_{V \cap W} : V \cap W \longrightarrow Y.$$

Then there exists a unique continuous map $f : X \longrightarrow Y$ such that for any $V \in \mathscr{V}$,

$$f|_V = f_V : V \longrightarrow Y.$$

Definition 2.2.7. — Suppose X a topological space, and suppose $x \in X$. Then a *fundamental system of neighborhoods* of $x \in X$ is a collection $\mathscr{V}_x$ of neighborhoods of x with the following property. For any neighborhood W of x, there is a neighborhood $V \in \mathscr{V}_x$ with $V \subset W$.

Definition 2.2.8. — Suppose X a topological space. Then a *base* for X is a collection $\mathscr{B} \subset \mathscr{O}p(X)$ such that for any point $x \in X$, the set

$$\mathscr{B}_x := \{U \in \mathscr{B} \mid x \in U\}$$

is a fundamental system of neighborhoods of $x \in X$.

2.2.9. — A base for a topological space is also a subbase for the topology $\mathscr{O}p(X)$.

Proposition 2.2.10. — *Suppose X is a topological space, and suppose $\mathscr{B}$ a base for the topology. The the following are equivalent for a subset $U \subset X$.*

(2.2.10.1) The set U is open in X.

(2.2.10.2) For any point $x \in U$, there exists an element $V \in \mathscr{B}_x$ such that $V \subset U$.

(2.2.10.3) The set U can be written as the union of elements of $\mathscr{B}$.

Proof. — Suppose U is open in X; then for any point $x \in U$, U is an open neighborhood of x. Since $\mathscr{B}_x$ is an open system of neighborhoods of $x \in X$, there is an element $V \in \mathscr{B}_x$ such that $V \subset U$. This shows that the first condition implies the second.

Suppose that, for any point $x \in U$, there exists an element $V \in \mathscr{B}_x$ such that $V \subset U$. Then (using the axiom of choice) one may select, for every point $x \in U$, one element $V_x \in \mathscr{B}_x$ such that $V_x \subset U$. The union $U' := \bigcup_{x \in U} V_x$ contains every point of U by definition, so that $U' \supset U$, and since $V_x \subset U$ for every $x \in U$, it follows that $U' \subset U$. Thus U is a union of elements of $\mathscr{B}$. This shows that the second condition implies the third.

Finally, suppose that U can be written as the union of elements of $\mathscr{B}$. By definition, every element of $\mathscr{B}$ is open, so by the axioms for a topological space, U is open as well. This completes the proof. $\square$

Proposition 2.2.11. — *Suppose X a topological space. Then a subset $\mathscr{B} \subset \mathscr{O}p(X)$ is a base for the topology if and only if the following conditions are satisfied.*

(2.2.11.1) The elements of $\mathscr{B}$ cover X.

(2.2.11.2) For any $U, V \in \mathscr{B}$ and any element $x \in U \cap V$, there is an element $W \in \mathscr{B}_x$ such that $W \subset U \cap V$.

Example 2.2.12. — For any totally ordered set X, the set

$$\{(a,b) \mid a,b \in X\}$$

is a base for the order topology on X. (Here, one writes

$$(a,b) := \{x \in X \mid a < x < b\}.)$$

Example 2.2.13. — Suppose $n \geq 0$ an integer. Then the set

$$\{B^n(x,\varepsilon) \mid x \in \mathbf{Q}^n,\ \epsilon \in \mathbf{Q}\}$$

is a base for the standard topology on euclidean space $\mathbf{R}^n$.

Exercise 19. — Suppose X a set with two topologies $\mathscr{O}p(X)$ and $\mathscr{O}p'(X)$; suppose $\mathscr{B}$ a base for $(X, \mathscr{O}p(X))$ and $\mathscr{B}'$ a base for $(X, \mathscr{O}p'(X))$. Show that $\mathscr{O}p(X)$ is finer than $\mathscr{O}p'(X)$ if and only if, for any $x \in X$ and any element $U \in \mathscr{B}'_x$, there exists an element $V \in \mathscr{B}_x$ such that $V \subset U$.

Example 2.2.14. — Here is a standard counterexample in topology; everyone must see it once in their lives. The *Sorgenfrey line* $\mathbf{R}_\ell$ is the set $\mathbf{R}$ with the *lower limit topology*. This is the topology generated by the half-open intervals $[a,b)$ for $a,b \in \mathbf{R}$. The set

$$\mathscr{B}_\ell := \{[a,b) \mid a,b,\in \mathbf{R}\}$$

is then a basis for $\mathbf{R}_\ell$. (One may also define the *upper limit topology* $\mathbf{R}_u$, which is homeomorphic to $\mathbf{R}_\ell$ via the map $x \longmapsto -x$.) One can check fairly easily that these basic open sets $[a,b)$ are both open and closed.

The previous exercise can be used to show that the Sorgenfrey line is finer than the standard (= order) topology on $\mathbf{R}$. Indeed, for any open interval (a,b), one may choose any point $a' \in (a,b)$, and then $[a',b) \subset (a,b)$.

A map $f : \mathbf{R}_\ell \longrightarrow \mathbf{R}$ is continuous if and only if, for any $x \in \mathbf{R}$,

$$f(x) = \lim_{\varepsilon \to 0^+} f(x+\varepsilon).$$

In particular, the floor function $\lfloor \cdot \rfloor : \mathbf{R}_\ell \longrightarrow \mathbf{R}$ is continuous.

CHAPTER 3

COMPACTNESS I

3.1. Compact spaces

Definition 3.1.1. — A topological space X is *compact* if any open covering has a finite subcovering. A compact topological space is sometimes called a *compactum*.

For any space X, write $\mathscr{C}p(X)$ for the set of all compact subsets of X.

Example 3.1.2. — Any finite space is compact, and the indiscrete topology on any set is compact.

Example 3.1.3. — The euclidean spaces $\mathbf{R}^n$ are not compact; indeed, the open balls $B^n(x,\varepsilon)$ form an open covering, but there is no finite subcovering. To see this, consider any finite covering of $\mathbf{R}^n$ by balls $B^n(x_i,\varepsilon_i)$ for $i=1,2,\ldots,m$, and suppose x a point of one of the balls. Then one may choose a real number r so that $B^n(x,r)\supset\bigcup_{i=1}^m B^n(x_i,\varepsilon_i)$.

Lemma 3.1.4. — *Suppose X and Y spaces, and suppose $f : X \longrightarrow Y$ a continuous function. Then the image $f(A)$ of any compact subspace $A\subset X$ is compact.*

Proof. — Suppose $\mathscr{U}$ an open covering of $f(A)$; then the set

$$\{f^{-1}U \mid U\in\mathscr{U}\}$$

is an open covering of A, which thus has a finite subcovering $\mathscr{V}$. Now the claim is that the subset

$$\{U\in\mathscr{U}\mid f^{-1}U\in\mathscr{V}\},$$

which is finite, is a covering of $f(A)$. Indeed, suppose $y=f(a)\in f(A)$ a point with $a\in A$. Then there exists some $U\in\mathscr{U}$ such that $f^{-1}U\in\mathscr{V}$ and $a\in f^{-1}U$, whence $f(a)\in U$. □

Lemma 3.1.5. — *A closed subspace of a compact space is compact.*

Proof. — Suppose X a compactum, and suppose $A\subset X$ closed. Suppose $\mathscr{U}$ an open covering of A. Since each element $U\in\mathscr{U}$ is open in A, one may choose a collection $\mathscr{U}'\subset\mathscr{O}p(X)$ such that

$$\mathscr{U}=\{U\in\mathscr{O}p(A)\mid U=U'\cap A \text{ for some} U'\in\mathscr{U}'\}.$$

Now the set $\mathscr{U}'\cup\{X-A\}$ is an open covering of A; hence there is a finite subcovering $\mathscr{V}'\subset\mathscr{U}'$. Now

$$\mathscr{V}=\{U\in\mathscr{O}p(A)\mid U=U'\cap A \text{ for some} U'\in\mathscr{V}'\}$$

is a finite subcovering of $\mathscr{U}$. □

Exercise 20. — Suppose we are given a finite collection $\{X_i\}$ of compact spaces, $i=1,2,\ldots,n$. Show that the product space $\prod_{i=1}^n X_i$ is compact. (This is the easy case of Tychonoff's Theorem, to which we shall return.)

Theorem 3.1.6. — *The following are equivalent for a subspace $A\subset\mathbf{R}^n$.*

(3.1.6.1) A is compact.
(3.1.6.2) [Bolzano–Weierstraß] Every sequence of points in A has a convergent subsequence.
(3.1.6.3) [Heine–Borel] A is closed and bounded.

Proof. — Let us show that the first of these conditions implies the second. Suppose A compact, and suppose $(x_i)_{i\geq 0}$ a sequence of points in A. If there were no convergent subsequence of $(x_i)_{i\geq 0}$, then the set $\{x_i\}_{i\geq 0}$ would be a closed subset of A, hence compact, and there would be a sequence ε_i of positive real numbers such that the balls $B^n(x_i,\varepsilon_i)$ would be disjoint. But then $\{B^n(x_i,\varepsilon_i)\}_{i\geq 0}$ would be an open cover of $\{x_i\}_{i\geq 0}$ with no finite subcover.

Let us show that the second property implies the third. Suppose A has the property that every sequence has a convergent subsequence. Then A must be bounded, since otherwise there exists a sequence $(x_i)_{i\geq 0}$ of points in A such that $||x_i|| \to \infty$. It must also be closed, since if $x \in \overline{A} - A$ is a point "at the boundary," one can construct a sequence of points of A converging to x.

Finally, let us show that the third condition implies the first. Suppose A is closed and bounded. Since A is bounded, it is contained in a box $[-a,a]^n$ for some $a > 0$. Since every closed subspace of a compact space is compact, it is enough to show that $[-a,a]^n$ is compact. By the previous exercise, it is enough to show that $[-a,a]$ is compact. Suppose $\mathscr{U}$ an open covering of $[-a,a]$; then let

$$c = \sup\{x \in \mathbf{R} \mid [-a,x] \text{ can be covered by finitely many elements of } \mathscr{U}\}.$$

Suppose, to generate a contradiction, that $c < a$; then let $U \in \mathscr{U}$ be an element of the open covering containing c. Then for some $\varepsilon > 0$, one has $B(c,\varepsilon) \subset U$. By the definition of c, one may cover $[-a,a-c/2]$ with finitely many elements of $\mathscr{U}$; adding in U to that set yields a finite open covering of $[-a,a+c/2]$, yielding the desired contradiction. □

Exercise 21. — Suppose X a compact space, and suppose $f : X \longrightarrow \mathbf{R}$ a continuous function. Then f attains both a maximum and a minimum value.

3.2. The Tychonoff Product Theorem

Definition 3.2.1. — A collection $\mathscr{F}$ of sets is said to be *of finite character* if the following condition holds: a set A is an element of $\mathscr{F}$ if and only if every finite subset of A is an element of $\mathscr{F}$.

***Lemma 3.2.2* (Tukey).** — *Suppose $\mathscr{F}$ a nonempty collection of sets. Then if $\mathscr{F}$ is of finite character, then every element $A \in \mathscr{F}$ is contained in a maximal set $M \in \mathscr{F}$, i.e., a set M such that no element of $\mathscr{F}$ properly contains M.*

Proof. — Give $\mathscr{F}$ a partial ordering by inclusion. The union of every chain of elements of $\mathscr{F}$ must also belong to $\mathscr{F}$, so by Zorn's Lemma, every element $A \in \mathscr{F}$ is contained in a maximal element. □

***Theorem 3.2.3* (Alexander Subbase).** — *Suppose X a topological space, and suppose $\mathscr{S}$ a subbase for $\mathscr{O}p(X)$ that is also an open cover of $\mathscr{O}p(X)$. If every subcover of $\mathscr{S}$ has a finite subcover, then X is compact.*

Proof. — The language here can get a tad confusing, so let's introduce some new terminology to make things a little simpler. Let's call a family $\mathscr{U} \subset \mathscr{P}(X)$ *exiguous* if it does *not* cover X, and let's call $\mathscr{U}$ *bourgeois* if any finite subset of $\mathscr{U}$ is exiguous.

Using our new terminology, we see that X is compact if and only if any bourgeois family of open sets is exiguous. We can also rewrite our assumption on the subbase $\mathscr{S}$: we are assuming that any bourgeois subfamily of $\mathscr{S}$ is exiguous.

Our aim is to show that any bourgeois family $\mathscr{A}$ of open sets is exiguous. Write

$$\mathscr{F} := \{\mathscr{U} \in \mathscr{P}(\mathscr{O}p(X)) \mid \mathscr{U} \text{ is bourgeois}\}.$$

This family is obviously of finite character. So suppose $\mathscr{A}$ a bourgeois family of open sets, and let $\mathscr{F}(\mathscr{A})$ be the collection of all bourgeois families of open sets of X that contain $\mathscr{A}$:

$$\mathscr{F}(\mathscr{A}) := \{\mathscr{U} \in \mathscr{F} \mid \mathscr{A} \subset \mathscr{F}\}.$$

Tukey's Lemma says that it contains a maximal element $\mathscr{M} \in \mathscr{F}(\mathscr{A})$. It is enough to show that $\mathscr{M}$ is exiguous, and this is now our goal.

Exercise 22. — Prove the following *Key Fact* about our maximal element $\mathscr{M}$:

> If $\{U_i\}_{i=1}^n$ is a finite collection of open sets of X such that the intersection $\bigcap_{i=1}^n U_i$ is contained in an element of $\mathscr{M}$, then one of the U_i is itself an element of $\mathscr{M}$.

(Hint: if U is an open set of X that is not contained in $\mathscr{M}$, show that there is a finite collection $\{V_j\}_{j=1}^m$ of elements of $\mathscr{M}$ such that $\{U\} \cup \{V_j\}_{j=1}^m$ is a covering of X.)

All right, now it's time to consider our subbase $\mathscr{S}$; again, we are assuming that any bourgeois subcollection of $\mathscr{S}$ is exiguous. Let's contemplate the intersection $\mathscr{M} \cap \mathscr{S}$; it is bourgeois, and thus (by our assumption on $\mathscr{S}$) exiguous.

Now I claim that

$$\bigcup_{U \in \mathscr{M}} U \subset \bigcup_{U \in \mathscr{M} \cap \mathscr{S}} U.$$

Suppose $x \in U \in \mathscr{M}$; then since $\mathscr{S}$ is a subbase, there is a finite intersection $W := \bigcap_{i=1}^n U_i$ of elements $U_i \in \mathscr{S}$ such that $x \in W$ and $W \subset U$. By the *Key Fact* above, one of the U_i is itself and element of $\mathscr{M}$. This means that $x \in U_i \in \mathscr{M} \cap \mathscr{S}$, just as we had hoped.

But remember the intersection $\mathscr{M} \cap \mathscr{S}$ is exiguous, so we conclude that $\mathscr{M}$ is also! □

Exercise* *23*. — Suppose X a totally ordered set. Let us say that X is *order-complete* if every subset of X has a supremum in X. Show that X is compact (with the order topology) if and only if it is order-complete.

***Theorem 3.2.4* (Tychonoff Product).** — *The product of compact spaces is compact.*

Proof. — Suppose $\{X_\alpha \mid \alpha \in A\}$ a collection of compact spaces, and give $\prod_{\alpha \in A} X_\alpha$ the product topology. This is the topology generated by the subbase

$$\mathscr{S} := \{\mathrm{pr}_\alpha^{-1}(U) \mid \alpha \in A,\ U \in \mathscr{O}p(X_\alpha)\}.$$

Let's use the Alexander Subbase Theorem. We need to show that any bourgeois subcollection $\mathscr{U}$ of $\mathscr{S}$ is exiguous. Suppose $\mathscr{U}$ a bourgeois subcollection of $\mathscr{S}$. For every $\alpha \in A$, set

$$\mathscr{U}_\alpha := \{U \in \mathscr{O}p(X_\alpha) \mid \mathrm{pr}_\alpha^{-1}(U) \in \mathscr{U}\}.$$

It is obvious that $\mathscr{U}_\alpha$ is bourgeois and hence exiguous (since X_α is compact). Hence for each $\alpha \in A$, we can choose a point $x_\alpha \in X_\alpha$ such that $x_\alpha \notin \bigcup_{U \in \mathscr{U}_\alpha} U$. But now the point

$$x := (x_\alpha)_{\alpha \in A}$$

is not a member of any element of $\mathscr{U}$. So $\mathscr{U}$ is exiguous. □

3.3. The compact-open topology

Definition 3.3.1. — Suppose X and Y two topological spaces. Denote by $\mathscr{C}(X,Y)$ the set of all continuous maps $X \longrightarrow Y$. For any subsets $A \subset X$ and $B \subset Y$, write

$$\mathscr{V}(A,B) := \{f \in \mathscr{C}(X,Y) \mid f(A) \subset B\}.$$

Now we topologize the set $\mathscr{C}(X,Y)$ in the following manner. Take the topology generated by the subbase

$$\{\mathscr{V}(K,U) \mid K \in \mathscr{C}p(X), \text{ and } U \in \mathscr{O}p(Y)\}.$$

This is called the *compact-open topology*.

Exercise 24. — Suppose X any compact space, and suppose (Y,d) a metric space. A sequence of continuous functions $f_n \in \mathscr{C}(X,Y)$ is said to *converge uniformly* to a continuous function $f \in \mathscr{C}(X,Y)$ if

$$\lim_{n\to\infty} \sup_{x \in X} d(f_n(x), f(x)) = 0;$$

in other words, for every $\varepsilon > 0$, there exists a natural number N such that for any $x \in X$ and any $n \geq N$,

$$d(f_n(x), f(x)) < \varepsilon.$$

Prove that a sequence f_n of continuous functions $X \longrightarrow Y$ converges uniformly if and only if it converges in $\mathscr{C}(X,Y)$ with the compact-open topology. Is the same true if X is not compact?

Exercise 25. — Suppose $X = S^\delta$ a discrete space. Show that the compact-open topology on $\mathscr{C}(X,Y)$ coincides with the product topology on $\prod_{s \in S} Y$.

CHAPTER 4

COMPACTNESS II

4.1. Forms of compactness

We have already seen that for subspaces of a euclidean space, compactness is equivalent to the Heine–Borel property (closed and bounded) of the Bolzano–Weierstraß property (the property every sequence has a convergent subsequence). This doesn't quite cut it for more general spaces. In this section, we'll discuss analogues and generalizations of these notions, and we'll also talk about the notion of a *net* or *Moore–Smith sequence*.

Definition 4.1.1. — Suppose X a topological space.

(4.1.1.1) One says that X is *limit point compact* if and only if every infinite subset A of X has a limit point (i.e., a point $x \in X$ that is contained in $\overline{A - \{x\}}$).

(4.1.1.2) One says that X is *countably compact* if and only if every countable open cover has a finite subcover.

(4.1.1.3) One says that X is *sequentially compact* if and only if every sequence $x : \mathbf{N} \longrightarrow X$ has a convergent subsequence.

4.1.2. — These three notions seem remarkably similar, and in $\mathbf{R}^n$ (and more generally, in any metrizable space) they are indistinct and, moreover, equivalent to compactness. However, they are different conditions in general, and they are very different from compactness. To see this, will contemplate some examples, including a key counterexample. First, let's see what implications we do have.

Exercise 26. — Show that a space X is countably compact if and only if every nested sequence $X \supset V_1 \supset V_2 \supset V_3 \supset \cdots$ of closed, nonempty subspaces of X has a nonempty intersection.

Proposition 4.1.3. — *Consider the following properties for a topological space X.*

(4.1.3.1) X is compact.

(4.1.3.2) X is sequentially compact.

(4.1.3.3) X is countably compact.

(4.1.3.4) X is limit point compact.

Then we have a diagram of implications

$$\begin{tikzcd}
(1) \arrow[dr, Rightarrow] & & (2) \arrow[dl, Rightarrow] \\
& (3) \arrow[d, Rightarrow] & \\
& (4) &
\end{tikzcd}$$

Proof. — It is obvious that a compact spaces is countably compact.

Let us now see that a sequentially compact space is countably compact. Suppose $\mathscr{U} = \{U_i\}_{i \geq 0}$ a countable open cover of X. Proceeding inductively, write $V_0 := U_0$, and let V_p be the first of the sequence of U_i's not covered by the $V_0 \cup V_1 \cup \cdots \cup V_{p-1}$. If this choice is impossible at any (finite) stage, then we have our required finite subcover; if this choice is always possible, then for every $n \in \mathbf{N}$ we may select an element $x_n \in V_n$ such that for any $i < n$,

$x_n \notin V_i$. Now sequential compactness guarantees that the sequence $(x_n)_{n\geq 0}$ has a convergent subsequence $(x_{n_i})_{i\geq 0}$; let y be the limit of this subsequence. Then $y \in V_n$ for some integer n, so for every $i \gg 0$ (in particular for i so that $n_i > n$), we have $x_{n_i} \in V_m$, a contradiction.

Let us now see that a countably compact space is limit point compact. Suppose X a countably compact space, and suppose $B \subset X$ a subset with no limit point. Let us show that every countable subset $A \subset B$ is finite, whence B is itself finite. Of course any countable subset $A \subset B$ has no limit point; so A is closed, and for every point $a \in A$, we can choose an open neighborhood U_a such that $U_a \cap A = \{a\}$. Now write

$$X = (X - A) \cup \bigcup_{a \in A} U_a;$$

this is an open cover of X, which must admit a finite subcover. But it admits no proper subcover at all, so A must be finite. □

Example 4.1.4. — To construct this example, we have to appeal to some basic facts about ordinals.

First, let's recall that a *strict well-ordering* on a set A is a well-ordering $<$ that is *irreflexive*, in the sense that for any element $a \in A$, it is not the case that $a < a$. Now we come to von Neumann's definition of *ordinal*: a set A is an *ordinal* if

(4.1.4.1) every element of A is also a subset of A, and
(4.1.4.2) A is strictly well-ordered with respect to set membership.

Put differently, an *ordinal* is a set A such that the following conditions hold.

(4.1.4.1') If $\beta \in A$, then $\beta \subsetneq A$.
(4.1.4.2') If $\beta, \gamma \in A$, then one of the following is true: (1) $\beta = \gamma$, (2) $\beta \in \gamma$, or (3) $\gamma \in \beta$.
(4.1.4.3') If $\varnothing \neq \beta \subsetneq A$, then there exists an element $\beta \in A$ such that $\gamma \cap \beta = \varnothing$.

With this definition, we *define* $0 := \varnothing$, the minimal ordinal. Now recursively, we define, for every integer $n > 0$ the ordinal

$$\mathbf{n} := \{0, \mathbf{1}, \dots, \mathbf{n-1}\}.$$

These are the finite ordinals.

Every ordinal α has a *successor*, namely

$$\alpha + \mathbf{1} := \alpha \cup \{\alpha\}.$$

Some ordinals are not successors; these are called *limit ordinals*. The minimal ordinal 0 is a limit ordinal.

The first infinite ordinal is also the first nonzero limit ordinal; it is

$$\omega := \{0, \mathbf{1}, \dots\}.$$

It has a successor $\omega + \mathbf{1}$, and its successor has a successor $\omega + \mathbf{2}$, etc.

If α and β are ordinals, then we say that $\alpha < \beta$ if $\alpha \in \beta$. Suppose A any set of ordinals; then A is well-ordered by this relation. Consequently, a set A of ordinals is itself an ordinal if and only if, for every element $\alpha \in A$, any ordinal $\beta < \alpha$ is also an element of A. Using this, there is a *supremum* of any set A of ordinals: it is simply the union of the elements of A.

We won't prove it here, but in fact every well-ordered set is order-isomorphic to a unique ordinal in a unique fashion. As a result, if we build a well-ordered set A, there is an ordinal $\mathrm{ord}(A)$ and an order-isomorphism $A \longrightarrow \mathrm{ord}(A)$. In particular, every ordinal α is equal to the set $[0, \alpha)$ of all ordinals less than α.

We can define operations of *ordinal arithmetic* in the following manner. Given two ordinals α and β, we can *concatenate* them by well ordering the disjoint union $\alpha \sqcup \beta$ so that the order within α and β is preserved, and every element of α is strictly less than every element of β. Now set

$$\alpha + \beta := \mathrm{ord}(\alpha \sqcup \beta).$$

Similarly, one can give the product $\alpha \times \beta$ the lexicographic ordering, so that $(a, b) < (a', b')$ if either $a < a'$ or else $a = a'$ and $b < b'$. This is well-ordered, and so one may set

$$\alpha \cdot \beta := \mathrm{ord}(\alpha \times \beta).$$

We can generalize the lexicographic order to a well ordering on the set $\mathrm{Map}(\alpha, \beta)$ of all maps $f : \alpha \longrightarrow \beta$ such that only finitely many elements of α are mapped to an element larger than $0 \in \beta$; one sets

$$\beta^{\alpha} := \mathrm{ord}(\mathrm{Map}(\alpha, \beta)).$$

Note that none of these operations is commutative, though both addition and multiplication are commutative.

If we apply these operations to the finite ordinals and to ω, we get a sequence of ordinals:

$$0, 1, 2, \dots, \omega, \omega+1, \omega+2, \dots, \omega \cdot 2, \omega \cdot 2+1, \omega \cdot 2+2, \dots, \omega^2, \dots, \omega^3, \dots, \omega^{\omega}, \dots, \omega^{\omega^{\omega}}, \dots, \omega^{\omega^{\omega^{\omega}}}, \dots, \varepsilon_0, \dots.$$

The last one listed there was studied by Cantor: it is the smallest ordinal satisfying the equation

$$\varepsilon_0 = \omega^{\varepsilon_0};$$

it can be expressed as the supremum of the sequence

$$\varepsilon_0 = \sup\{\omega, \omega^{\omega}, \omega^{\omega^{\omega}}, \omega^{\omega^{\omega^{\omega}}}, \dots\}.$$

The surprising fact is that these ordinals are all *countable*, because every supremum of a countable set of countable ordinals is countable. So the infinite ordinals all have the same cardinality, but they are non-isomorphic as well-ordered sets.

We can give any of these sets the order topology, but it might not be an interesting enterprise:

Exercise 27. — Show that for any ordinal α, the order topology and the discrete topology coincide if and only if α contains no limit ordinal apart from 0.

Now the supremum of all countable cardinals is an uncountable cardinal ω_1. This is the first uncountable cardinal, and it is going to give us a useful counterexample now and in the future.

Every increasing sequence in ω_1 converges to an element of ω_1. Hence the space ω_1 (when equipped with the order topology) is sequentially compact, but it cannot be compact, since it contains no maximal element. The set $\omega_1 + 1$ is, however, compact.

Exercise 28. — Consider the discrete space $\{0,1\}^{\delta}$. Show that the product

$$X := \prod_{x \in [0,1]} \{0,1\}^{\delta}$$

is compact, but not sequentially compact.

Exercise * ***29.*** — Find a space that is limit point compact, but not countably compact. (Hint: such a space must have the property that it contains a finite subspace that is not closed.)

4.2. Locally compact Hausdorff spaces

The next variant of compactness we want to consider, *local compactness*, is most interesting when coupled with a particular separation axiom, called *Hausdorffness*. We will discuss a variety of separation axioms later, but for now, let us concentrate on this one, since it appears so often.

Definition 4.2.1. — A space X is said to be *Hausdorff* if for any two distinct points $x, y \in X$, there are neighborhoods U of x and V of y such that $U \cap V = \varnothing$.

Proposition 4.2.2. — *A space X is Hausdorff if and only if, for any point $x \in X$, the intersection I_x of all closed neighborhoods of x is the singleton $\{x\}$.*

Proof. — Suppose X Hausdorff, and suppose $x \in X$. Surely $x \in I_x$, and I claim that any point $y \in X - \{x\}$ is not in I. Indeed, for any such point, one may find disjoint open neighborhoods U of x and V of y, whence $X - V$ is a closed neighborhood of x not containing y. Thus $I = \{x\}$.

Conversely, suppose that, for any point $x \in X$, one has $I_x = \{x\}$. Suppose x and y two distinct points of X. Then since $I_x = \{x\}$, there exists a closed neighborhood W of x not containing y, and now the interior W° and the complement $X - W$ are disjoint open neighborhoods of x and y, respectively. □

Exercise 30. — Show that a space X is Hausdorff if and only if the diagonal

$$\Delta_X := \{(x,x) \in X \times X\}$$

is a closed subspace of $X \times X$.

Example 4.2.3. — Any metric space (in particular any euclidean space $\mathbf{R}^n$) is Hausdorff. The cofinite topology on any infinite set is not.

Proposition 4.2.4. — *Any subspace of a Hausdorff space is Hausdorff, and any product of any collection of Hausdorff spaces is Hausdorff.*

Proposition 4.2.5. — *If X and Y are spaces, and Y is Hausdorff, then for any continuous maps $f, g : X \longrightarrow Y$, the* equalizer

$$E_{(f,g)} := \{(x,x') \in X \times X \mid f(x) = g(x')\}$$

is closed in $X \times X$, and

Proof. — The set $E_{(f,g)}$ is the inverse image of Δ_Y under the map $(f,g) : X \times X \longrightarrow Y \times Y$. □

Corollary 4.2.6. — *If S and T are spaces, and T is Hausdorff, then for any continuous map $h : S \longrightarrow T$, the* graph

$$\Gamma_h := \{(s,t) \in S \times T \mid h(s) = t\}$$

is closed in $S \times T$.

Proof. — Apply the proposition to the case when $X = S \times T$, $Y = T$, $f(s,t) = t$, and $g(s,t) = h(s)$ to obtain a closed subspace of $S \times T \times S \times T$, and intersect this with the diagonal $\Delta_{S\times T}$. □

Proposition 4.2.7. — *If X is a Hausdorff space, then any compact subspace $K \subset X$ is closed.*

Proof. — Suppose $K \subset X$ a compact subspace. Write $U = X - K$. Suppose $x \in U$. For any element $y \in K$, choose disjoint open neighborhoods V_y of x and W_y of y. The collection $\{W_y \mid y \in K\}$ is an open cover of K, so there is a finite collection of points $y_1, y_2, \ldots, y_n$ such that the subcollection $\{W_{y_i} \mid i \in \{1,2,\ldots,n\}\}$ covers n. Now the intersection $\bigcap_{i=1}^n V_{y_i}$ is an open neighborhood of x that is disjoint from K, i.e., completely contained in U. □

Exercise 31. — Suppose X a Hausdorff space, and suppose $K, L \subset X$ two disjoint compact subspaces. Show that there exist open sets $U, V \subset X$ such that $K \subset U$, $L \subset V$, and $U \cap V = \varnothing$.

Definition 4.2.8. — A space is said to be *locally compact* if and only if any point has a compact neighborhood.

Example 4.2.9. — Any compact space is of course locally compact. The euclidean spaces $\mathbf{R}^n$ are locally compact, and any discrete space is locally compact.

More exotically, suppose X any set, and suppose $x \in X$. Then the *particular point topology* for (X,x) is the topology in which

$$\mathcal{O}p(X) := \{U \in \mathcal{P}(X) \mid x \in U\} \cup \{\varnothing\}.$$

The particular point topology is always locally compact.

Example 4.2.10. — The space $\mathbf{Q}$ with the subspace topology is not locally compact, since its compact subspaces have empty interior.

Exercise 32. — Suppose X a locally compact Hausdorff space X, and suppose $K \subset X$ a compact subspace. Show that for any point $x \in X - K$, there exists a neighborhood U of x and an open set V containing K such that $U \cap V = \varnothing$, and the closure $\overline{V}$ is compact.

Lemma 4.2.11. — *Suppose X a locally compact Hausdorff space X, and suppose U an open subset that contains a compact subspace K. Then there exists an open set W whose closure is compact such that $K \subset W \subset \overline{W} \subset U$.*

Proof. — If $U = X$, then exercise 32 applies directly. If not, then for every point $x \in X - U$, the exercise 32 gives an open neighborhood V_x of x and an open set W_x whose closure is compact such that

$$K \subset V_x \qquad \text{and} \qquad V_x \cap \overline{W_x} = \varnothing.$$

Now the family $\{(X - U) \cap \overline{W_x} \mid x \in X - U\}$ is a family of compact subspaces of X whose intersection is empty. Consequently there are a finite number of elements $x_1, x_2, \ldots, x_n \in X - U$ such that the intersection $(X - U) \cap \bigcap_{i=1}^n \overline{W_x}$ is empty. Now $W := \bigcap_{i=1}^n W_x$ does the job. □

Proposition 4.2.12. — *Suppose X a Hausdorff space. Then the following are equivalent.*

(4.2.12.1) Every point $x \in X$ has a fundamental system of compact neighborhoods.
(4.2.12.2) Every point $x \in X$ has a closed compact neighborhood.
(4.2.12.3) X is locally compact.

Proof. — The first two conditions imply the third by definition. The third also implies the second, since compact subspaces of a Hausdorff space are closed. Thus it remains only to show that the third condition implies the first. So suppose X is locally compact, and suppose V an open neighborhood of x; I can apply the previous lemma to get an open neighborhood W of x such that $\overline{W}$ is compact and contained in V. □

4.2.13. — We often think of *locally compact and Hausdorff* as a single condition, which we abbreviate as LCH.

Lemma 4.2.14. — *A closed subspace of an LCH space is itself LCH.*

Proof. — The intersection of a closed set and a compact set is compact. (Why does our result follow from this?) □

Exercise 33. — More generally, a subspace of an LCH space X is itself LCH if and only if it can be written as the complement of a closed subspace inside another closed subspace.

4.2.15. — We now turn to the *Alexandroff* or *one-point compactification* of a space X.

Definition 4.2.16. — Suppose X a space. Then define $X^\star$ to be the set $X \sqcup \{\infty\}$, with the following topology: a set $U \subset X^\star$ will be said to be open if and only if either U is an open subset of X, or else U contains ∞, and $X - U$ is closed and compact.

Exercise 34. — Show that this indeed defines a topology on $X^\star$, and the resulting space is compact.

Theorem 4.2.17. — *Suppose X a space. We have the following facts about the Alexandroff compactification.*

(4.2.17.1) The inclusion map $i : X \longrightarrow X^\star$ is continuous and open.
(4.2.17.2) The image of i is a dense open set in X unless X is already compact.
(4.2.17.3) The Alexandroff compactification $X^\star$ is Hausdorff if and only if X is LCH.

Exercise 35. — Prove the previous theorem.

Exercise 36. — Suppose Y an LCH space, and suppose $Z \subset Y$ a closed subspace. Compare the Alexandroff compactification $(Y - Z)^\star$ and the quotient space Y/Z. Are they homeomorphic? Prove it or give a counterexample. If it's true, is the LCH condition necessary? If it's false, can you add conditions on Z and Y so that it will be true?

Exercise 37. — Describe the Alexandroff compactifications of the following spaces in terms of more familiar spaces.

(37.1) $\mathbf{R}^n$.
(37.2) A discrete space.
(37.3) The subspace $\{(n, x) \mid n \in \mathbf{Z} x \in \mathbf{R}\} \subset \mathbf{R}^2$.
(37.4) $\mathbf{Q}$ with the subspace topology.

CHAPTER 5

COUNTABILITY AND SEPARATION

5.1. Taxonomy of separation

Definition 5.1.1. — Suppose X a space. Then two points $x, y \in X$ are *topologically distinguishable* if they do not have the same collection of open neighborhoods. Otherwise, the points are said to be *topologically indistinguishable*.

Exercise 38. — Observe that for any space X, if $x, y \in X$ topologically distinguishable points, then $x \neq y$. Give an example to show that the converse need not hold.

Definition 5.1.2. — Suppose X a space. Then a *neighborhood* of a subset $A \subset X$ is a set containing whose interior contains A.

Definition 5.1.3. — Suppose X a space, and suppose $A, B \subset X$ two subspaces of X.

(5.1.3.1) The sets A and B are said to be *separated* if both $A \cap \overline{B}$ and $\overline{A} \cap B$ are empty.
(5.1.3.2) The sets A and B are said to be *separated by neighborhoods* if there exists disjoint neighborhoods U and V and A and B, respectively.
(5.1.3.3) The sets A and B are said to be *separated by closed neighborhoods* there exists disjoint closed neighborhoods U and V and A and B, respectively.
(5.1.3.4) The sets A and B are said to be *separated by a function* if there exists a continuous map $f : X \longrightarrow [0, 1]$ such that $f|_A = 0$ and $f|_B = 1$.
(5.1.3.5) The sets A and B are said to be *precisely separated by a function* if there exists a continuous map $f : X \longrightarrow [0, 1]$ such that $A = f^{-1}\{0\}$ and $B = f^{-1}\{1\}$.

5.1.4. — Suppose X a space, and suppose $A, B \subset X$. It is easy to see that each of these conditions follows from the next.

(5.1.4.1) A and B are separated.
(5.1.4.2) A and B are separated by neighborhoods.
(5.1.4.3) A and B are separated by closed neighborhoods.
(5.1.4.4) A and B are separated by a function.
(5.1.4.5) A and B are precisely separated by a function.

Definition 5.1.5. — Suppose X a space. Then we introduce the *Trennungsaxiome*.

T_0 One says that X is *Kolmogoroff* or T_0 if any two distinct points are topologically distinguishable.
T_1 One says that X is T_1 if any two distinct points are separated.
T_2 One says that X is *Hausdorff* or T_2 if any two distinct points are separated by neighborhoods.
$T_{2\frac{1}{2}}$ One says that X is *Urysohn* or $T_{2\frac{1}{2}}$ if any two points are separated by closed neighborhoods.
T_3 One says that X is *regular Hausdorff* or T_3 if it is T_0 and if any closed subset $Z \subset X$ and any point $x \notin Z$ are separated by neighborhoods.
$T_{3\frac{1}{2}}$ One says that X is *Tychonoff* or $T_{3\frac{1}{2}}$ if it is T_0 and if any closed subset $Z \subset X$ and any point $x \notin Z$ are separated by a function.

T_4 One says that X is *normal Hausdorff* or T_4 if it is T_1 and if any two closed subsets can be separated by neighborhoods.

T_5 One says that X is *completely normal Hausdorff* or T_5 if it is T_1 and if any two separated sets are separated by neighborhoods.

T_6 One says that X is *perfectly normal Hausdorff* or T_6 if it is T_1 and if any two closed sets are precisely separated by a function.

5.1.6. — For any space X, topological indistinguishability is an equivalence relation on the points of X. The quotient space under this equivalence relation is the *Kolmogoroff quotient* KX of X. This space is automatically Kolmogoroff. This is frequently done in analysis.

Lemma 5.1.7. — *The following are equivalent for a space X.*

(5.1.7.1) X is T_1.
(5.1.7.2) X is T_0, and any two topologially distinguishable points are separated.
(5.1.7.3) Points of X are closed.
(5.1.7.4) Every finite set in X is closed.

Theorem 5.1.8. — *For any $i, j \in \{0,1,2,2\frac{1}{2},3,3\frac{1}{2},4,5,6\}$, if $i \leq j$, then any T_j space is also T_i.*

Example 5.1.9. — Let us convince ourselves that these conditions are distinct.

(5.1.9.1) You have already given an example of a space that is not T_0.

(5.1.9.2) The Sierpiński space Σ_0 is T_0 but not T_1.

(5.1.9.3) The cofinite topology on an infinite set is T_1 but not T_2.

(5.1.9.4) Let $X = \{(x,y) \in \mathbf{Q}^2 \mid y \geq 0\}$, and fix an irrational number θ. Topologize this set in the following manner: for any point $(x,y) \in X$, and any $\varepsilon > 0$, let
$$N_\varepsilon(x,y) = \{(x,y)\} \cup \left[\left(\left(x + \frac{y}{\theta} - \varepsilon, x + \frac{y}{\theta} + \varepsilon\right) \cup \left(x - \frac{y}{\theta} - \varepsilon, x - \frac{y}{\theta} + \varepsilon\right)\right) \cap \mathbf{Q}\right],$$
and contemplate the topology generated by these sets; the sets $N_\varepsilon(x,y)$ form a base for this topology. The closure of each basic neighborhood $N_\varepsilon(x,y)$ contains the union of four strips of slope $\pm\theta$ extending from the intervals $(x + \frac{y}{\theta} - \varepsilon, x + \frac{y}{\theta} + \varepsilon)$ and $(x - \frac{y}{\theta} - \varepsilon, x - \frac{y}{\theta} + \varepsilon)$. Hence the closures of every two open sets intersect. This is a T_2 space that is not $T_{2\frac{1}{2}}$.

(5.1.9.5) Let H be the upper half-plane $\{(x,y) \in \mathbf{R}^2 \mid y \geq 0\}$. Topologize this set in the following manner. Begin with all the open sets in the set $P = \{(x,y) \in \mathbf{R}^2 \mid y > 0\}$, and add in sets of the form $\{p\} \cup (U \cap P)$, where $p = (x,0)$, and U is an open neighborhood of p in the topology on $\mathbf{R}^2$. This is a topology on H; and it is not difficult to see that it is $T_{2\frac{1}{2}}$. However, the complement of a neighborhood of $p = (0,0)$ intersects the closure of every neighborhood of p. So H is not T_3.

(5.1.9.6) The *Tychonoff corkscrew* is the standard example of a space that is T_3 but not $T_{3\frac{1}{2}}$, but it's an absurd space that just makes you sorry you ever tried to contemplate it. Here's one that's slightly simpler (though still pretty absurd), due to A. B. Raha. Let's call it the *Raha space*. For every even integer n, set $T_n := \{n\} \times (-1,1)$, and let $X_1 = \bigcup_{n \text{ even}} T_n$. Now let $(t_k)_{k \geq 1}$ be an increasing sequence of positive real numbers converging to 1. For every odd integer n, set $T_n := \bigcup_{k \geq 1}\{(x,y) \in \mathbf{R}^2 \mid (x-n)^2 + y^2 = t_k^2\}$, and let $X_2 = \bigcup_{n \text{ odd}} T_n$. Now let $X = \{a,b\} \cup \bigcup_{n \in \mathbf{Z}} T_n$. Topologize X so that: (1) every point of X_2 except the points (n, t_k) are isolated; (2) a neighborhood of (n,t_k) consists of all but finitely many elements of $\{(x,y) \in \mathbf{R}^2 \mid (x-n)^2 + y^2 = t_k^2\}$; (3) a neighborhood of a point $(n,y) \in X_1$ consists of all but a finite number of points of $\{(z,y) \mid n-1 < z < n+1\} \cap (T_{n-1} \cup T_n)$; (4) a neighborhood of a is a set U_c containing a and all points of $X_1 \cup X_2$ with x-coordinate greater than a number c; (5) a neighborhood of b is a set V_d containing b and all points of $X_1 \cup X_2$ with x-coordinate less than a number d. This is a space that is T_3, but every continuous map $f : X \longrightarrow \mathbf{R}$ has the property that $f(a) = f(b)$, so it is not $T_{3\frac{1}{2}}$.

(5.1.9.7) Consider the space $\mathrm{Map}(\mathbf{R},\mathbf{R})$ of all maps $\mathbf{R} \longrightarrow \mathbf{R}$ with the topology of pointwise convergence. This space is $T_{3\frac{1}{2}}$ but not T_4.

(5.1.9.8) The *Tychonoff plank* is the space $P := (\omega+1)\times(\omega_1+1)$. It is compact and Hausdorff, and hence it is T_4. It is not, however T_5, because the space $P - \{(\omega,\omega_1)\}$ is not normal. (So you walk the plank ... Get it? It's clever.)

(5.1.9.9) Suppose X an uncountable set, and suppose $x \in X$. Then we have the variant on the particular point topology, which is called the *Fort space topology*. The closed sets are declared to be those sets V such that either V is finite or $x \in V$. Then it is simple to verify that X is T_5, and a short argument shows it cannot be T_6.

The moral of the story is not, of course, to memorize exotic counterexamples. The point is to realize that there are examples that prevent the reversal of any of the implications $T_j \Rightarrow T_i$ for $i \leq j$. Also, in each case the example listed is the simplest one I could find. This suggests something interesting: the differences between $T_{2\frac{1}{2}}$ and T_2 and T_3 is pretty small, as is the difference between $T_{3\frac{1}{2}}$ and T_3. However, the difference between T_4 and $T_{3\frac{1}{2}}$ is pretty significant, contrary to what might have been your expectation. This shows us two relatively bright lines separating the "easy" separation axioms $T_0 - T_1$ from the "midrange" separation axioms $T_2 - T_{3\frac{1}{2}}$, and separating the "midrange" axioms from the "hard" separation axioms $T_4 - T_6$.

5.1.10. — There is another condition, called *sobriety*, that is better suited to applications in algebraic geometry.

Definition 5.1.11. — A space X is said to be *irreducible* (or *hyperconnected*) if no two nonempty open subsets of X are disjoint.

Exercise 39. — Show that the following are equivalent for a space X.
(39.1) X is irreducible.
(39.2) Every nonempty open set of X is dense.
(39.3) X cannot be written as the union of two proper nonempty closed subsets.
(39.4) The interior of every proper closed subset is empty.

Definition 5.1.12. — A point of a space X is said to be *generic* if it is dense in X.

Definition 5.1.13. — A space X is said to be *sober* if every irreducible closed subspace $Z \subset X$ has a unique generic point.

5.1.14. — Every Hausdorff space is sober, and every sober space is Kolmogoroff, but there is no comparison between sobriety and T_1.

Exercise* ***40.*** — Suppose X and Y sober spaces, and consider the posets $(\mathscr{O}p(X),\subset)$ and $(\mathscr{O}p(Y),\subset)$. Show that if $(\mathscr{O}p(X),\subset)$ and $(\mathscr{O}p(Y),\subset)$ are isomorphic as posets, then X and Y are homeomorphic.

5.2. Normality, Urysohn's lemma, and the Tietze extension theorem

Definition 5.2.1. — A space X is said to be *normal* if disjoint closed sets have disjoint neighborhoods.

5.2.2. — A normal T_1 space is normal Hausdorff, or T_4. The results of this section do not require the T_1 property, so we do not include it in our definition.

Lemma 5.2.3. — *Suppose $D \subset \mathbf{R}_{>0}$ a dense subset, and suppose that for each element $t \in D$, there is a subset F_t of a set X such that*
(5.2.3.1) if $t < s$, then $F_t \subset F_s$, and
(5.2.3.2) the sets F_t cover X.
Then for any $x \in X$, let $f(x) = \inf\{t \mid x \in F_t\}$. Then for any element $s \in \mathbf{R}$, one has
$$\{x \in X \mid f(x) < s\} = \bigcup\{F_t \mid t \in D \text{ and } t < s\}$$
and
$$\{x \in X \mid f(x) \leq s\} = \bigcap\{F_t \mid t \in D \text{ and } t > s\}.$$

Proof. — An exercise in set theory (not to be handed in). □

Lemma 5.2.4. — *Suppose now X a space, and suppose $D \subset \mathbf{R}_{>0}$ a dense subset. If for each element $t \in D$, there is an* open *subset $F_t \subset X$ such that*

(5.2.3.1) if $t < s$, then $\overline{F}_t \subset F_s$, and
(5.2.3.2) the sets F_t cover X,
then the function f defined by $f(x) = \inf\{t \mid x \in F_t\}$ is continuous.

Proof. — It's enough to check that the inverse image of the sets $(-\infty, s)$ and $(-\infty, s]$ are open and closed (respectively) for every $s \in \mathbf{R}$.

The previous lemma shows that the inverse images of sets of the first kind, $(-\infty, s)$, are the union of open sets, hence open.

It now follows from the previous lemma that the inverse images of sets of the second kind, $(-\infty, s]$, are closed (and hence we are done) if the following equality obtains:

$$\bigcap\{F_t \mid t \in D \text{ and } t > s\} = \bigcap\{\overline{F}_t \mid t \in D \text{ and } t > s\}.$$

It is clear that the first set is contained in the second. On the other hand, for any $t \in D$ with $t > s$, there is an element $r \in D$ with $s < r < t$, whence $\overline{F}_r \subset F_t$. This completes the proof. □

***Lemma 5.2.5* (Urysohn).** — *Suppose X a normal space, and suppose A and B two disjoint closed subsets of X. Then A and B are separated by functions, i.e., there is a continuous map $f : X \longrightarrow [0,1]$ such that $f|_A = 0$ and $f|_B = 1$.*

Proof. — We wish to apply the previous lemma with the set

$$D = \{x \in \mathbf{Q} \mid x = p2^{-q}, \text{ where } p, q \in \mathbf{N}\}.$$

We define our sets F_t for $t \in D$ thus.

(5.2.5.1) If $t > 1$, then let $F_t = X$.
(5.2.5.2) Let $F_1 = X - B$.
(5.2.5.3) Let F_0 be an open set containing A such that $\overline{F}_0$ is disjoint from B.
(5.2.5.4) For $0 < t < 1$, we proceed by recursion. Write $t = (2m+1)2^{-n}$ and let F_t be an open set containing $\overline{F}_{(2m)2^{-n}}$ whose closure $\overline{F}_t \subset F_{(2m+2)2^{-n}}$ (by normality).

It is plain to see that the conditions of the previous lemma apply, and it constructs our continuous function f. By construction, $f|_A = 0$ and $f|_B = 1$. □

***Theorem 5.2.6* (Tietze extension).** — *Suppose X is a normal space, and suppose A a closed subspace of X. Then any continuous map $A \longrightarrow [-1,1]$ can be extended to a continuous map $X \longrightarrow [-1,1]$.*

CHAPTER 6

METRIZATION THEOREMS

6.1. First and second countability

Definition 6.1.1. — Suppose X a space. Then X is said to satisfy the *first axiom of countability* or to be *first-countable* if for every point $x \in X$ there is a countable fundamental system of neighborhoods of x.

The space X is said to satisfy the *second axiom of countability* or to be *second-countable* if the topology on X has a countable base.

Lemma 6.1.2. — *The second axiom of countability implies the first.*

Example 6.1.3. — The ordinal ω_1 (with the order topology) satisfies the first but not the second axiom of countability, and the ordinal $\omega_1 + 1$ satisfies neither axiom of countability.

Lemma 6.1.4. — *Any subspace of a first-countable (respectively, second-countable) space is first-countable (resp., second-countable), as is any countable product of first-countable (resp., second-countable) spaces. Additionally, any continuous open image of a second-countable space is second-countable.*

Exercise 41. — Give an example demonstrating that an uncountable product of second-countable spaces need not be first-countable.

Exercise 42. — Suppose X a first-countable space. Show that for any subset $A \subset X$, a point $x \in X$ is contained in the closure $\overline{A}$ if and only if there is a sequence of points of A converging to x. Show also that a map $f : X \longrightarrow Y$ to a topological space Y is continuous if and only if, for any convergent sequence $x_n \to x$, the sequence $f(x_n)$ converges to $f(x)$. Finally, show that X is sequentially compact if and only if it is countably compact.

Example 6.1.5. — The real line $\mathbf{R}$ is second countable, as is any countable product of $\mathbf{R}$ with itself.

Exercise 43. — Suppose X a second-countable space; show that the set $\mathscr{O}p(X)$ has cardinality not greater than $\beth_1 = \#\mathbf{R}$ (i.e., the cardinality of the continuum).

Definition 6.1.6. — A space X is said to be *separable* if it has a countable dense subset. A space X is said to be *Lindelöf* if every open cover has a countable subcover.

Theorem 6.1.7. — *Second-countable spaces are both separable and Lindelöf.*

Proof. — Suppose $\mathscr{B} = \{B_1, B_2, \dots\}$ a countable base for a space X.

We show first that X is Lindelöf. Suppose $\mathscr{U}$ an open covering of X. For every natural number j, select, if possible, an element $U_j \in \mathscr{U}$ containing B_j. This defines a countable subset $\mathscr{U}' := \{U_{j_1}, U_{j_2}, \dots\} \subset \mathscr{U}$ indexed by a subset $J \subset \mathbf{N}$. I claim that this is a subcover. Indeed, for any point $x \in X$, an element $U \in \mathscr{U}$ contains x, and so for some j, we have $x \in B_j \subset U$. Hence $j \in J$, whence some $U_j \in \mathscr{U}'$ contains x.

Showing that X is separable is even easier. Simply select a point $x_j \in B_j$ for every $j \in \mathbf{N}$. This set is dense. (Why?) $\square$

Exercise 44. — Suppose X a metrizable space. Show that the following conditions are equivalent on X: (1) X satisfies the second axiom of countability; (2) X is separable; and (3) X is Lindelöf. Check that the Sorgenfrey line is first-countable, separable, and Lindelöf, but not second-countable.

***Lemma 6.1.8* (Tychonoff).** — *Every T_3 Lindelöf space is T_4.*

Proof. — Suppose A and B disjoint closed subsets of X. Since X is regular, for any point $a \in A$ there is a neighborhood of a whose closure fails to intersect B, and thus the family $\mathscr{U}$ of open sets whose closures do not intersect B cover A. Similarly, the family $\mathscr{V}$ of open sets whose closures do not intersect A cover B. Now we form a cover

$$\mathscr{U} \cup \mathscr{V} \cup \{X - (A \cup B)\}$$

of X.

It now follows from the Lindelöf property that there are countable subcovers $\{U_1, U_2, \dots\} \subset \mathscr{U}$ of A and $\{V_1, V_2, \dots\} \subset \mathscr{V}$ of B. Now set

$$U'_r := U_r - \bigcup_{m \leq r} V_m \qquad \text{and} \qquad V'_s := V_s - \bigcup_{m \leq n} U_m$$

for every natural number r and s. One verifies immediately that $U'_r \cap V'_s = \varnothing$ for every pair (r, s) of natural numbers, and so

$$U = \bigcup_{r \geq 1} U'_r \qquad \text{and} \qquad V = \bigcup_{s \geq 1} V'_s$$

are disjoint open neighborhoods of A and B, respectively. □

Exercise 45. — Show that any locally compact Hausdorff space is Tychonoff.

Example 6.1.9. — The product of two Lindelöf spaces need not be Lindelöf; indeed, the product of the Sorgenfrey line with itself is the *Sorgenfrey plane*, which is not Lindelöf. Similarly, a subspace of a Lindelöf space need not be Lindelöf.

6.2. Urysohn's metrization theorem

Exercise* 46. — Prove the following

***Theorem 6.2.1* (Embedding).** — *A space is a Tychonoff space if and only if it is homeomorphic to a subspace of a cube $[0,1]^A$ for some set A.*

Begin by proving the following

***Lemma 6.2.2* (Embedding).** — *Suppose X a space, and suppose $\mathscr{F}$ a family of continuous maps $f : X \longrightarrow Y_f$ (valued in possibly many different spaces).*

(6.2.2.1) The evaluation map is a continuous map $E : X \longrightarrow \prod_{f \in \mathscr{F}} Y_f$.
(6.2.2.2) The evaluation map E is open if for any closed subset $A \subset X$ and any point $x \in X - A$, there exists an element $f \in \mathscr{F}$ such that $f(x) \notin \overline{f(A)}$.
(6.2.2.3) The evaluation map E is injective if and only if for any two distinct point $x, y \in X$, there exists an element $f \in \mathscr{F}$ such that $f(x) \neq f(y)$.

***Theorem 6.2.3* (Urysohn metrization).** — *The following are equivalent for a space X.*

(6.2.3.1) X is second countable and T_3.
(6.2.3.2) X is homeomorphic to a subspace of the Hilbert cube $[0,1]^{\mathbf{N}}$.
(6.2.3.3) X is separable and metrizable.

Proof. — It is not hard to see that the Hilbert cube is separable and metrizable; hence so is any subspace. (Why?)

Any metrizable space is T_3, and any separable metrizable space is second countable. (Check this!)

Hence it remains to show that a T_3 space that satisfies the second axiom of countability is homeomorphic to a subspace of the Hilbert cube. Using the embedding lemma above, it is enough to construct a countable family $\mathscr{F}$ of continuous functions $X \longrightarrow [0,1]$ such that for any closed subset $A \subset X$ and any point $x \in X - A$, there exists an element $f \in \mathscr{F}$ such that $f(x) \notin \overline{f(A)}$. Suppose $\mathscr{B}$ a countable base for the topology of X, and set

$$\mathscr{A} := \{(U,V) \in \mathscr{B} \times \mathscr{B} \mid \overline{U} \subset V\}.$$

One sees that $\mathscr{A}$ is a countable set. Since X is T_3 and Lindelöf, it is also T_4 by Tychonoff's lemma. So, using Urysohn's lemma, choose, for any $(U,V) \in \mathscr{A}$, a Urysohn function $f_{(U,V)} : X \longrightarrow [0,1]$ such that $f|_U = 0$ and $f|_{X-V} = 1$; set

$$\mathscr{F} := \{f_{(U,V)} \mid (U,V) \in \mathscr{A}\}.$$

This is a countable set of functions, and it remains to show that it has the desired properties. Suppose $A \subset X$ a closed subset, and suppose $x \in X - A$. We may now choose a basic open $V \in \mathscr{B}$ such that $x \in U \subset X - A$ and a basic open $U \in \mathscr{B}$ such that $x \in \overline{U} \subset V$. (Why?) Since $(U,V) \in \mathscr{A}$, the corresponding function $f_{(U,V)}$ does the job. □

Exercise 47. — Fill in the details I skipped in the proof of the Urysohn metrization theorem, and write up a complete proof.

6.3. The Bing–Nagata–Smirnov metrization theorem

Definition 6.3.1. — A family $\mathscr{U}$ of subspaces of a topological space X is said to be *locally finite* if every point is contained in a neighborhood that intersects only finitely many elements of $\mathscr{U}$. One says that $\mathscr{U}$ is *discrete* if every point is contained in a neighborhood that intersects only *one* element of $\mathscr{U}$.

A family $\mathscr{V}$ of subspaces of a space X is said to be *σ-locally finite* (respectively, *σ-locally discrete*) if it can be written as the union of a countable collection of locally finite (resp., locally discrete) families of subspaces of X.

***Theorem 6.3.2* (Bing–Nagata–Smirnov metrization).** — *The following are equivalent for a space X.*

(6.3.2.1) X is metrizable.

(6.3.2.2) X is T_3, and there is a σ-locally finite base for X.

(6.3.2.3) X is T_3, and there is a σ-locally discrete base for X.

Outline of proof. — First, observe that the final condition implies the penultimate condition. The first two lemmas here will show that the second condition implies the first.

Lemma 6.3.3. — *A T_3 space whose topology has a σ-locally finite base is T_4.*

Lemma 6.3.4. — *A T_3 space whose topology has a σ-locally finite base is metrizable.*

Lemma 6.3.5. — *Any open cover of a metrizable space has a σ-discrete open refinement.*

It is now an immediate consequence that a metrizable space has a σ-discrete base. □

6.3.6. — One question that might naturally arise here is: so? I can show that if a space satisfies some abstract conditions, then it can be given a metric that gives rise to its topology. Let's list some of the purely topological corollaries of the metrizability of a space X.

(6.3.6.1) A subset $U \subset X$ is open if and only if no sequence of $X - U$ converges to a point of U.

(6.3.6.2) Dually, for any subspace $A \subset X$, a point $x \in X$ is a member of the closure $\overline{A}$ if and only if there is a sequence of points in A that converge to A.

(6.3.6.3) For any space Y, function $f : X \longrightarrow Y$ is continuous if and only if, for any convergent sequence $x_n \to x$ in X, the sequence $f(x_n) \to f(x)$.

(6.3.6.4) A subspace $A \subset X$ is compact if and only if it is countably compact, if and only if it is sequentially compact.

(6.3.6.5) X is T_6.

(6.3.6.6) X is *paracompact*, i.e., every open cover of X has a locally finite open refinement.

ENTR'ACTE

CHAPTER 7

SHEAF THEORY

7.1. Presheaves

Example 7.1.1. — Let us begin right away with a motivational example. Suppose X is a topological space, and let's contemplate, for any open set $U \in \mathscr{O}p(X)$, the set

$$\mathscr{O}_X^{\mathrm{cts}}(U) := \{f : U \longrightarrow \mathbf{R} \mid f \text{ is continuous}\}.$$

It is not difficult to see that if I have a continuous function f on U, then I may restrict it to a continuous function $f|_V$ on any open subset $V \subset U$. This defines a *restriction map*

$$\rho_{V\subset U} : \mathscr{O}_X^{\mathrm{cts}}(U) \longrightarrow \mathscr{O}_X^{\mathrm{cts}}(V).$$

I can, of course, be rather silly in how I restrict: I might have taken $V = U$; then my restriction map is simply the identity.

On the other hand, what if I consider three open sets $W \subset V \subset U$ in X? If I restrict a function f on U to a function on V, and thence to a function on W, I see that I get the same answer as if I had restricted f all the way from U to W. That is, we get an equality of maps

$$\rho_{W\subset V} \circ \rho_{V\subset U} = \rho_{W\subset U}.$$

We can now attempt to study, not the individual sets $\mathscr{O}_X^{\mathrm{cts}}(U)$ or the restriction maps $\rho_{V\subset U}$, but the *entire system* made up of these:

$$\mathscr{O}_X^{\mathrm{cts}} := \left((\mathscr{O}_X^{\mathrm{cts}}(U))_{U\in\mathscr{O}p(X)},\ (\rho_{V\subset U})_{V\subset U\in\mathscr{O}p(X)}\right).$$

Example 7.1.2. — There was nothing particularly special about the topological space $\mathbf{R}$ in the above example. I could try the same trick with any "target" topological space Y: for any open set $U \in \mathscr{O}p(X)$, write

$$\mathscr{O}_X^{Y}(U) := \{f : U \longrightarrow Y \mid f \text{ is continuous}\}.$$

Again, it's pretty clear that we can restrict continuous maps, and so we have a similar structure

$$\mathscr{O}_X^{Y} := \left((\mathscr{O}_X^{Y}(U))_{U\in\mathscr{O}p(X)},\ (\rho_{V\subset U})_{V\subset U\in\mathscr{O}p(X)}\right).$$

Example 7.1.3. — More interestingly, suppose $p : Y \longrightarrow X$ a continuous map. Then we can make a similar definition: for any open set $U \in \mathscr{O}p(X)$, set

$$\Gamma(p)(U) = \Gamma(X/Y)(U) := \{s : U \longrightarrow Y \mid s \text{ is continuous and } p\circ s = \mathrm{id}_U\}.$$

We say that $\Gamma(X/Y)(U)$ is the *set of sections of p over U*.

Definition 7.1.4. — Suppose X a topological space. Then a *presheaf* $\mathscr{F}$ on X is the following data:

(7.1.4.A) for any open set $U \in \mathscr{O}p(X)$, a set $\mathscr{F}(U)$, and

(7.1.4.B) for any open sets $U, V \in \mathscr{O}p(X)$ with $V \subset U$, a *restriction map* $\rho_{V\subset U} : \mathscr{F}(V) \longrightarrow \mathscr{F}(U)$.

These data are subject to the following conditions.

(7.1.4.1) For any open set $U \in \mathscr{O}p(X)$, the map $\rho_{U\subset U}$ is the identity map on $\mathscr{F}(U)$.

(7.1.4.2) For any three open sets $U, V, W \in \mathscr{O}p(X)$ with $W \subset V \subset U$, we have

$$\rho_{W\subset V} \circ \rho_{V\subset U} = \rho_{W\subset U} : \mathscr{F}(U) \longrightarrow \mathscr{F}(W).$$

In other words, the data of a presheaf is a pair

$$\mathscr{F} := \left((\mathscr{F}(U))_{U \in \mathscr{O}p(X)}, (\rho_{V\subset U})_{V\subset U \in \mathscr{O}p(X)} \right),$$

and these data are subject to the axioms $\rho_{U\subset U} = \mathrm{id}$, and $\rho_{W\subset V} \circ \rho_{V\subset U} = \rho_{W\subset U}$.

For an open set $U \in \mathscr{O}p(X)$, an element $s \in \mathscr{F}(U)$ is sometimes called a *section of $\mathscr{F}$ over U*.

7.1.5. — This sort of pattern appears all the time: to any "object" of some collection (in this case $\mathscr{O}p(X)$), you assign a set, and to any "relation" between two objects (in this case set inclusion), you assign a map between the corresponding sets. Sometimes that map goes in the same direction as your relation; sometimes (as here), the map goes in the opposite direction.

When you have these data, then you want to impose some conditions. These conditions always amount to the same thing: when the relation is trivial ($U \subset U$, for example), the corresponding map has to be the identity, and composing relations ($W \subset V \subset U$) gives rise to compositions of maps.

These data, with these conditions, are collectively called a *functor* in general. When the maps go in the same direction as the relation, then the functor is said to be *covariant*; when the maps go in the opposite direction, the functor is said to be *contravariant*. Thus presheaves are particular kinds of functors. We won't say more about this here, but you can expect functors to appear again in your mathematics education!

7.1.6. — One of the most challenging aspects of modern mathematics for beginners is to maintain a clear idea which of a list of bullet points are *data*, and which are *conditions*. In the case of a presheaf $\mathscr{F}$ on a space X, it's a little more subtle than anything we've seen so far: the data are the choice of all the sets $\mathscr{F}(U)$ *and* the choice of all the restriction maps $\rho_{V\subset U} : \mathscr{F}(V) \longrightarrow \mathscr{F}(U)$. It isn't enough to say that there are suitable maps out there; one has to specify them as part of the information of a presheaf.

Example 7.1.7. — Our definition is general enough to allow for some "degenerate" examples of presheaves. For any set S, one may form the *constant presheaf $\mathscr{P}_S$* at S, which assigns to any open set U the set S, and to any open sets $U, V \in \mathscr{O}p(X)$ with $V \subset U$ the identity map on S.

Example 7.1.8. — It's not difficult to see that for any topological space X, the objects $\mathscr{O}_X^Y$ (and hence the object $\mathscr{O}_X^{\mathrm{cts}}$) are presheaves on X.

Example 7.1.9. — Likewise, for any continuous map $p : Y \longrightarrow X$, the object $\Gamma(Y/X)$ is a presheaf on X. Let's contemplate this example in a few particular cases.

(7.1.9.1) Suppose p is a homeomorphism. Then $\Gamma(Y/X)$ is not so very interesting. For any open set $U \in \mathscr{O}p(X)$, there is exactly one section of p over U. It's not hard to see that, in this case, $\Gamma(Y/X)$ is the constant presheaf on one point.

(7.1.9.2) Now if p is not a homeomorphism, but a *local homeomorphism* (so that every point of Y in contained in a neighborhood V such that the restriction of p to V is open and injective), then the story becomes a lot more interesting. First, consider the folding map $\nabla : S^1 \sqcup S^1 \longrightarrow S^1$; in this case $\Gamma(\nabla)$ is somewhat more interesting: to any interval (a, b) it assigns a two-point set. But if the open set is more complicated, the value of $\Gamma(\nabla)$ is more complicated as well. If, for example, U is the disjoint union of two intervals, $\Gamma(\nabla)(U)$ is a four point set. And in general, $\Gamma(\nabla)(V)$ has cardinal 2^n, where n is the number of *connected components* of V. Note, in particular, that $\Gamma(\nabla)(S^1)$ is a two-point set also.

(7.1.9.3) But now consider the map $p : S^1 \longrightarrow S^1$ given by $\xi \longmapsto \xi^2$ (with respect to complex multiplication, thinking of $S^1 \subset \mathbf{C}$). This too is a local homeomorphism, but observe that the presheaf $\Gamma(p)$ has no global sections, i.e., $\Gamma(p)(S^1) = \varnothing$. One way to think of this is to think that there is no single, consistent way to extract square roots of complex numbers. This is quite striking: the fiber over each point of p is two points, just as in the previous example. But the presheaf $\Gamma(p)$ has shown us that there is something nontrivial about our map p.

(7.1.9.4) Here's yet another example in the same vein: consider the exponential map $\exp : \mathbf{C} \longrightarrow \mathbf{C}^\times := \mathbf{C} - \{0\}$. We can contemplate the presheaf $\Gamma(\exp)$, and we see immediately that over a sufficiently small open neighborhood U of any point in $\mathbf{C}^\times$, there are countably infinitely many elements of $\Gamma(\exp)(U)$; these correspond to the *branches* of the logarithm: if we write a complex number $z = r\exp(i\theta)$, then we might write $\log(z) = \log|r| + i\theta$; however, this is ambiguous, I can add $2k\pi$ to θ. Once again, however, one sees that one cannot make a consistent choice of a branch of the logarithm: there are no global sections of $\Gamma(\exp)$!

Example 7.1.10. — Suppose X a space, and suppose $A \in \mathscr{O}p(X)$ a particular *fixed* open set. We have an associated presheaf h_A defined by the rule

$$h_A(U) := \begin{cases} \star & \text{if } U \subset A \\ \varnothing & \text{else.} \end{cases}$$

In this case, the restriction maps are forced on us: for any open sets $V, U \in \mathscr{O}p(X)$ with $V \subset U$, there is necessarily only one map $h_A(U) \longrightarrow h_A(V)$. Observe that we cannot expect this in general: ordinarily, we'll have to put some effort into describing the restriction maps!

This gives us a way (in fact, a very powerful way) of "representing" open sets $A \in \mathscr{O}p(X)$ as presheaves; consequently, the presheaves h_A are known as *representable presheaves*, or presheaves *represented* by $A \in \mathscr{O}p(X)$.

Exercise 48. — Suppose X a space, $A \in \mathscr{O}p(X)$ an open subset. Suppose I attempted to define a presheaf h^A the other way around from our representable friend:

$$h^A(U) := \begin{cases} \varnothing & \text{if } U \subset A \\ \star & \text{else.} \end{cases}$$

What goes wrong here? Why does this not define a presheaf?

Example 7.1.11. — Suppose S is a set, and suppose X a space with a distinguished point $x \in X$. Then the *skyscraper presheaf at x with value S* is defined by the rule

$$S_x(U) := \begin{cases} S & \text{if } x \in U \\ \star & \text{else.} \end{cases}$$

7.1.12. — We can also ask for additional structure on a presheaf $\mathscr{F}$. If all the sets $\mathscr{F}(U)$ are equipped with a group structure, and all the restriction maps $\rho_{V\subset U}$ are group homomorphisms, then the presheaf $\mathscr{F}$ is said to be a *presheaf in groups*. Similarly, if all the sets $\mathscr{F}(U)$ are equipped with an abelian group (respectively, ring, vector space, algebra, ...) structure, and all the restriction maps $\rho_{V\subset U}$ are group (resp., ring, vector space, algebra, ...) homomorphisms, then the presheaf $\mathscr{F}$ is said to be a *presheaf in abelian groups* (resp., *rings*, *vector spaces*, *algebras*, ...). The key point here is that both the sets $\mathscr{F}(U)$ and the maps $\rho_{V\subset U}$ carry this added structure.

Example 7.1.13. — Suppose $X = \mathbf{R}^n$ with the euclidean topology. Then for any open set $U \subset \mathbf{R}^n$, we can contemplate the set

$$\mathscr{O}_X^{\mathrm{an}}(U) = \mathscr{C}^\omega(U) := \{f : U \longrightarrow \mathbf{R} \mid f \text{ is analytic}\}.$$

(Recall that a function f on an open set U is analytic if it is infinitely differentiable and, for any point $x \in U$, there is a neighborhood of x in U such that the Taylor series of f at x converges to f. This is surely a presheaf as well; moreover, since the sum, difference, and product of any two analytic functions is analytic, we can see that $\mathscr{O}_X^{\mathrm{an}}$ is a sheaf of rings (even of $\mathbf{R}$-algebras).

Definition 7.1.14. — A *morphism* or simply *map* of presheaves $\phi : \mathscr{F} \longrightarrow \mathscr{G}$ is the following data: for any open set $U \in \mathscr{O}p(X)$, a map (called the *component* of ϕ)

$$\phi_U : \mathscr{F}(U) \longrightarrow \mathscr{G}(U),$$

subject to the following condition: for any open sets $U, V \in \mathcal{O}p(X)$ with $V \subset U$, the following diagram commutes:

$$\begin{array}{ccc} \mathscr{F}(U) & \xrightarrow{\phi_U} & \mathscr{G}(U) \\ {\scriptstyle \rho_{V\subset U}}\downarrow & & \downarrow{\scriptstyle \rho_{V\subset U}} \\ \mathscr{F}(V) & \xrightarrow[\phi_V]{} & \mathscr{G}(V), \end{array}$$

that is, for any section $s \in \mathscr{F}(U)$, one has

$$\rho_{V\subset U}(\phi_U(s)) = \phi_V(\rho_{V\subset U}(s)).$$

Given morphisms of presheaves $\phi : \mathscr{F} \longrightarrow \mathscr{G}$ and $\psi : \mathscr{G} \longrightarrow \mathscr{H}$, we can form the composite $\psi \circ \phi : \mathscr{F} \longrightarrow \mathscr{H}$ in the following manner: for any open set $U \in \mathcal{O}p(X)$, set

$$(\psi \circ \phi)_U = \psi_U \circ \phi_U.$$

This defines a morphism of presheaves $\mathscr{F} \longrightarrow \mathscr{H}$ as desired.

A morphism of presheaves $\phi : \mathscr{F} \longrightarrow \mathscr{G}$ is said to be an *isomorphism* if there exists a morphism of sheaves $\psi : \mathscr{G} \longrightarrow \mathscr{F}$ such that both $\psi \circ \phi = \mathrm{id}$ and $\phi \circ \phi = \mathrm{id}$.

Exercise 49. — (49.1) Check that our description of the composite of two morphisms of presheaves is indeed again a morphism of presheaves.

(49.2) Verify that if we have three morphisms of presheaves $\phi : \mathscr{F} \longrightarrow \mathscr{G}$, $\psi : \mathscr{G} \longrightarrow \mathscr{H}$, and $\chi : \mathscr{H} \longrightarrow \mathscr{K}$, then one has

$$(\chi \circ \psi) \circ \phi = \chi \circ (\psi \circ \phi).$$

(49.3) Finally, show that a morphism of presheaves $\phi : \mathscr{F} \longrightarrow \mathscr{G}$ is an isomorphism if and only if, for any open set $U \in \mathcal{O}p(X)$, the components ϕ_U are all bijections.

Definition 7.1.15. — We say that a morphism $\mathscr{F} \longrightarrow \mathscr{G}$ of presheaves is *injective* if for every open set $U \in \mathcal{O}p(X)$, the corresponding map $\mathscr{F}(U) \longrightarrow \mathscr{G}(U)$ is injective. (We'll need to be a little more cautious when we define surjectivity.)

Notation 7.1.16. — For any space X, and any presheaves $\mathscr{F}$ and $\mathscr{G}$, write $\mathrm{Mor}_X(\mathscr{F}, \mathscr{G})$ for the set of all morphisms of presheaves $\mathscr{F} \longrightarrow \mathscr{G}$.

Exercise 50. — Show that for any space X, and any pair of open sets $U, V \in \mathcal{O}p(X)$ with $V \subset U$, there is a *unique* morphism of presheaves $h_V \longrightarrow h_U$.

Exercise* 51. — Show that for any space X, any open set $U \in \mathcal{O}p(X)$, and any presheaf $\mathscr{F}$ on X, there is a natural bijection

$$\Phi_U : \mathscr{F}(U) \cong \mathrm{Mor}_X(h_U, \mathscr{F}).$$

This word *natural* gets thrown around a lot. In this case, we can say exactly what we mean by it: for any pair of open sets $U, V \in \mathcal{O}p(X)$ with $V \subset U$, we have our unique morphism of presheaves $j : h_V \longrightarrow h_U$. Composition with j induces a map $- \circ j : \mathrm{Mor}_X(h_U, \mathscr{F}) \longrightarrow \mathrm{Mor}_X(h_V, \mathscr{F})$. (How?) For an extra challenge, show that in this case, the following diagram of sets commutes:

$$\begin{array}{ccc} \mathscr{F}(U) & \xrightarrow{\Phi_U} & \mathrm{Mor}_X(h_U, \mathscr{F}) \\ {\scriptstyle \rho_{U\subset V}}\downarrow & & \downarrow{\scriptstyle -\circ j} \\ \mathscr{F}(V) & \xrightarrow[\Phi_V]{} & \mathrm{Mor}_X(h_V, \mathscr{F}). \end{array}$$

Exercise 52. — Show that for any two spaces X and Y, the presheaf $\mathcal{O}_X^Y$ of local Y-valued functions is isomorphic to the presheaf $\Gamma(Y \times X/X)$ of local sections of the projection map $Y \times X \longrightarrow X$.

7.2. Stalks

Definition 7.2.1. — Suppose X a space, and suppose $\mathscr{F}$ a presheaf on X. Then for any point $x \in X$, consider the set
$$\coprod_{x \in U \in \mathscr{O}p(X)} \mathscr{F}(U) = \{(U,s) \mid x \in U \in \mathscr{O}p(X),\ s \in \mathscr{F}(U)\}.$$
On this set we may impose an equivalence relation $\sim$ in the following manner. For any two elements (U,s) and (V,t), we say that $(U,s) \sim (V,t)$ if and only if there exists an open neighborhood $W \subset U \cap V$ of x such that $\rho_{W \subset U}(s) = \rho_{W \subset V}(t)$. Now define the *stalk* of $\mathscr{F}$ at x to be the set
$$\mathscr{F}_x := \left(\coprod_{x \in U \in \mathscr{O}p(X)} \mathscr{F}(U) \right) \Big/ \sim .$$
The equivalence class of a section s under this equivalence relation is called the *germ* of s, and is dented s_x.

Example 7.2.2. — Suppose $p : Y \longrightarrow X$ a *covering space*, i.e., a surjective continuous map with the property that every point $x \in X$ is contained in an open neighborhood U such that $p^{-1}U$ is a disjoint union $\coprod_{\alpha \in \Lambda} V_\alpha$ of open subsets $V_\alpha \subset Y$ such that $p|_{V_\alpha} : V_\alpha \longrightarrow U$ is a homeomorphism. It is easy to see that such a map must be a local homeomorphism, but here we have required much more: we have required that we can choose the neighborhoods of various points in a fiber in a consistent manner.

For any point $x \in X$, we can compute the stalk $\Gamma(Y/X)_x$ easily: there exists an open neighborhood U of x whose inverse image is homeomorphic (via p) to a disjoint union of copies of U. But now it's easy to see that a section of $\Gamma(Y/X)$ over this open set U is completely determined by where it sends out point x. For every sub-neighborhood $V \subset U$, we'll be in the same situation, so the equivalence relation doesn't change anything from this point on as the neighborhoods get smaller. Hence the stalk $\Gamma(Y/X)_x$ is simply the fiber $p^{-1}(x)$ again.

We will undertake a more systematic study of coverings spaces very soon.

Exercise 53. — Suppose $p : Y \longrightarrow X$ *any* continuous map. Is it true that for any point $x \in X$, the stalk $\Gamma(Y/X)_x$ is the fiber $p^{-1}(x)$? Prove or give a counterexample.

Example 7.2.3. — Let us consider the case of two topological spaces X and Y, and let us consider the presheaf $\mathscr{O}_X^Y$ of local continuous functions on X valued in Y. What is the stalk of this presheaf at a point $x \in X$? It's difficult to say in familiar language. The elements of $\mathscr{O}_{X,x}^Y$ are equivalence classes of continuous functions defined in a neighborhood of x, where we declare two such functions to be equivalent when there is a neighborhood of x on which they agree. These equivalence classes are known as *germs* of continuous functions.

Example 7.2.4. — Suppose $X = \mathbf{R}^n$, and consider the presheaf $\mathscr{O}_X^{\mathrm{an}}$ of local analytic functions on X. For any point $x \in X$, the stalk $\mathscr{O}_{X,x}^{\mathrm{an}}$ is the set of all power series at x with a positive radius of convergence. This is a ring, sometimes denoted $\mathbf{R}\{x_1, x_2, \ldots, x_n\}$.

7.3. Sheaves

Example 7.3.1. — To motivate the sheaf axiom, let's briefly recall a little lemma we proved in Exercise 18 some time back. If X is a space with an open covering $\mathscr{V}$, and if for any $V \in \mathscr{V}$ we are given a continuous function $f_V : V \longrightarrow Y$ with the property that for any elements $V, W \in \mathscr{V}$, one has
$$f_V|_{V \cap W} = f_W|_{V \cap W} : V \cap W \longrightarrow Y,$$
then there exists a unique continuous map $f : X \longrightarrow Y$ such that for any $V \in \mathscr{V}$,
$$f|_V = f_V : V \longrightarrow Y.$$

7.3.2. — Suppose $\mathscr{F}$ a presheaf on a space X. For any open set U, and any cover $\{U_\alpha\}_{\alpha \in \Lambda}$ of U, we can contemplate the following diagram:
$$\mathscr{F}(U) \longrightarrow \prod_{\alpha \in \Lambda} \mathscr{F}(U_\alpha) \rightrightarrows \prod_{\beta,\gamma \in \Lambda} \mathscr{F}(U_\beta \cap U_\gamma). \tag{7.3.2.1}$$

Let us describe these maps in more detail. The first map is the assignment

$$\begin{aligned} \mathscr{F}(U) &\longrightarrow \prod_{\alpha\in\Lambda} \mathscr{F}(U_\alpha) \\ s &\longmapsto \left(\rho_{U_\alpha\subset U}(s)\right)_{\alpha\in\Lambda}. \end{aligned}$$

The other two maps are the assignments

$$\begin{aligned} \prod_{\alpha\in\Lambda} \mathscr{F}(U_\alpha) &\longrightarrow \prod_{\beta,\gamma\in\Lambda} \mathscr{F}(U_\beta\cap U_\gamma) \\ (s_\alpha)_{\alpha\in\Lambda} &\longmapsto \left(\rho_{U_\beta\cap U_\gamma\subset U_\beta}(s_\beta)\right)_{\beta,\gamma\in\Lambda} \end{aligned} \qquad \text{and} \qquad \begin{aligned} \prod_{\alpha\in\Lambda} \mathscr{F}(U_\alpha) &\longrightarrow \prod_{\beta,\gamma\in\Lambda} \mathscr{F}(U_\beta\cap U_\gamma) \\ (s_\alpha)_{\alpha\in\Lambda} &\longmapsto \left(\rho_{U_\beta\cap U_\gamma\subset U_\gamma}(s_\gamma)\right)_{\beta,\gamma\in\Lambda} \end{aligned}$$

We can ask whether the diagram (7.3.2.1) is an *equalizer*, i.e., whether the map $\mathscr{F}(U)\longrightarrow\prod\mathscr{F}(U_\alpha)$ is an injective map whose image consists of *exactly* those elements that have the same values under the maps $\prod\mathscr{F}(U_\alpha)\longrightarrow\prod\mathscr{F}(U_\beta\cap U_\gamma)$. In other words, (7.3.2.1) is an equalizer if and only if the following condition is satisfied: for every $\alpha\in\Lambda$ there is a section $s_\alpha\in\mathscr{F}(U_\alpha)$ such that for any pair $\beta,\gamma\in\Lambda$, the restrictions

$$\rho_{U_\beta\cap U_\gamma\subset U_\beta}(s_\beta)=\rho_{U_\beta\cap U_\gamma\subset U_\gamma}(s_\gamma)$$

coincide, then there exists a unique section $s\in\mathscr{F}(U)$ such that $\rho_{U_\alpha\subset U}(s)=s_\alpha$.

Definition 7.3.3. — A presheaf $\mathscr{F}$ on a space X is said to be a *sheaf* if, for any open set $U\in\mathscr{O}p(X)$ and any open cover $\{U_\alpha\}_{\alpha\in\Lambda}$ of U, the diagram

$$\mathscr{F}(U)\longrightarrow\prod_{\alpha\in\Lambda}\mathscr{F}(U_\alpha)\rightrightarrows\prod_{\beta,\gamma\in\Lambda}\mathscr{F}(U_\beta\cap U_\gamma).$$

is an equalizer in the sense above.

7.3.4. — This is a little tricky, so let's discuss this. A sheaf $\mathscr{F}$ on a space X is a presheaf that has a powerful *local-to-global* property.

To illustrate this local-to-global property, consider the case of an open cover with just two open sets. Suppose that $U\subset X$ is an open set, which we have written as a union $U=V\cup W$ for some open sets $V,W\subset X$. Then the sheaf condition for $\mathscr{F}$ says that if we are given sections $s_V\in\mathscr{F}(V)$ and $s_W\in\mathscr{F}(W)$ that *match up* on $V\cap W$ (so that $\rho_{V\cap W\subset V}(s_V)=\rho_{V\cap W\subset W}(s_W)$), then there is one and only one set $s\in\mathscr{F}(U)$ who restrictions to V and to W are s_V and s_W, respectively.

What does this say? It says that *sections glue uniquely*. In other words, specifying a section of $\mathscr{F}$ over an open set $U\subset X$ is the same specifying a compatible family of sections of $\mathscr{F}$ over each element of an open cover of $\mathscr{F}$.

Example 7.3.5. — Our first example of this section shows that, indeed, for any two topological spaces X and Y, the presheaf $\mathscr{O}_X^Y$ is a sheaf, called the *sheaf of local continuous functions on X with values in Y*.

Example 7.3.6. — More generally, one can see that for any continuous map $p:Y\longrightarrow X$, the associated presheaf $\Gamma(Y/X)$ is indeed a sheaf, called the *sheaf of local sections of p*.

Example 7.3.7. — For an space X, any point $x\in X$, and any set S, the skyscraper presheaf S_x is a sheaf as well.

Exercise 54. — Show that for any sheaf $\mathscr{F}$, the map

$$\mathscr{F}(U)\longrightarrow\textstyle\prod_{x\in U}\mathscr{F}_x$$

is injective. Give a counterexample to show that this is not true of presheaves.

Lemma 7.3.8. — *Suppose $\mathscr{F}$ a sheaf on a space X. Then $\mathscr{F}(\varnothing)=\star$.*

Lemma 7.3.9. — *Suppose $\mathscr{F}$ and $\mathscr{G}$ sheaves on a space X. Then if $f,g:\mathscr{F}\longrightarrow\mathscr{G}$ are two morphisms such that for every point $x\in X$, the induced maps $f_x,g_x:\mathscr{F}_x\longrightarrow\mathscr{G}_x$ on stalks coincide (so that $f_x=g_x$), then $f=g$.*

Proof. — This follows immediately from the previous exercise. $\square$

Proposition 7.3.10. — *Suppose $\mathscr{F}$ and $\mathscr{G}$ sheaves on a space X. Then a morphism $\phi : \mathscr{F} \longrightarrow \mathscr{G}$ is an isomorphism, or an injection, if and only if, for every point $x \in X$, the map on stalks $\mathscr{F}_x \longrightarrow \mathscr{G}_x$ is so.*

Proof. — The injectivity statement follows from the fact that if the composite of two maps is injective, then the first of the two must also be injective.

Now suppose $U \in \mathscr{O}p(X)$ any open set, and consider the component $\phi_U : \mathscr{F}(U) \longrightarrow \mathscr{G}(U)$. The claim is that it is surjective. To demonstrate this, suppose $s \in \mathscr{G}(U)$ a section. For every point $x \in U$, the element $(U, s) \in \mathscr{G}_x$ has an inverse image in $\mathscr{F}_x$, i.e., for each point $x \in U$ there is a pair (V_x, t_x), consisting of an open neighborhood $V_x \subset U$ of x along with a section $t_x \in \mathscr{F}(V_x)$ such that $\phi_{V_x}(t_x) = \rho_{V_x \subset U}(s)$. Since ϕ is injective, for any two points $x, y \in U$, we have

$$\rho_{V_x \cap V_y \subset V_x}(t_x) = \rho_{V_x \cap V_y \subset V_y}(t_y).$$

The sheaf axiom applied to the open cover $\{V_x \mid x \in U\}$ of U now implies that there is a global section $t \in \mathscr{F}(U)$ such that $t_x = \rho_{V_x \subset U}(t)$. And now $\phi(t) \in \mathscr{G}(U)$ is the unique section whose restriction to any V_x is $\phi_{V_x}(t_x) = \rho_{V_x \subset U}(s)$, namely, s itself! □

Warning 7.3.11. — Note that what we have *not* said is that two sheaves $\mathscr{F}$ and $\mathscr{G}$, all of whose stalks are isomorphic, are themselves isomorphic. This is *false*. We must have a morphism $\mathscr{F} \longrightarrow \mathscr{G}$ inducing the stalkwise isomorphism in order to make this conclusion. This sounds like a minor point, but it is extremely important.

Indeed, recall our two examples of *double covers* of the circle: we were given two local homeomorphisms $\nabla : S^1 \sqcup S^1 \longrightarrow S^1$ and $p : S^1 \longrightarrow S^1$. The former map is simply the folding map, and the second is $z \longmapsto z^2$. We saw that, while $\Gamma(\nabla)$ was simply the constant sheaf on a two point set, $\Gamma(p)$ was a more subtle object, with no global sections. In this case, the stalk over any point $x \in S^1$ is always the same: simply two points. But the two sheaves differ dramatically! The point is that there is no morphism of sheaves between $\Gamma(\nabla)$ and $\Gamma(p)$ inducing bijections on the stalks. This situation is quite typical, and it's one that sheaves are particularly well-adapted to.

Exercise 55. — Suppose n a positive integer. Consider the continuous map $p_n : S^1 \longrightarrow S^1$ given by $z \longmapsto z^n$ (where again we think of S^1 as embedded in $\mathbf{C}$, where we may use complex multiplication). Consider also the n-fold folding map $\nabla^n : S^1 \sqcup S^1 \sqcup \cdots \sqcup S^1 \longrightarrow S^1$. Prove that $\Gamma(p_n)$ and $\Gamma(\nabla^n)$ are both sheaves, that their stalks at any point are isomorphic, but the sheaves themselves are not isomorphic.

7.4. The *espace étalé* and sheafification

Definition 7.4.1. — A continuous map $p : Y \longrightarrow X$ is said to be a *local homeomorphism* if every point $y \in Y$ is contained in a neighborhood V such that p is open and injective.

7.4.2. — Suppose X a space and $\mathscr{F}$ a presheaf on X. Consider the set $\acute{\mathrm{E}}\mathrm{t}(\mathscr{F}) := \coprod_{x \in X} \mathscr{F}_x$; there is an obvious map $p_{\mathscr{F}} : \acute{\mathrm{E}}\mathrm{t}(\mathscr{F}) \longrightarrow X$ whose fibers are precisely the stalks of $\mathscr{F}$. For any open set U and any section $s \in \mathscr{F}(U)$, there is a corresponding map $\sigma_s : U \longrightarrow \acute{\mathrm{E}}\mathrm{t}(\mathscr{F})$, $x \longmapsto s_x$, such that $p \circ s = \mathrm{id}$. (This is the origin of the term "section.")

Definition 7.4.3. — Suppose X a space and $\mathscr{F}$ a presheaf on X. The *espace étalé* of $\mathscr{F}$ is the set $\acute{\mathrm{E}}\mathrm{t}(\mathscr{F})$ equipped with the finest topology such that for any section $s \in \mathscr{F}(U)$, the corresponding map $\sigma_s : U \longrightarrow \acute{\mathrm{E}}\mathrm{t}(\mathscr{F})$ is continuous. That is, we declare a subset $V \subset \acute{\mathrm{E}}\mathrm{t}(\mathscr{F})$ to be open if and only if, for any open set $U \in \mathscr{O}p(X)$ and any section $s \in \mathscr{F}(U)$, the inverse image $\sigma_s^{-1}(V)$ is open in U.

Exercise 56. — Show that for any space X and any presheaf $\mathscr{F}$ on X, the natural morphism $p_{\mathscr{F}} : \acute{\mathrm{E}}\mathrm{t}(\mathscr{F}) \longrightarrow X$ is a local homeomorphism.

Exercise 57. — Suppose S a set, and suppose $\mathscr{F}_S$ is the constant presheaf at S on a space X. Then show that the éspace étalé of $\mathscr{F}_S$ is the fold map $\nabla^{(S)} : X^{\sqcup(S)} \longrightarrow X$, where $X^{\sqcup(S)}$ is the S-fold disjoint union of X with itself:

$$X^{\sqcup(S)} := \coprod_{s \in S} X = X \times S^{\delta}.$$

Lemma 7.4.4. — *Suppose $p : Y \longrightarrow X$ a local homeomorphism. Then the éspace étalé of the sheaf $\Gamma(Y/X)$ is canonically homeomorphic* over *X to Y. That is, there is a unique homeomorphism $Y \longrightarrow \acute{\mathrm{E}}\mathrm{t}(\Gamma(Y/X))$ such that the diagram*

$$\begin{tikzcd} \acute{\mathrm{E}}\mathrm{t}(\Gamma(Y/X)) \arrow[rr, "w"] \arrow[dr, "p_{\Gamma(Y/X)}"'] & & Y \arrow[dl, "p"] \\ & X & \end{tikzcd}$$

commutes.

Proof. — Suppose $s \in \acute{\mathrm{E}}\mathrm{t}(\Gamma(Y/X))$ a germ of a section near $x \in X$. Then its value $w(s) = s(x)$ at x is well-defined. This defines a map $w : \acute{\mathrm{E}}\mathrm{t}(\Gamma(Y/X)) \longrightarrow Y$ such that $p \circ w = p_{\Gamma(Y/X)}$.

In the other direction, suppose $y \in Y$, and suppose $x = p(y)$. Then there exist open neighborhoods U of x and V of y such that p restricts to a homeomorphism $V \longrightarrow U$; denote by $s : U \longrightarrow V$ its inverse, and define $v(y) \in \Gamma(Y/X)_x$ as the germ of the section s. The class $v(y)$ does not depend upon the choice of U or V. This defines an inverse $v : Y \longrightarrow \acute{\mathrm{E}}\mathrm{t}(\Gamma(Y/X))$ to the map w.

It remains to show that w is continuous and open. We leave this to the reader. ☐

Definition 7.4.5. — Suppose X a space and $\mathscr{F}$ a presheaf on X. The *sheafification* of $\mathscr{F}$ is the sheaf $a\mathscr{F} := \Gamma(\acute{\mathrm{E}}\mathrm{t}(\mathscr{F})/X)$ of sections of the projection map $p : \acute{\mathrm{E}}\mathrm{t}(\mathscr{F}) \longrightarrow X$. The canonical morphism of presheaves $\mathscr{F} \longrightarrow a\mathscr{F}$ that assigns to any section $s \in \mathscr{F}(U)$ the local section $x \longmapsto s_x$ is called the *unit morphism*.

Proposition 7.4.6. — *For any presheaf $\mathscr{F}$ on a space X, the natural morphism $\mathscr{F} \longrightarrow a\mathscr{F}$ induces an isomorphism on all stalks.*

Proof. — By 7.4.4, the stalks of $a\mathscr{F}$ correspond under the unit morphism to the fibers of $\acute{\mathrm{E}}\mathrm{t}(\mathscr{F}) \longrightarrow X$, which in turn are the stalks of $\mathscr{F}$. ☐

Corollary 7.4.7. — *For any sheaf $\mathscr{F}$ on a space X, the unit morphism $\mathscr{F} \longrightarrow a\mathscr{F}$ is an isomorphism.*

Corollary 7.4.8. — *There is a bijective correspondence between isomorphism classes of sheaves on X and homeomorphism classes of local homeomorphisms $Y \longrightarrow X$.*

Exercise 58. — Recall that on a space X, the presheaf $\mathscr{O}_X^{\mathrm{bdd}}$ of local bounded continuous functions (valued in $\mathbf{R}$) may not be a sheaf. Describe its sheafification.

Example 7.4.9. — The *constant sheaf* $\mathscr{F}_S$ at a set S on a space X is the sheafification of the constant presheaf $\mathscr{P}_S$ at S. Observe that the constant sheaf is not really constant: it can take many different values on an open set $U \subset X$.

Theorem 7.4.10. — *Suppose X a space, suppose $\mathscr{F}$ a presheaf on X, and suppose $\mathscr{G}$ a sheaf on X. Then any morphism $\mathscr{F} \longrightarrow \mathscr{G}$ factors uniquely through the unit. In other words, for any morphism $\mathscr{F} \longrightarrow \mathscr{G}$, there exists a unique morphism $a\mathscr{F} \longrightarrow \mathscr{G}$ such that the diagram*

$$\begin{tikzcd} \mathscr{F} \arrow[rr] \arrow[dr] & & \mathscr{G} \\ & a\mathscr{F} \arrow[ur] & \end{tikzcd}$$

commutes.

Proof. — Let us show that unicity is guaranteed once the existence of the morphism $a\mathscr{F} \longrightarrow \mathscr{G}$ satisfying the diagram above is confirmed. Indeed, any two morphisms $a\mathscr{F} \longrightarrow \mathscr{G}$ satisfying the diagram above must agree an all stalks, since the unit morphism $\mathscr{F} \longrightarrow a\mathscr{F}$ is a bijection on all stalks. Now we may employ 7.3.9 to see that the morphisms $a\mathscr{F} \longrightarrow \mathscr{G}$ coincide.

To verify existence, observe that we can simply apply a to the morphism $\mathscr{F} \longrightarrow \mathscr{G}$ to get a morphism of sheaves $a\mathscr{F} \longrightarrow a\mathscr{G}$; since the unit morphism $\mathscr{G} \longrightarrow a\mathscr{G}$ is an isomorphism, we can compose $a\mathscr{F} \longrightarrow a\mathscr{G}$ with the inverse $a\mathscr{G} \longrightarrow \mathscr{G}$ to get the desired morphism $a\mathscr{F} \longrightarrow \mathscr{G}$. ☐

7.4.11. — There is a more sophisticated way of talking about all this. The construction of the *espace étalé* is a functor $\acute{\mathrm{E}}\mathrm{t}$ from the category $\mathscr{P}re\mathscr{S}h(X)$ of presheaves on a space X to the category $(\mathscr{L}\mathscr{H}/X)$ of local homeomorphisms to X. Now $(\mathscr{L}\mathscr{H}/X)$ is actually equivalent (via the sections sheaf functor Γ) to the category $\mathscr{S}h(X)$ of sheaves on X. Now the composite $a = \Gamma \circ \acute{\mathrm{E}}\mathrm{t}$ is the sheafification functor, which is left adjoint to the inclusion $\mathscr{S}h(X) \subset \mathscr{P}re\mathscr{S}h(X)$.

7.5. Direct and inverse images of sheaves

Definition 7.5.1. — For any two continuous maps $f : X \longrightarrow Y$ and $h : Z \longrightarrow Y$, the *fiber product* $X \times_Y Z$ is the set

$$X \times_Y Z := \{(x,z) \in X \times Z \mid f(x) = h(z)\} \subset X \times Z,$$

equipped with the subspace topology. (Note that the notation here is ambiguous, since it makes no explicit mention of the maps f and h.) The (continuous) projection $X \times_Y Z \longrightarrow X$ is called the *pullback* of the map h.

Exercise 59. — For any map $f : X \longrightarrow Y$, contemplate the continuous map $\Delta_f : X \longrightarrow X \times_Y X$ given by the assignment $x \longmapsto (x,x)$. Show that f is a local homeomorphism if and only if both f and Δ_f are open maps.

Exercise 60. — Show that the pullback of a local homeomorphism is a local homeomorphism.

Definition 7.5.2. — Suppose $f : X \longrightarrow Y$ a continuous map.

(7.5.2.1) For any sheaf $\mathscr{F}$ on X, define the *direct image* $f_\star\mathscr{F}$ of $\mathscr{F}$ as the sheaf that assigns to any open set $V \in \mathscr{O}p(Y)$ the set

$$f_\star\mathscr{F}(V) := \mathscr{F}(f^{-1}V).$$

(7.5.2.2) For any sheaf $\mathscr{G}$ on Y, we define the *inverse image* $f^\star\mathscr{G}$ as the sheaf of sections of the pullback $X \times_Y \acute{\mathrm{E}}\mathrm{t}(\mathscr{G}) \longrightarrow X$ of the map $p_{\mathscr{G}} : \acute{\mathrm{E}}\mathrm{t}(\mathscr{G}) \longrightarrow Y$

Example 7.5.3. — Suppose $A \subset X$ a subspace of a topological space X. Then for any sheaf $\mathscr{F}$ on X, if i denotes the inclusion map, the sheaf $i^\star\mathscr{F}$ on A is denoted $\mathscr{F}|_A$ and is called the *restriction* of $\mathscr{F}$ to A. If, in particular, A is an open set, then the restriction $\mathscr{F}|_A$ assigns to any open set $U \subset A$ the set $\mathscr{F}(U)$. On the other hand, if $A = \{x\}$, then $\mathscr{F}|_{\{x\}}$ is simply the stalk of $\mathscr{F}$ at x, regarded as a sheaf on $\{x\}$.

Exercise 61. — Any sheaf on the one-point space $\star$ is uniquely determined by a set S. Show that the direct image of such a sheaf along the inclusion $\star \longrightarrow X$ of a point x of a space X is the skyscraper sheaf S_x. Show that the inverse image of such a sheaf under the unique map $X \longrightarrow \star$ is the constant sheaf $a\mathscr{F}_S$.

ACT II

AN INTRODUCTION TO HOMOTOPY THEORY

CHAPTER 8

COVERING SPACE THEORY

8.1. Connectedness and π_0

Warning 8.1.1. — The definitions we give here do not coincide with definitions in many other texts. The definitions agree in "good cases," i.e., in all the cases with which we will have to deal in the sequel; however, there are important ways in which our definitions differ from those found elsewhere in the literature.

8.1.2. — For any set S, consider the constant *presheaf* $\mathscr{P}_S$ on a space X; its associated sheaf is the constant sheaf $\mathscr{F}_S$. The unit morphism $\mathscr{P}_S \longrightarrow \mathscr{F}_S$ induces a canonical map $S \longrightarrow \mathscr{F}_S(X)$ on global sections. If S is the one-point set, then this map is an isomorphism. If S is empty, then this map is a bijection unless X is empty.

An alternate description of the sheaf $a\mathscr{F}_S$ is the following. It is the sheaf $\mathscr{O}_X^{S^\delta}$ of local continuous functions valued in the discrete space S^δ.

Definition 8.1.3. — A space X is said to be *connected* if for any set S, the obvious map $S \longrightarrow \mathscr{F}_S(X)$ (which assigns to any element $s \in S$ the global section $x \longmapsto (x, s)$) is a bijection.

8.1.4. — Thus a space X is connected if and only if it is nonempty and every global section of the constant sheaf at S is the inclusion $X \times \{s\}^\delta \subset X \times S^\delta$ for some $s \in S$. Equivalently, X is connected if and only if the only continuous maps $X \longrightarrow S^\delta$ are constant.

Definition 8.1.5. — Suppose X a space. Then $\pi_0^{\text{ét}} X$ is a set along with a global section $u \in \mathscr{F}_{\pi_0^{\text{ét}} X}(X)$ such that for any set S and any global section $\sigma \in \mathscr{F}_S(X)$, there exists a unique morphism of sheaves $\widetilde{\sigma} : a\mathscr{F}_{\pi_0^{\text{ét}} X} \longrightarrow a\mathscr{F}_S$ such that $\widetilde{\sigma}(u) = \sigma$.

8.1.6. — Note that $\pi_0^{\text{ét}} X$ is a set with a continuous map $X \longrightarrow (\pi_0^{\text{ét}} X)^\delta$ such that for any set S and any continuous map $X \longrightarrow S^\delta$, there exists a unique map $\pi_0^{\text{ét}} X \longrightarrow S$ such that the diagram

$$\begin{tikzcd} X \arrow[rr] \arrow[dr] & & S^\delta. \\ & (\pi_0^{\text{ét}} X)^\delta \arrow[ur] & \end{tikzcd}$$

8.1.7. — Note that a space X is connected if and only if $\pi_0^{\text{ét}} X = \star$.

Warning 8.1.8. — Note that this definition of π_0 is a little strange. In particular, the definition is nonconstructive in the sense that it is not clear from the definition alone that $\pi_0^{\text{ét}} X$ exists for a given space X. In fact, it may very well not exist: in the case of $\mathbf{Q}$ with the subspace topology, for instance, there is no $\pi_0^{\text{ét}}$.

However, it's much easier to compute this version of π_0 than it is to compute other variants, at least once you get used to checking universal properties. This is primarily because it involves finding maps *out* of your space.

Exercise 62. — Recall that for any collection of spaces $\{V_\alpha\}_{\alpha\in\Lambda}$, the *topological disjoint union* $\coprod_{\alpha\in\Lambda} V_\alpha$ is the set-theoretic disjoint union $\coprod_{\alpha\in\Lambda} V_\alpha$ equipped with the following topology: an open set $U \subset \coprod_{\alpha\in\Lambda} V_\alpha$ is open if and only if each intersection $U \cap V_\alpha$ is open in V_α.

Now for any space V, a *decomposition* of V is a collection of subspaces $V_\alpha \subset V$ — called *summands* — such that V can be written as $\coprod_{\alpha\in\Lambda} V_\alpha$. Show that there is a bijective correspondence between decompositions $V = \coprod_{\alpha\in\Lambda} V_\alpha$ and continuous maps $V \longrightarrow \Lambda^\delta$. Let us call a decomposition $V = \coprod_{\alpha\in\Lambda} V_\alpha$ *nondegenerate* if and only if the corresponding map $V \longrightarrow \Lambda^\delta$ is surjective, or, equivalently, if no summand V_α is empty.

Show that a space X is connected if and only if it is nonempty and admits no nondegenerate decomposition. Show that if X is a space such that $\pi_0^{\text{ét}}X$ exists, then the continuous map $X \longrightarrow \pi_0^{\text{ét}}X$ is surjective. Conclude that $\pi_0^{\text{ét}}X$ exists if and only if there exists a *maximal* nondegenerate decomposition of X, i.e., a decomposition of X such that each summand is connected.

8.1.9. — Let us consider $\mathbf{Q}$ with the subspace topology. The only connected subspaces of $\mathbf{Q}$ are points. (Proof: if $A \subset \mathbf{Q}$ is a connected subspace containing two distinct points a and b, choose an irrational number $a < c < b$. Now consider the map $g : A \longrightarrow \{0,1\}^\delta$ given by the rule

$$g(x) := \begin{cases} 0 & \text{if } x < c; \\ 1 & \text{else.} \end{cases}$$

This is continuous, since $g^{-1}\{0\} = (-\infty, c) \cap A$ and $g^{-1}\{1\} = (c, \infty) \cap A$.) But it is not the case that $\mathbf{Q}$ is the disjoint union of its points, for then points would be open in $\mathbf{Q}$. Hence there is no maximal nondegenerate decomposition of $\mathbf{Q}$, so $\pi_0^{\text{ét}}\mathbf{Q}$ does not exist.

Exercise 63. — Compute $\pi_0^{\text{ét}}$ of the following spaces. (In all these cases, $\pi_0^{\text{ét}}$ exists, so there are no tricks here.)

(63.1) a set S with the discrete topology,
(63.2) the sphere S^n for any $n > 0$,
(63.3) the set $\mathbf{GL}_n\,\mathbf{R}$ of $n \times n$ invertible matrices with real entries, with the subspace topology from the inclusion $\mathbf{GL}_n\,\mathbf{R} \subset \mathbf{R}^{n^2}$, for $n > 0$,
(63.4) the set $\mathbf{SL}_n\,\mathbf{R}$ of $n \times n$ matrices with real entries of determinant 1, with the subspace topology from the inclusion $\mathbf{SL}_n\,\mathbf{R} \subset \mathbf{R}^{n^2}$, for $n > 0$, and
(63.5) the set $\mathbf{GL}_n\,\mathbf{C}$ of $n \times n$ invertible matrices with complex entries, with the subspace topology from the inclusion $\mathbf{GL}_n\,\mathbf{C} \subset \mathbf{C}^{n^2}$, for $n > 0$.

Exercise 64. — Consider the letters of the (English) alphabet A–Z as subspaces of the plane $\mathbf{R}^2$. (Consider them in their simplest possible form, without any thickness of any lines, and without any serifs!) Can you classify all of them up to homeomorphism? (Hint: if two spaces X and Y are homeomorphic, then their $\pi_0^{\text{ét}}$ agree, and if $f : X \longrightarrow Y$ is a homeomorphism, then so is $f : X - \{x\} \longrightarrow Y - \{f(x)\}$ for any $x \in X$.)

8.1.10. — The construction $X \longmapsto \pi_0^{\text{ét}}X$ is *functorial* in the sense that for any continuous map $f : X \longrightarrow Y$, there is a corresponding map $\pi_0^{\text{ét}}f : \pi_0^{\text{ét}}X \longrightarrow \pi_0^{\text{ét}}Y$ if both $\pi_0^{\text{ét}}X$ and $\pi_0^{\text{ét}}Y$ exist. Indeed, the definition implies that there exists a unique map $\pi_0^{\text{ét}}f : \pi_0^{\text{ét}}X \longrightarrow \pi_0^{\text{ét}}Y$ such that the diagram

$$\begin{tikzcd} X \arrow[r, "f"] \arrow[d] & Y \arrow[d] \\ (\pi_0^{\text{ét}}X)^\delta \arrow[r, "\pi_0^{\text{ét}}f"'] & (\pi_0^{\text{ét}}Y)^\delta \end{tikzcd}$$

commutes.

8.2. Covering spaces and locally constant sheaves

Definition 8.2.1. — A *covering space* over a topological space B (called the *base*) is a continuous map $p : E \longrightarrow B$ such that every point $b \in B$ is contained in an open neighborhood U for which there is a homeomorphism

$$\gamma_U : p^{-1}U \longrightarrow U \times (p^{-1}b)^{\delta}$$

such that the diagam

$$\begin{tikzcd} p^{-1}U \arrow[rr, "\gamma_U"] \arrow[dr, "p|_{p^{-1}U}"'] & & U \times (p^{-1}b)^{\delta} \arrow[dl, "\mathrm{pr}_1"] \\ & U & \end{tikzcd}$$

commutes.

Exercise 65. — Show that a covering space is a local homeomorphism.

Exercise 66. — Check that the pullback $E \times_B B' \longrightarrow B'$ of a covering space $E \longrightarrow B$ along any continuous map $B' \longrightarrow B$ is a covering space.

Definition 8.2.2. — A *morphism of covering spaces* from $p' : E' \longrightarrow B$ to $p : E \longrightarrow B$ (covering spaces over the same base) is simply a continuous map $f : E' \longrightarrow E$ such that the diagram

$$\begin{tikzcd} E' \arrow[rr, "f"] \arrow[dr, "p'"'] & & E \arrow[dl, "p"] \\ & B & \end{tikzcd}$$

commutes.

Example 8.2.3. — (8.2.3.1) For any space X and any set S, the map $\nabla^{(S)} : X^{\sqcup(S)} \cong X \times S^{\delta} \longrightarrow X$ is a covering space. These covering spaces are known as the *trivial covering spaces*. The corresponding locally constant sheaves are constant.

(8.2.3.2) Consider again the continuous map $p_n : S^1 \longrightarrow S^1$ given by $z \longmapsto z^n$. This is a covering space too. The corresponding sheaf assigns to any open set $U \subset S^1$ the set of n-th root functions defined on U.

(8.2.3.3) The exponential map $\exp : \mathbf{C} \longrightarrow \mathbf{C}^{\times} := \mathbf{C} - \{0\}$ is a covering map. One can pull this example back to $S^1 \subset \mathbf{C}^{\times}$ to get the covering space $\mathbf{R} \longrightarrow S^1$ given by $\theta \longmapsto e^{2\pi\theta\sqrt{-1}}$. The corresponding sheaf assigns to any open set U a logarithm defined on U.

Exercise 67. — Consider the subspace $(\mathbf{C} - \{0\})^{\times n} \subset \mathbf{C}^n$. Show that for any n-tuple $(k_1, \dots, k_n)$ of nonnegative integers, the map

$$\begin{aligned} p_{(k_1,\dots,k_n)} : (\mathbf{C} - \{0\})^{\times n} &\longrightarrow (\mathbf{C} - \{0\})^{\times n} \\ (z_1, \dots, z_n) &\longmapsto (z_1^{k_1}, \dots, z_n^{k_n}) \end{aligned}$$

is a covering space. Pull this back to get a covering space on the n-dimensional torus $\mathbf{T}^n$. Now contemplate the fiber over $(1, \dots, 1)$. Can you draw a picture of this when $n = 2$?

Definition 8.2.4. — A sheaf $\mathscr{F}$ on a space X is *locally constant* if every point of X is contained in an open neighborhood U such that the restriction $\mathscr{F}|_U$ is a constant sheaf.

Theorem 8.2.5. — *A local homeomorphism $p : E \longrightarrow B$ is a covering map if and only if the sheaf $\Gamma(E/B)$ is locally constant; correspondingly, a sheaf $\mathscr{F}$ is locally constant if and only if the espace étalé* Ét$(\mathscr{F})$ *is a covering space.*

8.2.6. — Under our bijective correspondence between homeomorphism classes of local homeomorphisms and isomorphism classes of sheaves, we see that locally constant sheaves correspond (bijectively) to covering spaces. As we have seen, constant sheaves correspond to trivial covering spaces.

To put this more formally, for any space B, let $(\mathscr{LC}/B)$ be the category of locally constant sheaves on B, and let $\mathscr{C}ov(B)$ denote the category of covering spaces with base B. Then our equivalence of categories $(\mathscr{LH}/B) \simeq \mathscr{S}h(B)$ restricts to an equivalence of categories $(\mathscr{LC}/B) \simeq \mathscr{C}ov(B)$.

***Exercise*** **68**.* — Show that the only possible locally constant sheaves on the interval $[0,1]$ are constant. (Hint: Suppose $\mathscr{F}$ a locally constant sheaf on $[0,1]$. Show that there is a finite open cover of $[0,1]$ comprised of intervals U on which $\mathscr{F}|_U$. Now induct on the size of $\mathscr{U}$ in order to show that $\mathscr{F}$ must be constant. Why doesn't this argument work for S^1?)

Using this, suppose X any space, and suppose γ a *path* in X — i.e., a continuous map $\gamma : [0,1] \longrightarrow X$. For any locally constant sheaf $\mathscr{G}$ on X, the pullback $\gamma^\star\mathscr{G}$ is a constant sheaf. More precisely, for each $t \in [0,1]$, show that there is a unique isomorphism between $\gamma^\star\mathscr{G}$ and the constant sheaf $\mathscr{F}_{S(t)}$ at the stalk $S(t) := \mathscr{G}_{\gamma(t)}$. Now if γ is a *loop* in X — i.e., a path such that $\gamma(0) = \gamma(1) = x \in X$ —, we have specified two isomorphisms:

$$\mathscr{F}_x = \mathscr{F}_{S(0)} \longrightarrow \gamma^\star\mathscr{G} \longleftarrow \mathscr{F}_{S(1)} = \mathscr{F}_x,$$

so we have an automorphism of the stalk $\mathscr{F}_x$. Give an example to show that this autmorphism need not be the identity.

***Definition 8.2.7**.* — We say that a covering space $p : E \longrightarrow X$ is *finite* if the fibers of p are finite. If $\pi_0^{\text{ét}}X$ exists, then we say that it is *of finite type* if $\pi_0^{\text{ét}}E$ exists, and the fibers of $\pi_0^{\text{ét}}p : \pi_0^{\text{ét}}E \longrightarrow \pi_0^{\text{ét}}X$ are finite. Let us denote by $(\mathscr{C}ov^{\text{ft}}/X)$ the set of all isomorphism classes of covering spaces of finite type over X.

8.3. The étale fundamental group

8.3.1. — We are now interested in classifying all locally constant sheaves/covering spaces over a connected space X. In order to do this, we will begin by constructing, for any point $x \in X$, the *étale fundamental group* $\pi_1^{\text{ét}}(X,x)$ out of the collection of all covering spaces of finite type over X.

***Definition 8.3.2**.* — Suppose X a space, and suppose $x \in X$. For any covering space $E \longrightarrow X$, a *monodromy element for E at x* is an automorphism of the fiber E_x.

A *uniform monodromy element at x* is an element

$$\gamma = (\gamma_E) \in \prod_{E \in (\mathscr{C}ov^{\text{ft}}/X)} \operatorname{Aut}(E_x)$$

such that for any morphism $\phi : E' \longrightarrow E$ of covering spaces of finite type,

$$\begin{tikzcd} E'_x \arrow[r, "\phi_x"] \arrow[d, "\gamma_{E'}"'] & E_x \arrow[d, "\gamma_E"] \\ E'_x \arrow[r, "\phi_x"'] & E_x. \end{tikzcd}$$

For any two uniform monodromy elements γ and γ', one may compose them to obtain

$$\gamma' \circ \gamma := (\gamma'_E \circ \gamma_E) \in \prod_{E \in (\mathscr{C}ov^{\text{ft}}/X)} \operatorname{Aut}(\mathscr{F}_x)$$

The *étale fundamental group* $\pi_1^{\text{ét}}(X,x)$ is the set of uniform monodromy elements.

***Lemma 8.3.3**.* — *The set $\pi_1^{\text{ét}}(X,x)$ is a group under composition.*

Proof. — Since $\pi_1^{\text{ét}}(X,x)$ is given as a subset of a group, we only have to check that $\pi_1^{\text{ét}}(X,x)$ contains the identity (obvious) and that it is closed under composition. This latter point follows from inspection of the diagram

$$\begin{tikzcd} E'_x \arrow[r, "\phi_x"] \arrow[d, "\gamma_{E'}"'] & E_x \arrow[d, "\gamma_E"] \\ E'_x \arrow[r, "\phi_x"'] \arrow[d, "\gamma'_{E'}"'] & E_x \arrow[d, "\gamma'_E"] \\ E'_x \arrow[r, "\phi_x"'] & E_x. \end{tikzcd}$$

□

8.4. Homotopy classes of loops and lifts

Lemma 8.4.1. — *Consider the following commutative diagram of spaces*

$$\begin{tikzcd} \star \arrow[r] \arrow[d] & E \arrow[d, "p"] \\ Y \arrow[r, "f"'] & B \end{tikzcd}$$

in which $p : E \longrightarrow B$ is a covering space and Y is connected space. Then a lift of f, if it exists, is unique.

Proof. — Suppose g and h two lifts. Define a map $r : Y \longrightarrow \{0,1\}^{\delta}$ by

$$r(y) = \begin{cases} 1 & \text{if } g(y) = h(y) \\ 0 & \text{if not.} \end{cases}$$

I claim r is continuous. Indeed, suppose $y \in Y$; let U be an open neighborhood of $f(y)$ such that $p^{-1}U \cong U \times E_{f(y)}$. Let $e_g, e_h \in E_{f(y)}$ be the projections of the element $g(y), h(y)$ to $E_{f(y)}$; then the set

$$V := g^{-1}(U \times \{e_g\}) \cap h^{-1}(U \times \{e_h\})$$

is open in Y, and one verifies that $V \subset r^{-1}(f(y))$. This shows that r is continuous, and the result follows from the connectedness of Y. □

Lemma 8.4.2. — *Unique path lifting ...*

Lemma 8.4.3. — *In the situation above, assume that Y is locally path connected. Then a lift exists if and only if the image of $\pi_1(Y)$ in $\pi_1(X)$ is contained in the image of $\pi_1(E)$.*

Proof. — Necessity is obvious. For sufficiency, define a lift $\phi : Y \longrightarrow E$ in the following manner. For any point $y \in Y$, choose a path γ from the basepoint to y. Then there is a unique lift γ' of this path to E; define $\phi(y) = \gamma'(1)$. The assumption guarantees that this is independent of the choice of path. □

8.5. Universal covering spaces

Definition 8.5.1. — Suppose X a space, and suppose $x \in X$. Then a *universal locally constant sheaf* on (X,x) is a locally constant sheaf $\mathscr{E}$ on X along with a germ of a section $\upsilon \in \mathscr{F}_x$ such that for any locally constant sheaf $\mathscr{F}$ and any germ $\sigma \in \mathscr{F}_x$, there is a unique morphism $\phi : \mathscr{E} \longrightarrow \mathscr{F}$ of sheaves such that $\phi_x(\upsilon) = \sigma$.

Exercise 69. — Suppose X a space, and suppose $x \in X$. Show that the universal locally constant sheaf $(\mathscr{E}, \upsilon)$ on (X,x), if it exists, has the following property: for any locally constant sheaf $\mathscr{F}$ on X, the map $\mathrm{Mor}_X(\mathscr{E}, \mathscr{F}) \longrightarrow \mathscr{F}_x$, $\phi \longmapsto \phi_x(\upsilon)$, is a bijection. Deduce that if a universal locally constant sheaf on (X,x) exists, then it is unique up to a unique isomorphism.

8.5.2. — There is a natural candidate for the universal locally constant sheaf on a pointed space (X,x).

Definition 8.5.3. — Suppose X a space, and suppose $x \in X$. Then a *universal covering space* over (X,x) is a covering space $u : \widetilde{X} \longrightarrow X$ along with a point $\widetilde{x} \in u^{-1}(x)$ such that for any covering space $p : E \longrightarrow X$ and any point $e \in p^{-1}(x)$, there exists a unique covering space map $f : \widetilde{X} \longrightarrow E$ with the property that $f(\widetilde{x}) = e$.

Definition 8.5.4. — A space X is said to be *well connected* if it is path connected, it admits a basis comprised of path-connected neighborhoods, and any point x is contained in a neighborhood U such that the homomorphism $\pi_1(U,x) \longrightarrow \pi_1(X,x)$ is trivial.

Theorem 8.5.5. — *Suppose X a well connected space. Then X admits a universal covering space. Moreover, it is the unique connected pointed covering space $(\widetilde{X},\widetilde{x})$ of X such that $\pi_1(\widetilde{X},\widetilde{x})$ is trivial.*

Proof. — Let us construct the corresponding locally constant sheaf. Begin by contemplating the based path space $\Pi_x(X) \subset \mathscr{C}(I,X)$, whose points are those paths γ such that $\gamma(0) = x$, equipped with the subspace topology. Evaluation at 1 gives a continuous map $e : \Pi_x(X) \longrightarrow X$ Consider now the sheaf of sections $\Gamma(e)$; for any open set U of X, define $\mathscr{G}(U)$ as the quotient $\Gamma(e)(U)/\sim$, where we say that two sections $s,t \in \Gamma(e)(U)$ are equivalent if they are homotopic as maps $U \times I \longrightarrow X$ rel $U \times \{0,1\}$. The result is a presheaf $\mathscr{G}$; the claim is that its sheafification is locally constant. To this end, observe that since X is well connected, it admits a base consisting of path connected open sets U such that the induced homomorphism $\pi_1(U) \longrightarrow \pi_1(X)$ is trivial. Over these open sets, the espace étalé of $\mathscr{G}$ is a trivial covering space.

So now we have the corresponding covering space $\widetilde{X}$. It is the space of homotopy (rel $\{0,1\}$) classes of paths starting at x; evaluation at 1 gives the structure map $\widetilde{X} \longrightarrow X$. Write $\widetilde{x}$ for the constant map at x. If now we can check that $\widetilde{X}$ is path connected and that $\pi_1(\widetilde{X},\widetilde{x})$ is trivial, then we are done by 8.4.3.

To show that $\widetilde{X}$ is path connected, consider an equivalence class α of paths. Let f be a representative thereof. Now for any real number t, let $f_t(s) = f(st)$. This gives a path from $\widetilde{x}$ to α.

To show that $\pi_1(\widetilde{X},\widetilde{x})$ is trivial, ... □

Theorem 8.5.6. — *Suppose X admits a universal covering space. Then there is a bijective correspondence between $(\mathscr{C}ov^{\mathrm{ft}}/X)$ and the set of isomorphism classes of $\pi_1^{\text{ét}}(X,x)$-sets with finitely many orbits, and under this correspondence, the connected covering spaces correspond to the transitive $\pi_1^{\text{ét}}(X,x)$-sets.*

CHAPTER 9

COMPUTING THE ÉTALE FUNDAMENTAL GROUP

CHAPTER 10

HOMOTOPY THEORY

CHAPTER 11

THE COMPARISON THEOREM

11.1. The two fundamental groups