4 Topology: A First Pass at Space

In which we cover all the elements of general (point-set) topology needed for sheaves on a topological space, introduce the relevant categories of interest, and raise some questions (properly addressed in the appendix) regarding some of the fundamental ingredients of topology.

The classical presentation of sheaves begins by looking at sheaves on topological spaces. As we will explore in the final chapters of the book, there are powerful elaborations of sheaves to more general settings than topological spaces. However, the right place to start to introduce the sheaf notion is with sheaves on topological spaces. While a few important topics of basic category theory remain to be covered, we already have the main category-theoretic ingredients for defining sheaves in the context of topological spaces. In order to take our first look at sheaves, then, all that remains is to cover the necessary elements of general topology. That is the task of the present chapter. In the next chapter, we take a first look at sheaves. Chapters 6 and 7 then return to two last important topics and ingredients of basic category theory, matters worth exploring in their own right but also ones that will allow us to significantly enrich the treatment of sheaves that will emerge in later chapters.

4.1 Motivation

The standard story told, when introducing point-set topology, presents the concepts of topology as abstractions of features of metric spaces, specifically those nearest and dearest to the heart of analysts: the Euclidean space $\mathbb{R}^k$, especially $\mathbb{R}$, the real line, and $\mathbb{R}^2$, the plane. It is true that point-set topology has its roots in analysis, and while today we develop these notions on their own terms, historically the staples of topology developed in order to meet the needs of analysis, not as part of a separate subject. A metric space is basically just a set of points and a relation on those points that acts as a certain quantitative measure of the degree of closeness or nearness of pairs of points. Equipped with a notion of how far any two points are, the decisive concerns of the analyst—such as continuity and approximation, among others—can be easily characterized. But as far as continuity and approximation (together with a number of other key notions of analysis) are concerned, one is never really concerned with points per se, but rather with regions that include everything nearby a point of interest. For now, think of these as *regions of approximation*.

In the general algebra of sets, we work with the set-theoretic operations of union, intersection, and complementation—the fundamental operations for relating and combining existing sets, allowing us to generate new ones from old. By attending to the behavior of the regions of approximation in relation to one another, as we take their union and intersection, it can be observed that these objects—whatever their shape or whatever notion of distance we use—always behave in a characteristic way: they are stable under intersection and union, but in an asymmetrical fashion as concerns the finiteness of these operations.

The passage from metric spaces to topological spaces in general was largely aided by the realization that if we get rid of the distance function that helps define a metric space and gives us our very sense of nearness but retain the properties concerning how the pertinent objects are stable under intersection and union (in an asymmetrical way with respect to finiteness), not only can we recover the same metric notions, but we are left with a new and more general notion that can capture a wide variety of further structures of interest.

Given a set, we can observe certain things about it, as far as its constituent parts are concerned—for instance, we can look at its cardinality: How *many* elements are in it? But we might care less about something like the number of parts and more about the relations between the parts. In asking questions about the interrelations between the parts of a set, we would appear to need an answer to the questions:

- What kind of parts do we allow?
- What sorts of relations between those parts are allowed?

As it turns out, in the more general setting, questions about what kind of objects we have (e.g., how the regions of approximation are characterized) can be reduced to specifying how they relate to one another. A topology fundamentally consists of a collection of subsets of a set X, a certain structure endowing the constituent objects with some coherence, meaning that it makes sense to determine when things are nearby or close together; this can be articulated entirely in terms of how these subsets satisfy certain conditions, specifically conditions on how the subsets relate to one another.

The standard story told to motivate general topological notions and arrive at the key axioms of a topology proceeds by

1. observing certain features of approximating objects native to Euclidean space; then
2. abstracting from Euclidean space to metric spaces in general; then
3. abstracting again from metric spaces to certain properties of general *open sets*, using these as the axioms determining a space in general.

The standard objects or parts we work with in topology are *open* and *closed* sets. The canonical example of an open set, from standard metric space settings, is an open interval or open disk—where this subset characteristically contains none of its boundary. The canonical example of a closed set is a closed interval or closed disk, where this by contrast contains all of its boundary. In general, when we take the intersection of two sets, say A and B, denoted $A \cap B$, we end up with the set that contains all elements of A that also belong to B (or, equivalently, all elements of B that also belong to A). Suppose we are working with some closed intervals. Clearly, the intersection of any closed intervals will itself be a closed set, and this can be extended to arbitrarily many such closed sets. By contrast, if we start with open intervals, notice how the intersection of any two open intervals will itself be an open interval—yet this *cannot* be extended to an arbitrary number of open intervals. After all, the intersection of an infinite number of shrinking open intervals centered on a point will just be the point itself—yet points are boundary-like objects and cannot be considered open without things becoming rather degenerate. In a similar way, we can consider what happens when we take unions of sets. Suppose we have open intervals. The union of any number of open intervals will itself be an open interval. But if we start instead with closed intervals, it is not the case that the union of an arbitrary number of closed intervals will be closed—so, to ensure that we are left with another closed set after we deploy the operation of union on closed sets, we must restrict the union to the finite case.

At bottom, these two dual conditions are what define a topology:

- for open sets: stability under finite intersection and arbitrary union;
- for closed sets: stability under arbitrary intersection and finite union.

Open and closed sets are related to one another in a particularly nice way, so in principle it does not really matter which sorts of object we use. Rather, the essence of the topology notion is the characteristic asymmetrical finiteness conditions on how subsets behave when unioned and intersected. These conditions on the interrelations of subsets are what ultimately constitute openness versus closedness, rather than any intrinsic property of some "open" or "closed" entity in itself.

There are aspects of the standard story—where topological spaces are presented as a natural generalization of conspicuous properties of especially familiar metric spaces—that can appear compelling to newcomer and seasoned mathematician alike. However, the confidence with which this story is typically told tends to leave unanswered—or, worse, even obfuscate—a number of legitimate and lingering questions. These important (if somewhat more philosophical) questions are raised at the end of this chapter, but treatment of them is relegated to the appendix, in order not to distract from the march onward toward sheaves. The appendix is perhaps best read after chapter 7 on adjunctions—or, if the reader is already familiar with such matters and is willing to postpone arriving at sheaves, after the present chapter.

In addition to being vital to the first presentation of the sheaf concept, general topology is one of the fundamental branches of modern mathematics, next only to set theory. The concepts one meets in point-set topology provide us with a framework for expressing ideas that extend to nearly all branches of mathematics. For these reasons, this chapter devotes considerable time to developing the core notions of general topology (henceforth, in this chapter, we'll just say "topology"), and refers to the appendix for a firmer understanding of the "essence" of topology by exploring potential answers to the three questions raised at the end of this chapter.

The rest of this chapter is structured as follows. Section 4.2 first motivates and describes, in a more informal fashion, via a dialogue, a number of the decisive concepts explored and dealt with in general topology. Section 4.3 then explores these in a more formal and detailed fashion, providing a self-contained account rooted in examples and worked-out exercises, of all that is needed from topology for the purposes of sheaf theory. The Philosophical Pass at the end of the chapter (section 4.4) raises three lingering questions about topological matters, questions whose treatment is relegated to the appendix.

4.2 A Dialogue Introducing the Key Notions of Topology

Suppose you have in front of you a sheet of paper on which a blot of ink has been dropped:

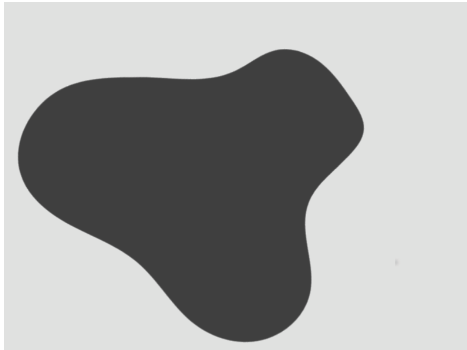

A bright but mischievous young student points to somewhere well within the ink blob, announcing that the paper is black there. You agree. The student asks you,

"But how do you know that the point I pointed to is black?"

You take them to be asking,

"How do you know that the point (call it *b*) is within the region of the ink blob (call it *R*)?"

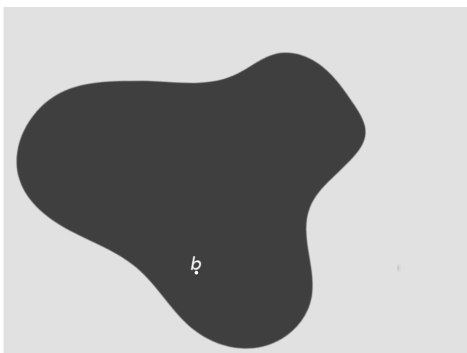

How will you answer them?

Perhaps you are feeling rather lazy at first, and you are not "on duty" at the moment, and so you attempt to meet their queries with some reference to how one can "just see" that the point in question has to be black, adding some tautological remark about the color black while gesturing to the point. But the student is not deterred—sensing that some evasion is happening, they grow even more eager to address the issue head on:

STUDENT: You said that I can observe that *b* is black (in *R*). But *what* exactly am I being asked to observe? I have no problem agreeing that, whatever the black region *R* is, *R* is itself black—this is trivial. But I wanted to know about *b*.

MATHEMATICIAN: Yes, well *b* is in *R*, as I said.

STUDENT: Yes, and as I said: *How* do you know that? You cannot say that I can observe *b* itself—I cannot. It seems to me that whatever is given to observation will be extended. Plus, in my geometry classes they say that points have no size, no extension. So even though we use dots to represent points, isn't this just an idealization, just like the mythical "instant"—there is no such extended point, so there can be no features *at the point*; in particular, I can observe no color there.

MATHEMATICIAN: Interesting. And you seem to be suggesting that this would not just be a limitation on your part: If it is assumed that points are not extended, but observables are extended (in space or time or both, say), then you would be correct in asserting that points would (at the very least) not be the primary bearers of spatial properties or relations—no property, such as color, could be observed of, or primarily ascribed to, a point itself. No matter how fine-grained an observation could get—no matter how much you could "zoom in"—it could not zoom in *exactly onto the point* itself and directly observe a property of it.

STUDENT: Yes, that is basically what I was thinking.

MATHEMATICIAN: In that case, it seems like we have two options: (1) give up on saying anything about points at all or (2) declare that whatever holds at a point really pertains to what holds of some extended region or some non-point-like part of a space *surrounding* the point, so of the *nearby* points (where perhaps we can then use what is nearby to say something about our point).

With the second option, we might even initially relinquish the idea that there are zero-sized points or instants—instead, we might take ourselves to be working with nonzero or nondimensionless regions, while allowing that such regions could be made *arbitrarily small*. The wager here is that, even though we are apparently relinquishing the idea of points, we may be able to recover everything we might want to say about points via such regions.

STUDENT: The first option seems silly. After all, we say things about points all the time, and it seems unnecessarily pessimistic to prohibit ourselves from making claims about points. And I was being sneaky before. I *do* see that some points of the paper are black and others gray (the color of the paper). And I see that where I pointed earlier, at b in particular, is clearly within R—I just want to know how to articulate in a less suspect or vague way this word "clearly."

MATHEMATICIAN: Indeed. And the second option seems sensible in its own right: after all, we have been asking about how b relates to its environs.

This also gets at an important feature of the overall problem, one that relates to your observations about points. In settings where mathematicians first clearly formulated the notions needed to address our problem, along the lines of route 2, a primary concern was with *continuity* and related matters. Strictly speaking, the sorts of things that can be continuous—motions, velocities, and so on—cannot be seen as an *intrinsic* property of some object like a point or instant. Motion, for instance, is something that involves a relationship between the state of an object at one instant and the state of that object at another instant or collection of instants. In concerning ourselves with such things, and whether or not they are *continuous*, we are concerning ourselves with something that ultimately involves a relationship between *multiple* moments, instants, or points.

Before we get too carried away with any of this, though, I bet we can start to get a better handle on your earlier questions by returning to this idea of "zooming in," since I suspect that we will find that, even if we cannot directly observe a property exactly at a point, there is something we can use in this notion of zooming.

STUDENT: So we are thinking about having something like lenses or magnifying glasses, each one of which gives us a distinct magnification?

MATHEMATICIAN: Yes, that works. Each lens's distinct magnification could then just be uniquely described by the radius of the resulting window of observation that we would get, upon looking through the lens. In that way, we can give each lens with a distinct magnification its own name, depending on the radius of its window of observation.

STUDENT: Okay, so for a given point, we can center a lens on that point and take a look at it with that lens. For instance, for our point *b*, a medium-sized lens might show us

while a lens that can zoom in further might instead show us

MATHEMATICIAN: That is the right idea. And what do these lenses tell you?

STUDENT: Well, so far, I'm not really sure they tell me anything. Not anything definitive, at least. I would like to say that the second one is somehow "better" though.

MATHEMATICIAN: If it is better, as you say, you could try zooming in even further.

STUDENT: That is exactly what I was thinking. Suppose we end up with a small enough region of observation such that all around us we see only black,

MATHEMATICIAN: It seems that, with this lens, we can now declare without hesitation that *b* itself is black, that is, inside the ink blob, no?

STUDENT: I agree. And I don't have to zoom in any further, right?

MATHEMATICIAN: Let's give this a name.

Paradigm/Principle of Truth Continuity: an observable property will be true of a point whenever it continues to be true in all better approximations.

STUDENT: And this would have worked just as well on, say, the point *a*, where this is decidedly in the gray paper region. It seems that whenever, around a point *a*, there is a small enough region of observation such that all around us we see only gray, we can likewise declare without hesitation that *a* itself must be gray, that is, in the paper region.
It seems like I can use these lenses to tell when any point is black or not black.

MATHEMATICIAN: Hold on. Consider

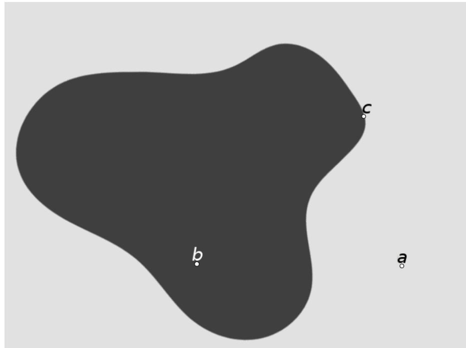

For a and b the concepts we are building would indeed suffice to tell us unequivocally. Imagine zooming in further and further to b. Clearly we can zoom in close enough such that "we only see black." For a, similarly, we can zoom in to a small enough region around a such that "we only see gray." But what about c?

STUDENT: Well, as I zoom in further and further to c, *if we are assuming that the region R includes all its outer edges*, so that c is exactly *on* the edge of R, there is evidently no way of zooming in *close enough* that we are surrounded by only black. Likewise, there is no way of zooming in close enough that we are only surrounded by only gray. However small we make our zoomed-in region around c, we will see both black and not-black!

MATHEMATICIAN: Yes, exactly, and we can use such trouble cases to define a new notion: that of the *boundary*.

STUDENT: Doesn't that mean, by the way, that what we have been calling "lenses" are being thought of as *not* including their boundary? After all, if they had included their boundary, then there would be points within its observation window that we could zoom in on with more and more refined lenses and yet for which we could never say unequivocally whether or not they were black or not-black. It seems like this could further force us to give up on your Paradigm of Truth Continuity, which seemed sensible enough to me.

MATHEMATICIAN: That is exactly right. Incidentally, what we have been attempting to name with our "lenses"—or, really, the various windows of observation that result from using such lenses—could instead be thought of in terms of what mathematicians would call an *open disk* in a plane (or, if we had been restricting our attention to the real line, an *open interval*).

STUDENT: Yes, I know about these. In $\mathbb{R}$, we're thinking about intervals of the form $(a,b)=\{x\in\mathbb{R} \mid a<x<b\}$:

And in terms of the plane $\mathbb{R}^2$, we are thinking of objects like the following:

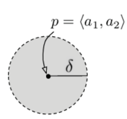

that is, the set of all the points within some fixed nonzero distance. For instance, with such disks, we have all the points inside a circle with center p and radius $\delta > 0$.

MATHEMATICIAN: That is right. Here, of course, we have been implicitly using a particular distance. In the plane, for instance, $d(p, q)$ denotes the usual distance between two points $p = \langle a_1, a_2 \rangle$ and $q = \langle b_1, b_2 \rangle$ in $\mathbb{R}^2$, that is,

$$d(p, q) = \sqrt{(a_1 - b_1)^2 + (a_2 - b_2)^2}$$

STUDENT: Okay, so it seems like we have been saying that, for a given region (say R), in examining some point x, if we can find such a disk D of some radius $\delta > 0$ around x such that the disk itself D is contained in the region R, then we should be able to confidently assert of the point x itself that it is "inside" R.

MATHEMATICIAN: Yes, you have effectively described here what mathematicians call an *interior point*: for a subset A of $\mathbb{R}^2$, a point $p \in A$ is an interior point of A precisely when p belongs to some open disk D_p that is contained in A, that is, $p \in D_p \subseteq A$.

STUDENT: And is there a name for the inverse of the relation "x in an interior point of A"?

MATHEMATICIAN: Yes, you were simultaneously describing the notion of a *neighborhood*: for a point $x \in \mathbb{R}^2$, a subset N of $\mathbb{R}^2$ is a neighborhood of x precisely when it contains an open disk that contains x, that is, $p \in D \subseteq N$.

STUDENT: Wait, that's all great, but getting back to R and what we were talking about a moment ago, how do these notions apply?

MATHEMATICIAN: Well, let's look closer at c. If c is assumed to be on the boundary of the region R, will we able able to form an open disk around it such that that open disk is entirely contained in R?

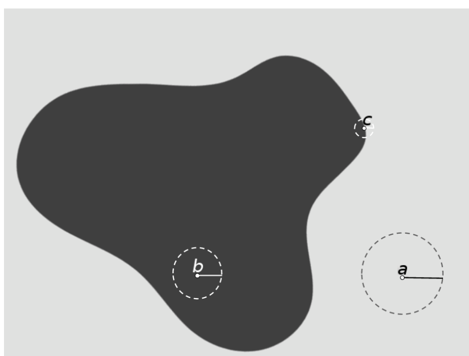

STUDENT: No. Any open disk around c would have to include something that was *not* R. So c cannot be an interior point of R.

MATHEMATICIAN: Exactly.

STUDENT: But b is surely an interior point—and a, for its part, may be an interior point of the gray paper region (whatever we want to call that).

MATHEMATICIAN: That is right. Suppose, then, that *each* of the points of R is an interior point. In this case, mathematicians would call R itself an *open set*.

STUDENT: It seems like this is all very wrapped up with the notion of boundary. For instance, to say that R is open—so that every point of R is an interior point—seems to just say that R does not include any of its boundary.

MATHEMATICIAN: Right! You have hit on another formulation of this notion of openness.

STUDENT: Okay, but when we move from the real line $\mathbb{R}$ to the plane $\mathbb{R}^2$, the same notions seem to carry over without a problem. I was wondering: What happens as we move to other spaces? And I have heard that there are still other notions of distance. What happens then? Also, is there something special about disks?

MATHEMATICIAN: Great observations! The construction we have been using works for any notion of nearness or distance, as long as we have specified what this is. The relevant notion of nearness is specified by the underlying "metric." So far, we have really just been working with a particularly familiar metric, and the familiar space over which this reigns, so we have not bothered to be explicit about this. But let's correct that now.

Definition 73 A *metric space* (X, d) consists of X a nonempty set, the elements of which are called "points," and a function $d : X \times X \to \mathbb{R}_{\geq 0}$ called a *metric* or *distance* that associates to any two points $x, y \in X$ a point $d(x, y)$ in such a way as to satisfy the following properties for all $x, y, z \in X$:

1. $0 \leq d(x, y)$ (or just $d(x, x) = 0$);
2. if $d(x, y) = 0$, then $x = y$;
3. $d(x, y) = d(y, x)$;
4. $d(x, y) + d(y, z) \geq d(x, z)$.

We can then use this notion of a metric to define a more general version of our open disks (and open intervals) from before.

Definition 74 For (X, d) a metric space, with $a \in X$ and $r > 0$, an *open ball around a of radius r* is the set

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

As you might have expected, these open balls can then be used to define the following:

Definition 75 For (X, d) a metric space, and $x \in X$, a subset $A \subseteq X$ is called a *neighborhood* of x if there exists an $\epsilon > 0$ such that

$$B_\epsilon(x) \subseteq A,$$

that is, provided an open ball around x can be contained in A.

Finally,

Definition 76 For (X, d) a metric space, subset $O \subseteq X$ is said to be *open* in (X, d) whenever for every $a \in O$, there exists an $\epsilon > 0$ such that

$$B_\epsilon(a) \subseteq O.$$

STUDENT: So what we have been calling lenses, and thinking of as "wiggle rooms," in the context of $\mathbb{R}^2$ with the usual sense of distance, are just particular instances of this general notion of *open balls*, which is like the distance-agnostic definition of the same thing?

MATHEMATICIAN: Precisely. Now observe how as we have been thinking about open sets thus far, in terms of in our real line $\mathbb{R}$ or plane, with the usual sense of distance between points, we have been thinking that

> A set is open if whenever it contains a number a, it also contains all numbers "sufficiently close" to a.

In other words, a subset A of $\mathbb{R}$ is open if for every $a \in A$, there exists a number $\epsilon > 0$ such that the open interval

$$(a - \epsilon, a + \epsilon)$$

is a subset of A, where of course the interval $(a - \epsilon, a + \epsilon)$ just consists of all numbers within (the usual) distance ϵ of a, that is,

$$\{x \in \mathbb{R} \mid |x - a| < \epsilon\}.$$

But the point is that we need not confine ourselves to the usual notion of what counts as close—as long as there is a notion of distance, we can use that and the same objects will be available to us. A set A is open if, for each point in the set, we can find a little "wiggle room"—defined in terms of the prescribed distance—around the point, without having to leave A.

STUDENT: And, on this more general account, we need not confine attention to discs or open sets of a certain shape?

MATHEMATICIAN: That's correct. They need not be (literal) balls or circles at all. Even in the plane, especially as we introduce different metrics, open sets will include a variety of things, like open discs $D = \{\langle x, y\rangle \mid (x - a_1)^2 + (y - a_2)^2 < \delta^2\} = \{q \in \mathbb{R}^2 \mid d(p, q) < \delta\}$, open half-planes $\{\langle x, y\rangle \mid y < a\}$, open rectangles, and so on.

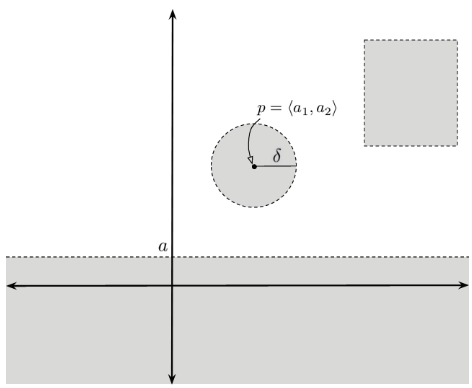

The only thing we require, for a set to be open, is that it be a neighborhood of *each of* its points.

Observe, though, that an open ball is itself an open set. And since each open ball is itself an open set, it is routine to show that a set U is open precisely when it is a union of open balls.

STUDENT: It seems like so much of this depends fundamentally on how we treat the boundary. None of these so-called open sets includes their boundaries.

MATHEMATICIAN: Yes, as I mentioned before, this gets at an important alternative characterization of open sets.

STUDENT: What if we had decided to treat boundaries differently?

MATHEMATICIAN: Great question. If we instead work with objects that include their boundaries—think of the circle including its boundary, or an interval including its end points—we are exhibiting instances of the notion of a *closed set*. And we could just as well have told this story using them!

STUDENT: Okay. Now we have a lot of words to describe when parts of a space equipped with a notion of distance are "near," which we can also use to home in on points. What's the big deal?

MATHEMATICIAN: Actually, we do not even need a notion of distance to develop such ideas.

STUDENT: What?! Explain!

MATHEMATICIAN: So far, we have basically been assuming that we already understand the underlying "space," the points of which we were comparing for nearness, interiority, and so on. But let us step back a bit.

Let us take these open set objects as our primitives and see what happens when we relate them. The most obvious thing we can try, for any two open sets, is to take their union or their intersection.

As for union, if we take two open disks and union them together, there is clearly no way of introducing any boundary, so the result will itself be open. But we can go further than this: we can union together any two open disks, and as we take the union of an arbitrary number of open disks, still no boundary can be introduced, so the union of any open disks must itself remain open. Moreover, trivial as it seems, notice the degenerate case as well: if we take the empty union—the union of no open disks, which is still technically a union!—this can only be the empty set $\emptyset$ itself. So this must also be open! Moreover, we can accordingly describe a set as open precisely when it is a union of open balls.

When taking intersections, the result will include the parts of the space that all of them have in common. First of all, notice the degenerate case here as well: if we take an empty intersection—the intersection of *no sets*—the result will be the entire space! In our case of the plane, this is all of $\mathbb{R}^2$, making the entire plane itself open.

Suppose we now take the intersection of two open disks:

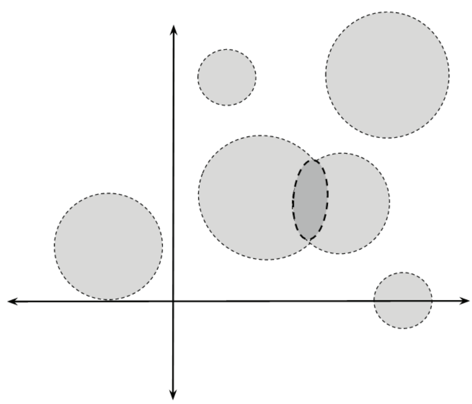

Observe how, for a point $p_0 \in D_1 \cap D_2$ in the intersection of the two open disks D_1 and D_2—the ones that are pictured as overlapping—we will have that $D(p_1, p_0) < \delta_1$ and $d(p_2, p_0) < \delta_2$. So if we just set $r > 0$ as

$$r = \min\{\delta_1 - d(p_1, p_0), \delta_2 - d(p_2, p_0)\},$$

and let

$$D = \{q \in \mathbb{R}^2 \mid d(p_0, q) < \frac{1}{2}r\},$$

then clearly $p_0 \in D \subseteq D_1 \cap D_2$, making p_0 an interior point of the intersection $D_1 \cap D_2$. This shows that the intersection of any two open disks will also be open.

STUDENT: And, as we could do with unions of opens, if we intersect an arbitrary number of open sets, then the result will be open?

MATHEMATICIAN: I was just about to speak to this. We cannot! This can easily be seen by looking at open intervals of the real line (exactly the same sort of argument applies to the plane).

Suppose, for some point $x \in \mathbb{R}$, I tell you that I can supply you with a "magical lens" (call it L_∞)

$$\bigcap_{n=1}^{\infty} \left(x - \frac{1}{n}, x + \frac{1}{n}\right).$$

For concreteness, let's just consider the point $x = 0$, so that this becomes

$$\bigcap_{n=1}^{\infty} \left(\frac{-1}{n}, \frac{1}{n}\right).$$

As n gets larger, this is of course like zooming in closer and closer, giving smaller and smaller windows of observation, around the point 0. As the windows of observation around the point 0 are getting arbitrarily small, the only thing that can be found in *every one* of the abstract lenses of the form $\left(\frac{-1}{n}, \frac{1}{n}\right)$ is the *point* $\{0\}$ itself. If *this*—the point $\{0\}$ itself—was allowed to be declared open as well, as you were suggesting, then we would have some problems.

There is no such problem in the case of restricting ourselves to *finite* intersections, for any finite intersection of open balls will be itself open.

STUDENT: I see. It does seem like it would be strange to countenance the idea that points themselves would be open, as there is no "room" in a point at all, let alone wiggle room.

MATHEMATICIAN: Right.

STUDENT: So, I see that, focusing on the plane for concreteness: any union of open sets will be open, and finite intersections of open sets will be open.

MATHEMATICIAN: Exactly. And recall the degenerate cases of union and intersection as well, which informed us that the whole plane $\mathbb{R}^2$ and the empty set must also be open.

STUDENT: Sure. But what is the point of all this again? How does this answer my question about how to develop the notion of open sets even without any notion of distance?

MATHEMATICIAN: Well, in order not to prejudice things, forget everything you know about "open discs" and "open intervals" and all that.

STUDENT: I'll try.

MATHEMATICIAN: Suppose only that we have some featureless objects—they are sets, but for now don't attribute any properties to them. To insist on this, let us just call such things *blahs*. Here is the punch line. Suppose

arbitrary unions of blahs results in a blah; and
finite intersections of blahs results in a blah.

In precisely this case—needing nothing more, knowing nothing of distances—we agree to call such blahs *open sets*.

STUDENT: Fine. Of course our familiar open intervals and balls from before will be accommodated by this definition. But so what?

MATHEMATICIAN: Well, two things are of importance to note. First, observe that we have now given a definition that does not rely on any intrinsic properties of certain special objects, but the objects in question are themselves entirely determined by the sorts of relations they entertain (specifically with respect to the operations of union and intersection).

STUDENT: Like recovering the space of the plane itself just by looking at how the approximating lenses must relate to one another?

MATHEMATICIAN: That is a good way of thinking about it.

And this lends itself to the second point. In our earlier special cases of the real line and plane, the open sets were characterized in terms of the distance function we were assuming. And such metrics in fact induced the structure of the opens, according to which the two conditions were met. One can appreciate this by considering, for instance, how our sketch of a proof of the finite intersection property made constant use of the distance function.

As it turns out, despite the interest we have in Euclidean spaces, like the plane, there are a number of interesting and conspicuous examples of structures of subsets of a given set that meet the two defining conditions supplied above. And some of these have no notion of distance at all. For such things, the structure of the open sets is not induced by any metric.

Mathematicians call any such structure of subsets of a set X that act as open sets, in the sense that the two conditions above are met—even where there is no attendant notion of distance—a *topology*, and the set X together with this topology forms what is called a *topological space*.

STUDENT: So even though we began with a notion of distance as we formulated our approximating lenses—which happened to be useful for telling us when things are near, and so useful ultimately for formulating other things, presumably like a viable concept of continuity—with this new and more general formulation, we see that the notion of distance must be auxiliary even to the notion of continuity (even though we first learn to think about it in terms of "sending nearby points to nearby points").

MATHEMATICIAN: Yes. But there is a great deal more. Many topological spaces are massively important in their own right, as any student of any branch of mathematics must come to appreciate.

STUDENT: Okay, I can believe that. But what if we had used closed sets instead of open sets, that is, if we wanted to explicitly reason with boundaries?

MATHEMATICIAN: Good question. Earlier, I hinted that closed sets and closed set topologies can just be described in a dual fashion, and this will require that their general definition be in terms of any subsets that respect

arbitrary intersections of closed sets results in a closed set; and
finite unions of closed sets results in a closed set.

STUDENT: I see. So, thinking in terms of boundaries, you could say

open set version: *without introducing boundaries, an arbitrary number of boundaryless objects can be joined together (without restriction), while only finite intersections of boundaryless objects are permitted (otherwise boundaries can be introduced where there were none before).*

closed set version: *without introducing boundarylessness, an arbitrary number of boundary-containing objects can be intersected together (without restriction), while only finite unions of boundary-containing objects are permitted (otherwise boundarylessness can be introduced where there were none before).*

MATHEMATICIAN: That is an interesting way to put things.

STUDENT: In dealing with a topological space, does it matter, then, whether we use open sets or closed sets to describe it?

MATHEMATICIAN: Yes and no. We can convert—in a purely formal way—back and forth between open and closed sets, appropriately dualizing the pairs of axioms as we do so, and arriving at the alternative definitions. So while closed and open sets will in general be different—and one must of course respect that difference—any topology that we can describe in terms of open sets can be given an alternative closed set formulation.

On the other hand, this is hardly the end of the story. And there are some structures that seem much more amenable to description in terms of one or the other. Moreover, some mathematicians seem to act as if they believe that, for the most part, open sets might be better or more natural to work with than closed—in part, perhaps because it can appear more useful, in certain conspicuous contexts (such as when proving important theorems of analysis and of familiar topological spaces), to have arbitrary unions of open sets being open instead of having arbitrary intersections of closed sets being closed. It might in the end have something to do with the nature of boundaries—but this is a more complicated matter that we should postpone for the time being.

STUDENT: It does seem like all this is really about the structure of approximations and the implications of certain decisions about how to treat boundaries. But it is okay with me if we leave that aside for now. There is something else I have been wondering about anyway, and the "dual" description in terms of closed sets only makes me wonder even more.

MATHEMATICIAN: Shoot.

STUDENT: If all we need to formulate the notion of topologies and topological spaces—in all the full glory of their generality—is a pair of axioms regarding behavior of the subsets with respect to their union and intersection; and if topologies really are so integral to so much of modern mathematics; can you help me better see *why* these governing axioms, as opposed to some other axioms describing some other governing features of collections of subsets, are the decisive ones?

MATHEMATICIAN: Well, recall how the open intervals and open discs of our familiar and beloved spaces behaved. It was there that we first observed these features, and there that so much of advanced inquiry into continuous phenomena using such approximating tools first flourished.

STUDENT: Hmm. But I feel that this doesn't really address my question.

MATHEMATICIAN: [*Awkward silence*]

STUDENT: For instance, in my physics class, we learned about light, how it behaves in certain characteristic ways—for instance, as it moves from one medium to another, it changes direction, and so on. Later, we learned that we had been regarding light as a wave, and that this was part of a bigger story about *waves in general* and *wavelike behavior*. While light exhibited wavelike behavior, we learned about other waves that were not light waves and began to look at this in a more general way.

Later, as our teacher was describing waves and considering some of the laws governing the behavior of waves in general, one student asked the teacher to justify, or at least to better clarify, the representation of such phenomena in terms of waves, the fully general formulation of waves and why the defining wavelike behavior was the way it was.

If our teacher had simply appealed back to the example of light and how light's behavior can be observed as conforming to the general pattern of wavelike behavior, I would have felt that this answer had failed to address what the student had been wondering.

MATHEMATICIAN: Hmm.

STUDENT: Open sets or closed sets, I still don't see why these two properties—the ones involving the asymmetry in the finiteness of the conditions on stability with respect to unions and intersections—are the key defining features of such an important concept. It does not seem that the *fact* that such features govern the usual objects of our familiar Euclidean metric spaces offers much in the way of justification or clarification of these axioms, especially as they are apparently governing for such a wide spectrum of structures, well beyond the confines of our familiar plane.

MATHEMATICIAN: [*Somewhat flustered*] I see. Well, we may just have to come back to this question. Before we take it up, let us take a closer and more structured look at all these, and some further, notions of topology. Then, we may return to address your very legitimate questions.[60]

4.3 Topology and Topological Spaces More Formally

As we have started to see, a topology on a set X is just a collection of subsets of X satisfying certain properties. There are several equivalent definitions of the notion of a topology, but the following is the one most commonly used:

Definition 77 Given a set X, a *topology* on X is a collection τ of subsets of X—that is, a subset $\tau \subseteq \mathbb{P}(X)$ of the power set—that satisfies the following properties:

1. The empty set $\emptyset$ and X are in τ.
2. Any union of elements in τ is also in τ, that is, whenever $\{U_i\}_{i \in I}$ is a family (finite or no) of subsets of X such that $U_i \in \tau$ for all $i \in I$, then $\bigcup_{i \in I} U_i \in \tau$.
3. Any finite intersection of elements in τ is also in τ, that is, whenever $U_1, U_2 \in \tau$, then $U_1 \cap U_2 \in \tau$ (and this can be extended to the finite case).

Elements of the topology τ are called *open sets*, while a set is called *closed* if and only if its complement is open.

A *topological space* is then just a pair (X, τ), where X is a set equipped with τ a topology on X.[61]

Observe that, strictly speaking, we do not actually need to write the first condition, in the above definition. Regarding the redundancy of requiring $\emptyset \in \tau$: consider the second axiom (stability under arbitrary union), and take the degenerate case of the empty family of subsets, that is, $I = \emptyset$, so that $\bigcup_{i \in \emptyset} U_i$, where this is just the set of all points x such

60. The interested reader will find these (and a few other) questions raised in the final section of this chapter, and a dedicated treatment of them in the appendix.

61. When the context is clear, it is not uncommon to refer to a topological space by just its carrier set, X, a practice we occasionally adopt.

that $x \in U_i$ for some $i \in \emptyset$. But *there are no such points x*—thus, we already have that $\bigcup_{i\in\emptyset} U_i = \emptyset$, which amounts to ensuring that $\emptyset$ is in our collection.

As for $X \in \tau$: consider intersections $\bigcap_{i\in I} U_i$ of subsets $U_i \subseteq X$, where these are again indexed by the empty set $I = \emptyset$. The intersection of the empty family of subsets of X is just the set of all $x \in X$ such that $x \in U_i$ for $i \in \emptyset$—but *every* x satisfies this property, so $\bigcap_{i\in\emptyset} U_i = X$.

In short, then, a topology on a set X can be defined as a collection of subsets of X that is stable under arbitrary unions and finite intersections.[62]

For open sets A, B, we might intuitively read the relationship $A \subseteq B$ as "*A approximates B*"—in other words, A is only a "partial specification" of what B specifies.

Example 78 The so-called *usual (or standard) topology* on $\mathbb{R}^n$ is the topology we get by taking for open sets all possible unions of open balls. So, specializing for instance to the real line $\mathbb{R}$, this is given by the collection of intervals of the form (a, b) along with arbitrary unions of such intervals.

Recall that an open ball is defined using d, the usual Euclidean metric, the notion of distance we experience every day. Accordingly, this topology is also sometimes called the *Euclidean topology*.

Example 79 The two examples described here represent extreme cases, and might seem uninteresting at first but they are rather important nonetheless.

First, take the collection $\mathbb{P}(X)$ of all subsets of X. This forms a topology—called the *discrete topology*—on X. In other words, with the discrete topology, *every* subset of X is open.

Applied to $X = \mathbb{R}$, the discrete topology can be associated with a particular metric (very different from the usual Euclidean metric!), namely the *discrete metric*, defined thus:

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

For the second example, at the other extreme, consider the following. By the first property in the definition of a topology, we know that a topology on X must contain both X and $\emptyset$. But notice that the set $\{\emptyset, X\}$, on its own, already forms a topology on X, called the *indiscrete* (or *trivial*) *topology*. In other words, with the indiscrete topology, the *only* sets counted as open sets are X and $\emptyset$. In particular, observe that this entails that the entire set X must be the only open set containing any point $p \in X$. In terms of the intuition, this latter consequence means that all points of the space are smushed together in such a way that, from the perspective of the topology, they cannot be distinguished.

Observe that the trivial topology on a set X will always be the topology with the *least* possible number of open sets—after all, to be a topology at all, the empty set and the entire set must be open, and the trivial topology just takes these as its *only* opens. By contrast, the discrete topology on a set X is the topology that has the *greatest* possible number of open sets—this is the topology that makes *every* subset open!

62. Instead of saying "stable," it is common to say that it is *closed* under arbitrary union and finite intersection. This terminology is entirely sensible, but it has nothing to do with "closed" in the topological sense, so in order to avoid confusing the reader this terminology is generally avoided in this chapter.

As the previous example already suggested, occasionally two topologies on the same set will be comparable. Topologies need not be comparable, but when they are, the following language is useful:

Definition 80 For two topologies τ and τ' on the same set, if $\tau \subseteq \tau'$, we say that the topology τ is *coarser* (smaller) than τ'—or, what is the same, that the topology τ' is *finer* (larger) than τ.

Here, to remember the terminology—where "coarse" means a smaller number of subsets in the collection, and "finer" means a larger number of subsets in the collection—topologists sometimes propose the following sort of analogy: think of some sort of grinder breaking up the pieces of something, like coffee beans. If the grinder is set to "fine" grind, then it will break things up into a great number of pieces; on the other hand, if one uses a "coarse" setting, one will be left with a smaller number of pieces.

Exercise 1 Take $X = \{a, b, c, d, e\}$. Establish whether or not each of the following collections of subsets of X forms a topology on X:

1. $\tau_1 = \{X, \emptyset, \{a\}, \{a, b\}, \{a, c, d\}, \{a, b, c, d\}\}$.
2. $\tau_2 = \{X, \emptyset, \{a, b, c\}, \{a, b, d\}, \{a, b, c, d\}\}$.

Solution

1. Yes, this is a topology, since it satisfies the three axioms of the definition, as you can check manually.
2. No, this is not a topology, since $\{a, b, c\}$ and $\{a, b, d\}$ each belong to τ_2, yet their intersection does not, that is, $\{a, b, c\} \cap \{a, b, d\} = \{a, b\} \notin \tau_2$, violating the third axiom.

Exercise 2 Take $X = \mathbb{Z}$, the set of integers. Then the collection $\mathscr{C}$ of all finite subsets of the integers plus $\mathbb{Z}$ itself is *not* a topology. Why not?

Solution Well, in particular, the union of all finite subsets of $\mathbb{Z}$ not containing zero or any negative numbers

$$\{1\} \cup \{2\} \cup \cdots \cup \{n\} \cup \cdots = \{1, 2, \ldots, n, \ldots\}$$

will not be finite (and so cannot be in the given collection), and it is also not all of $\mathbb{Z}$. Thus, while each of the members of the above union are in $\mathscr{C}$, the union itself does not belong to $\mathscr{C}$. Thus, $\mathscr{C}$ does not satisfy axiom 2—the property of stability under arbitrary union—disqualifying it from being a topology.

The dialogue in the previous section hinted that the notion of a topological space was truly more general than that of a metric space. We speak of a space as *metrizable* if there is a metric that induces the topology, that is, the topology can be regarded as coming from a particular metric. If not, then it is said to be *nonmetrizable*. To show the topology notion to be a proper generalization, we need only exhibit a space that is nonmetrizable. There are in fact a great many nonmetrizable spaces!

Exercise 3 Take any set X that consists of more than one element. Take the indiscrete topology on X. Is the resulting space metrizable?

Solution It is not! For instance, take the set $X = \{1, 2\}$ and let τ be the indiscrete topology on X. We can show that τ is not induced by any metric one could put on X (and the same would be true for any arbitrary set with two or more elements).

Suppose we had a metric d on X. We can set $r = d(1, 2) > 0$, from which we will have that the open ball $B(1, r) = \{1\}$. Thus, in the topology induced by the metric d, the set $\{1\}$ will have to be open in X. Yet $\{1\}$ is not open in the indiscrete topology!

On the other hand, if X had only one element, the space would be trivially metrizable.

Exercise 4 Find the smallest topological space that is neither trivial nor discrete.

Solution On a set with only one element $X = \{x\}$, there is only one topology it can admit, namely $\{\emptyset, \{x\}\}$—and here, the trivial and discrete topologies coincide.

So take a two-point set $\{0, 1\}$. Take for the collection of open sets

$$\{\emptyset, \{1\}, \{0, 1\}\}.$$

This forms a topology, and one that is neither trivial nor discrete.

This particular space is rather special, and usually goes under the name of the *Sierpiński space* $\mathbb{S}$.

Example 81 Let X be the set $\mathbb{R}$, but for open sets all possible unions of (what, in the usual context of $\mathbb{R}$, we typically think of as) "half-open" intervals of the form

$$[a, b)$$

for $a, b \in X$.

This forms a topology called the *lower limit topology*, yielding a space sometimes called the *Sorgenfrey line*. Observe that this topology is finer, that is, has more open sets, than the usual topology on $\mathbb{R}$ (generated by the open intervals as basis).

Notice that, with respect to the usual topology on $\mathbb{R}$, with $a < b, a, b \in \mathbb{R}$, these intervals $[a, b)$ are neither open nor closed! However, with respect to the Sorgenfrey space, they are open (and closed!).

This example helps to reinforce an essential idea: *openness* (*closedness*) depends entirely on the topology in question. We should not treat these notions as involving some "inherent properties" of a particular set. A single set can carry many different topologies, and strictly speaking it is important to remember that there is no such thing as an "open set" in itself—only an open subset, that is, a set open in relation to some set. A given subset U of a space X is (or is not) a subset *open in* (or *open with respect to*) X. In practice though, which topological space we are working with will generally be obvious or clearly stated, so we will often simply speak (somewhat misleadingly) of certain sets as being "open."

Let us now investigate a little more closely a few other notions that arose in the dialogue of the previous section.

Definition 82 For p a point in a topological space X, we say that a subset $N \subseteq X$ is a *neighborhood* of p if and only if N is a superset of an open set G that contains p, that is,

$$p \in G \subseteq N \text{ where } G \text{ is an open set.}$$

Exercise 5 Suppose, on $X = \{a, b, c, d, e\}$, we have the following topology:

$$\tau = \{X, \emptyset, \{a\}, \{a, b\}, \{a, c, d\}, \{a, b, c, d\}, \{a, b, e\}\}.$$

Supply all the neighborhoods of

1. the point e;
2. the point c.

Solution

1. What are the open sets that contain e? These are $\{a, b, e\}$ and the entire set X. A neighborhood of e is just any superset of an open set that contains e. Thus, we must look for the supersets of $\{a, b, e\}$ and X.

   The supersets of $\{a, b, e\}$ are $\{a, b, e\}$, $\{a, b, c, e\}$, $\{a, b, d, e\}$, and X, while the only superset of X is of course X. Thus, the neighborhoods of e are given by

$$N_e = \{\{a, b, e\}, \{a, b, c, e\}, \{a, b, d, e\}, X\}.$$

2. What are the open sets that contain c? These are $\{a, c, d\}$, $\{a, b, c, d\}$, and X. Again, supersets of these sets will give us the neighborhoods, that is,

$$N_c = \{\{a, c, d\}, \{a, b, c, d\}, \{a, c, d, e\}, X\}.$$

Instead of considering the relation

N is a neighborhood of point p,

we might look at things from the point's perspective. Doing so gives us an inverse relation

p is an *interior point* of N.

We can accordingly define a point x of a subset U of a topological space X to be an *interior point* for U if and only if U is a neighborhood of x. By taking the set of all interior points of U—that is, taking the union of all open subsets of U—we get the *interior of* U, which is denoted $\mathbf{int}(U)$.

Exercise 6 Take the topology from the previous exercise. Find the interior points of the subset $A = \{a, b, c\} \subseteq X$.

Solution Since each of

$$a, b \in \{a, b\} \subseteq A,$$

with $\{a, b\}$ an open set that is contained in A, a and b are each interior points of A. What about c? Well, c does not belong to any open set contained in A, and so it is not an interior point of A. Thus, $\mathbf{int}(A) = \{a, b\}$.

Observe that the interior of a set A is the union of all open subsets of A. Moreover, one can show that

- $\mathbf{int}(A)$ is itself open;
- $\mathbf{int}(A)$ is the *largest* open subset of A; and
- A is open if and only if $A = \mathbf{int}(A)$.

The third fact is especially decisive. After all, a topology is just determined by those sets declared open. Thus, the fact that a set will be open precisely when it is equal to its interior supplies us with an alternative way of construing a topology. It is important enough that we set it apart:

Theorem 83 A subset U of a topological space X is open precisely when $\mathbf{int}(U) = U$.

In this way, open set topologies can be determined by *taking the interior*, where there will be certain conditions on how this must behave. Specifically, construing things in these terms, we get an alternative (but ultimately equivalent) definition of a topological space.

Definition 84 A topological space is a pair $(X, \mathbf{int})$, for X a nonempty set and $\mathbf{int} : \mathbb{P}(X) \to \mathbb{P}(X)$ an operation satisfying the four so-called Kuratowski axioms:

1. (**i1**) $\mathbf{int}(X) = X$ (it preserves the total space);
2. (**i2**) $\mathbf{int}(A) \subseteq A$ (it is intensive);
3. (**i3**) $\mathbf{int}(\mathbf{int}(A)) = \mathbf{int}(A)$ (it is idempotent);
4. (**i4**) $\mathbf{int}(A \cap B) = \mathbf{int}(A) \cap \mathbf{int}(B)$ (it preserves binary intersections).

Observe that the family of open sets can then just be defined by setting $\tau = \{A \subseteq X : \mathbf{int}(A) = A\}$. We will have much more to say about this operator throughout this chapter and the appendix.

For now, though, let us continue to introduce further essential notions of basic topology. In the elementary examples presented thus far, we could generally specify a topology by simply explicitly describing the entire collection of open sets. But such a specification will typically be unfeasible for many topologies you might meet out in the wild. In most cases in practice, one instead specifies a smaller collection of subsets of X that then "generate" the topology in question and so can be used to define that topology. And even when we could explicitly specify the open sets of a topology, it is sometimes easier to understand, and work with, a description in terms of the smaller collection. The following notion serves such purposes.

Definition 85 For X a set, a *basis* for a topology on X is a collection $\mathcal{B} \subseteq \mathbb{P}(X)$ of subsets of X, called the *basis elements*, satisfying

1. ($\mathcal{B}$ "covers" X) for each $x \in X$, there is at least one basis element $B \in \mathcal{B}$ that contains x—in other words, $X = \bigcup\{B : B \in \mathcal{B}\}$; and
2. if x belongs to the intersection of two basis elements, that is, $x \in A \cap B$ where $A, B \in \mathcal{B}$, then there is at least one basis element $C \in \mathcal{B}$ such that $x \in C \subseteq A \cap B$.

Then, the topology τ *generated* by the basis $\mathcal{B}$ is defined as the coarsest topology containing $\mathcal{B}$.

Another way of writing the first of the conditions placed on a basis of a topology τ on a set X would be to say that every open set $U \in \tau$ is the union of members of $\mathcal{B}$—so that, in particular, $X = \bigcup\{B : B \in \mathcal{B}\}$. In short, a basis specifies a topology by taking unions, that is, starting with a basis on X, by adding to it all possible unions of basis elements, the collection we end up with will be a topology on X. A basis for a topological space X is thus just a collection $\mathcal{B}$ of open subsets of X, such that every open subset of X is a union of sets in $\mathcal{B}$.

Note, though, that the expression for an open U as a union of basis elements is not unique—a topological space can have several different bases. Thus, the terminology of "basis" here is not to be conflated with the use of the same term in linear algebra, for instance, where the expression of a vector as a linear combination of basis vectors is indeed unique.

Example 86 The basis that consists of the open intervals in $\mathbb{R}$ (or open discs in $\mathbb{R}^2$) generates the usual topology on $\mathbb{R}$ (on $\mathbb{R}^2$).

For any metric space X, the open balls also form a basis for the induced topology on X.

Example 87 The collection of all the singletons (one-element subsets) of a set X is a basis for the discrete topology on X: after all, they are all open in this topology, and an arbitrary open set is the union of its singleton subsets.

Example 88 Let $\mathscr{B} = \{[a, b) \subseteq \mathbb{R} : a < b\}$. This forms a basis on $\mathbb{R}$. The topology that it generates is the lower limit topology on $\mathbb{R}$, presented in example 81—resulting in the topological space called the Sorgenfrey line.

Exercise 7 Let $X = \{a, b, c\}$ and consider the collection $\mathscr{B} = \{\{a, b\}, \{b, c\}\}$. Is $\mathscr{B}$ a basis for a topology on X?

Solution No. While of course $\{a, b\} \cup \{b, c\} = \{a, b, c\}$, if $\mathscr{B}$ were a basis, then $\{a, b\}$ and $\{b, c\}$ would each have to be open, and thus their intersection

$$\{a, b\} \cap \{b, c\} = \{b\}$$

would also be open. Yet $\{b\}$ is not the union of any members of $\mathscr{B}$.

Finally, given a topological space (X, τ) and a basis $\mathscr{B}$ on X, it is often convenient to use the following to check whether $\mathscr{B}$ indeed generates τ:

Corollary 89 For (X, τ) a topological space, and $\mathscr{B}$ a basis on X, $\mathscr{B}$ generates τ if and only if

1. $\mathscr{B} \subseteq \tau$; and
2. for every open set $U \in \tau$ and every $x \in U$, there exists a $B \in \mathscr{B}$ such that $x \in B \subseteq U$.

Finally, it is worth mentioning that we do need not restrict the above notions only to open sets—they apply just as well to closed sets, which we take up more explicitly in the section to follow.

Exercise 8 What topology τ on the real line $\mathbb{R}$ is generated by the collection $\mathscr{A}$ of all closed intervals $[a, a+1]$ of length 1?

Solution Consider any point $p \in \mathbb{R}$. Then the closed intervals $[p-1, p]$ and $[p, p+1]$ both belong to $\mathscr{A}$. Thus, their intersection $[p-1, p] \cap [p, p+1] = \{p\}$ must also belong to the topology τ. Thus, all singleton sets $\{p\}$ are open in this topology, giving us again the discrete topology on $\mathbb{R}$.

4.3.1 Closed Sets

In discussing topologies, we have confined ourselves thus far, as is common, to specifying the *open* sets. But if we appropriately dualize things, we can arrive at an alternative description of things in terms of closed sets and closed set topologies. Fundamentally, open and closed set topologies are formally distinct only in how nonfinite collections of elements of the topology are treated. Just building on what we already know of open sets,

Definition 90 A subset A of a topological space (X, τ) is *closed* if and only if its complement

$$X \setminus A$$

(also denoted A^c when the overall space is understood) is an open subset of (X, τ), that is, when $A^c \in \tau$.

Example 91 Since $X \setminus \emptyset = X$ and $X \setminus X = \emptyset$, both $\emptyset$ and X are always closed, in any topology.

Example 92 For the real number line $\mathbb{R}$ with the usual topology, the canonical closed intervals $[a, b]$ of real numbers

$$[a, b] = \{x \in \mathbb{R} \mid a \leq x \leq b\}$$

give us closed sets. To see this, using the definition just presented, one need only recognize that the complement

$$\mathbb{R} \setminus [a, b] = (-\infty, a) \cup (b, \infty) = \{x \in \mathbb{R} \mid x < a \text{ or } x > b\}$$

is open. Yet this can be expressed as a union of open intervals, for example, $(-\infty, a) = \bigcup_{n \in \mathbb{N}} (a - n, a)$, and similarly for (b, ∞).

Moreover, with such closed intervals, we can also take $a = b$, resulting in the closed interval $[a, a]$ or the single point $\{a\}$, making $\{a\}$ itself a closed subset of $\mathbb{R}$. Indeed, for most topological spaces, it is typical to find that single points are closed subsets. In particular, in $\mathbb{R}$ with the usual topology, each singleton set $\{a\}$ is closed, for the complement of $\{a\}$ is the union of the two open sets $(-\infty, a)$ and (a, ∞), which is open.

Similarly, the set $\mathbb{Z}$ of all integers is a closed subset of $\mathbb{R}$—which fact can be seen by acknowledging that the complement of $\mathbb{Z}$ is the union $\bigcup_{n=-\infty}^{\infty} (n, n+1)$ of the open subsets of the form $(n, n+1)$, which is itself open in $\mathbb{R}$.

Example 93 Given $X = \{a, b, c, d, e\}$ with the topology

$$\tau = \{X, \emptyset, \{a\}, \{c, d\}, \{a, c, d\}, \{b, c, d, e\}\},$$

the closed subsets of this topological space are given by the complements of the open subsets of the space. Thus, for closed sets we have

$$\emptyset, X, \{b, c, d, e\}, \{a, b, e\}, \{b, e\}, \{a\}.$$

As you can observe, there are subsets of X, like $\{b, c, d, e\}$, that are both open and closed; while there are subsets, like $\{a, b\}$, that are neither open nor closed.

Example 94 Take X to be a discrete topological space, that is, one where every subset of X is open. But then, every subset of X will also be closed, since its complement will automatically be open. Thus, in a discrete space, all subsets of X are both open and closed.

Exercise 9 Given the set $X = \{a, b, c\}$, come up with a topological space that is not discrete, where the closed sets and open sets are the same sets.

Solution Here is one:

$$\tau = \{X, \emptyset, \{a, b\}, \{c\}\},$$

which you can manually check has the stipulated properties.

The example that follows is an unusual one, meant only to suggest that things don't always behave as one might expect, especially in relation to distances.

Example 95 One might expect that closed sets that were disjoint would have to be some distance apart. But we can describe two disjoint closed sets in the plane, that are at zero distance apart.[63] Let $C_1 = \{\langle x, r\rangle \mid xy = 1\}$ and $C_2 = \{\langle x, y\rangle \mid y = 0\}$ = the x axis. Then $C_1 \cap C_2 = \emptyset$ (they are disjoint), and both are closed. Now, for any $\epsilon > 0$, the points $\left(\frac{2}{\epsilon}, \frac{\epsilon}{2}\right) \in C_1$ and $\left(\frac{2}{\epsilon}, 0\right) \in C_2$ are at distance $\frac{1}{2}\epsilon < \epsilon$.

Because we have merely used the operation of complement to arrive at the corresponding notions of closed sets, we should not be surprised that, by some elementary facts of set theory, we can arrive at a "dual" definition of a topology in terms of closed sets. In particular, first recall the following:

Lemma 96 (deMorgan's Laws: *taking complements turns intersections into unions and unions into intersections*) Let X be a set and $(U_i)_{i\in I}$ a collection of subsets of X. Then,

$$X \setminus \left(\bigcup_{i\in I} U_i\right) = \bigcap_{i\in I}(X \setminus U_i)$$

and

$$X \setminus \left(\bigcap_{i\in I} U_i\right) = \bigcup_{i\in I}(X \setminus U_i).$$

Given a topology built from open sets, then, by taking complements of open sets and using deMorgan's laws, we can define a topology in terms of closed sets, a *closed set topology*.

Theorem 97 Take X a topological space. Then the collection of closed subsets of X will have the following properties:

1. The empty set $\emptyset$ and X are closed sets.
2. The intersection of any number of closed sets is closed.
3. The union of any finite number of closed sets is closed.

Observe how unions became intersections and intersections unions, yielding alternate finiteness conditions than what we had for open sets.

Exercise 10 It is fairly routine to check the statements made above, for instance that the intersection of any number of closed sets is closed. For the condition on unions: you should at least try to give an example of an infinite union of closed sets that is not closed, helping you see why we must restrict to finite unions when dealing with closed sets.

Solution One such example is given by taking the union of all closed intervals of the form $\left[\frac{1}{n}, 1 - \frac{1}{n}\right]$. Observe that the infinite union equals $(0, 1)$, which is surely not a closed set.

Another example could be given by taking an infinite union of singleton sets $\left\{\frac{1}{n}\right\}$, which you can check will not be closed.

63. This example was taken from Gelbaum and Olmsted 1962.

With such a result, we can give an alternative definition of a topology, now in terms of closed sets: namely as a set X together with a collection τ of closed subsets of X. Moreover, because $X \setminus (X \setminus U) = U$, so that $(\tau^c)^c = \tau$, if you know what the closed sets are, you also know what the open sets are (and conversely).

In certain cases, however, it may be more natural to check the dual topology axioms using closed sets instead of open sets; and in some settings, it may simply be more natural to work with, or define things in terms of, closed sets.

Example 98 In general, a subset A of a set X that is such that its complement in X is a finite set is called a *cofinite* subset. For any set X, we can define a topology that takes for its open sets the empty set and all cofinite subsets of X, that is,

$$\tau = \{A \subseteq X \mid A = \emptyset \text{ or } X \setminus A \text{ is finite}\}.$$

In this topology—called the *cofinite topology*—all finite sets are closed (and they, together with all of X, are the only subsets that are closed). In other words, a subset $Z \subseteq X$ is closed in the cofinite topology if and only if it is finite or equal to X. And it is straightforward to verify that

1. the empty set is finite, and X is equal to X;
2. an arbitrary intersection of finite sets is finite;
3. a finite union of finite sets is finite.

Accordingly, such a topology appears to have a more natural or direct description, and the topology axioms are more easily verified, when we work with closed sets instead of open sets.

This topology is rather interesting and does indeed show up a fair amount in the wild. Moreover, it is the coarsest topology that satisfies the natural condition that singleton sets are regarded as closed. Before moving on, let us explore a few other features of this topology.

First of all, observe that the definition of a cofinite topology does *not* stipulate that every topology on X that has X and the finite subsets of X as closed is automatically the cofinite topology, as can be appreciated by considering that in the discrete topology on a set X, the set X and all finite subsets of X will be closed (yet many—all!—other subsets of X are closed as well). Rather, the definition says that X and the finite subsets of X have to be the *only* closed sets.

Second, it is not the case that infinite subsets will necessarily be open sets, as the following illustrates. Taking $\mathbb{N}$ the set of all positive integers, then of course sets such as

$$\{1\}, \{4, 6, 8\}, \{5, 6, 7, 8\}$$

are finite and thus closed in the cofinite topology. Therefore, their complements

$$\{2, 3, 4, 5, \dots\}, \{1, 3, 5, 7, 9, \dots\}, \{1, 2, 3, 4, 9, 10, \dots\}$$

will be open sets in the cofinite topology. However, the set of even positive integers, for instance, is not a closed set in this topology, since it is not finite. Thus, its complement, the set of odd positive integers, is *not* an open set in this topology. In other words, while it is true that all finite sets are closed in this topology, not all infinite sets will be open.

This is a portion of the eBook at doi:10.7551/mitpress/12581.001.0001

Again reinforcing a point we made earlier, observe how a set such as $\{n \mid n \geq 12\}$ is open in the cofinite topology on $\mathbb{N}$. However, this same set is *not* open in the indiscrete topology, for instance. Similarly, the set of even natural numbers is open in the discrete topology on the natural numbers, but it is *not* open in the cofinite topology. Again, it is easy to forget that the notions of "open" and "closed" entirely depend on the topology—we should not treat these notions as involving some "inherent properties" of a particular set.

Moreover, it is important to realize that, for a given topology, sets can be both open and closed, neither open nor closed, open but not closed, or closed but not open. So one should be sure to appreciate that we cannot prove that a set is open by proving that it is not closed.

As we did for open sets, we can define further notions corresponding to those we had for open sets. For instance, while open sets had interior points, closed sets have *limit* (or *accumulation*) *points*—which notion gives us another way of conceptualizing being closed.

Definition 99 We call a point x of a subset A of X a *limit* (or *accumulation*) *point* of A iff every neighborhood of x contains at least one point of A different from x. In other words, a point $p \in X$ is a limit point of a subset A of X iff

$$G \text{ open}, p \in G \text{ implies } (G - \{p\}) \cap A \neq \emptyset.$$

Thus, to show that a point y is *not* a limit point of A, it is enough to find even just one open set that contains y but no other point of A.

Example 100 Take the topological space (X, τ), where $X = \{a, b, c, d, e\}$ with the topology

$$\tau = \{X, \emptyset, \{a\}, \{c, d\}, \{a, c, d\}, \{b, c, d, e\}\}.$$

Consider the set $A = \{a, b, c\}$. The points a and c are *not* limit points of A—for, the set $\{a\}$ is open in τ and contains no other point of A, and the set $\{c, d\}$ is open in τ and contains c but no other point of A. On the other hand, b, d, and e are each limit points of A. To show that a point is a limit point of A, we just have to show that every open set containing that point contains a point of A other than it. So, for example, b is a limit point of A, as the only open sets containing b are X and $\{b, c, d, e\}$, each of which contains another element of A. Notice, finally, that d and e are indeed limit points of A, even though they are not in A.

Example 101 If (X, τ) is a discrete space, and A any subset of X, then A will have no limit points. Observe that for each $x \in X$, the singleton $\{x\}$ is an open set that contains no point of A different from x.

Example 102 In $\mathbb{R}$ with the usual topology, every element in the interval $[a, b]$ is a limit point of $[a, b]$.

Similar to what happened with open sets and their interior points, this notion of "limit point" gives us a useful way of characterizing which sets are closed.

Corollary 103 A subset A of a topological space (X, τ) is closed if and only if A contains all of its limit points.

The set of limit points of a set A is sometimes called the *derived set* of A, denoted $A^{'}$. Using this notation, the above just says that a set A is closed if and only if $A^{'} \subseteq A$.

Example 104 The set $[a,b]$ is closed in $\mathbb{R}$ with the usual topology, since all the limit points of $[a,b]$ are in $[a,b]$. By contrast, the set $[a,b)$ is not closed in $\mathbb{R}$, since b is a limit point, yet $b \notin [a,b)$.

Here is a strange and somewhat counterintuitive example.

Example 105 Suppose X is any set with more than one point. As we observed in example 79, if we let (X,d) be the metric space equipped with the discrete metric

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

this will induce the discrete topology on X. Now suppose x is any point of X, and let O be the open ball (in this space), with center x and radius 1, while C is the closed ball with center x and radius 1. Observe that we must then have that $O = \{x\}$, while $C = X$. But, as the topology is necessarily discrete, $\mathbf{cl}(O) = O \neq B$.

Altogether, this shows that we can produce a space for which we have open and closed balls, O and C, each with the same center and same radius, such that C is *not* equal to the closure of O!

In many other familiar spaces, such an example would be impossible.

Observe how on the real number line, there is a natural notion of "closeness." For instance, each point in the sequence

$$0.1, 0.01, 0.001, 0.0001, \ldots$$

is closer to the point 0 than the previous point in the sequence. 0 is clearly a limit point of this sequence, so an interval such as $(0,1]$ cannot be closed in $\mathbb{R}$, since it does not contain the limit point 0. However, in general topological spaces, we have seen that they need not be accompanied by a distance function. Thus, we make use of the notion of "limit point" that does not rely on distances. Yet even with this more general definition, we can still ensure that, in cases such as the above, the point 0 will remain a limit point of $(0,1]$, as we would have expected.

Just as, earlier, we were able to describe an interior operator **int** and use this to present an alternative definition of a topology, there is a dual "closure" operator, **cl**, where for any $A \subseteq X$, $\mathbf{cl}(A)$ is defined as the union of A and its limit points, that is,

$$\mathbf{cl}(A) = A \cup A'.$$

It should be apparent that $\mathbf{cl}(A)$ is itself a closed set. An important observation is that for S, T nonempty subsets of a topological space (X, τ), with $S \subseteq T$, if p is a limit point of S, it is straightforward to show that p must also be a limit point of the set T. But this means that every closed set containing A must also contain A'. Thus, $A \cup A' = \mathbf{cl}(A)$ must in fact be the *smallest* closed set containing A—in particular, $\mathbf{cl}(A)$ will be the intersection of all closed sets containing A. Altogether, a set A will then be closed (relative to a topology) precisely when $\mathbf{cl}(A) = A$. In this manner, we can begin to appreciate that closed set topologies can be determined by closure operators, just as we were able to give an alternative definition of an open set topology in terms of the interior operator.

Exercise 11 Recall the definition of an interior operator. I have suggested that there is an operator dual to the interior operator—and this is the closure operator. Referring back to the four axioms governing the interior operator, introduced in definition 84, how might you define the dual notion of a closure operator?

Solution Recalling that one can convert between open and closed sets by taking complements, we arrive at the following definition:

Definition 106 A *closure operator* on a set X is a mapping $\mathbf{cl} : \mathbb{P}(X) \to \mathbb{P}(X)$ that satisfies the four so-called Kuratowski closure axioms, for all $A, B \subseteq X$:

1. (**k1**) $\mathbf{cl}(\emptyset) = \emptyset$ (it preserves the empty set);
2. (**k2**) $A \subseteq \mathbf{cl}(A)$ (it is extensive);
3. (**k3**) $\mathbf{cl}(A) = \mathbf{cl}(\mathbf{cl}((A)))$ (it is idempotent);
4. (**k4**) $\mathbf{cl}(A \cup B) = \mathbf{cl}(A) \cup \mathbf{cl}(B)$ (it preserves binary unions).

Moreover, a fact we alluded to earlier, that

- (**k4′**) $A \subseteq B \Rightarrow \mathbf{cl}(A) \subseteq \mathbf{cl}(B)$,

is a consequence of the fourth axiom.

Exercise 12 Returning to example 93, where the closed sets were

$$\emptyset, X, \{b, c, d, e\}, \{a, b, e\}, \{b, e\}, \{a\},$$

what is the closure of the sets $\{b\}$, the set $\{a, c\}$, and the set $\{b, d\}$?

Solution The closure $\mathbf{cl}(A)$ of any set A is the intersection of all closed supersets of A. To find the closure of a particular set, we need only find all the closed sets containing that set and then pick the smallest. Thus, in our present case, $\mathbf{cl}(\{b\}) = \{b, e\}$, while $\mathbf{cl}(\{a, c\}) = X$, and $\mathbf{cl}(\{b, d\}) = \{b, c, d, e\}$.

Example 107 In any discrete space, as every set is closed (and open!), every set is equal to its closure.

Exercise 13 Let $S = \left\{1, \frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \ldots, \frac{1}{n}, \ldots\right\}$. Show that S is not closed in the usual topology on $\mathbb{R}$.

Solution 0 is a limit point of S, but $0 \notin S$, so S cannot be closed. If, instead, we had considered the same set, but now with 0 added, then it would of course be closed in $\mathbb{R}$.

Note also that S is not open either. Check that you can see why.

As should be evident from some of the examples just considered, the closure of a set U is in general bigger than U itself. As such, it is natural to want to consider what is in the closure of a set without being in the set itself. This gives rise to an operator ∂, called the *boundary*:

$$\partial(U) := \mathbf{cl}(U) \setminus U.$$

Note that for a set in a topology, we could also say that

$$\partial(U) := \mathbf{cl}(U) \setminus \mathbf{int}(U)$$

or

$$\partial(U) := \mathbf{cl}(U) \cap \mathbf{cl}(U^c),$$

where this last formulation describes the boundary of a set U as consisting of those points interior neither to U nor to its complement U^c (i.e., $X \setminus U$). As such, the intuition here is that the boundary of a subset $S \subseteq X$ consists of those points in X that are neither "fully in" S nor "fully not in" S. We typically think of the boundary of a set as the exterior surface or skin of a body. Yet the boundary of a set may be larger than the set itself—for example, since it can be shown that the rationals $\mathbb{Q}$ are dense in $\mathbb{R}$, $\partial\mathbb{Q} = \mathbb{R}$.

Another valuable way of defining the boundary of a subset S of a topological space X—now seeing things from the perspective of points—is as consisting of the set of points $p \in X$ such that every neighborhood of p contains at least one point of S *and* at least one point not of S. Here, such an element of the boundary of S is called a *boundary point* of S.

Example 108 Taking $\mathbb{R}$ with the usual topology,

$$\partial((0,3)) = \partial([0,3)) = \partial((0,3]) = \partial([0,3]) = \{0,3\},$$

just as one might have imagined.

While one might typically think of a boundary in the way the previous example supports, or imagine the boundary in intuitive terms of the edges of a circle or figure in the plane, boundaries can be strange, as the following example illustrates.

Example 109 One can construct three disjoint subsets of the plane, where these share a common boundary. (In fact, this works for any finite number of disjoint regions.)[64]

A few other things are worth observing at this point. In general, a set U may have limit points x where $x \notin U$, so $\partial(U)$ may in general be nonempty. Moreover, the boundary of a set is always closed, and a closed set will contain its own boundary. Altogether, the notion of boundary gives us yet another characterization of a closed set.

Definition 110 A set is closed if and only if it contains all of its boundary points.

Note also that the closure of a set can thus be expressed as the union of the set with its boundary,

$$\mathbf{cl}(S) = S \cup \partial(S),$$

the smallest closed set that contains S.

The above characterization of closed sets in terms of boundary points moreover informs us that a set is open if and only if it is disjoint from its boundary. Altogether, this characterization leaves us with a very useful way of thinking about closed versus open sets, one that was already anticipated in the earlier dialogue: a set is closed iff it contains its boundary, and thus open iff it is disjoint from its boundary.

At this point, the reader might be wondering: we have seen that we call a set closed if it contains all its limit points, and we have just seen that we can also express a set as closed if it contains all its boundary points—does this mean that these are two words for the same thing? No! Limit points and boundary points are not reducible to one another. In particular, a limit point can be a boundary point, but it need not be a boundary point, as a limit point can also be an interior point (recall that a boundary point is not an interior point). One can

64. See Gelbaum and Olmsted (1962, 138) for a description of this, and a nice story to help with the (counter)intuition.

also appreciate the difference by considering the interval $S = [0, 1]$ with the usual topology. Here, each element of $[0, 1]$ is in fact a limit point of S, while there are only two boundary points—namely, 0 and 1.

Likewise, in general, a certain point may be a boundary point, yet not a limit point. For instance, in $\mathbb{R}$ with the usual topology, the point 0 is a boundary point of the set $\{0\}$, but it is not a limit point of that set. Note that there may also be points that are both a limit point and a boundary, and points that are neither. However, again in general, it is often useful to realize that a limit point of a set that is not an element of the set itself will always be a boundary point.

In short, while a point $p \in X$ is a limit point of a subset A of X iff every neighborhood of p also contains a point of A other than p itself (which is something that can hold for an interior point as well), a point $p \in X$ is a boundary point of a subset A of X iff every neighborhood of p contains at least one point of A and at least one point not of A. While the notions are not the same, we can characterize closed sets in terms of either notion.

4.3.2 Covers

There is a final basic notion of topology, one that will be especially important for us in the development of sheaves: that of a *cover*. A common thing to do in mathematics is to approximate complicated objects or structures by means of simpler, more basic ones. The topological notion of a cover is a powerful way of doing so, allowing us to shift from a topological perspective to a more combinatoric perspective. Moreover, the notion of a cover will play a key role in the definition of sheaf, so we close out the treatment of the basic notions of general topology by devoting some attention to this notion.

A *cover* of a given subset S of a topological space X is any collection of subsets of X whose union contains S. In other words, a cover for a set S is just a particular bunch of (possibly overlapping) sets such that S is completely contained in that bunch of sets. Purely in terms of set theory, a cover of a set X is thus just a family $\mathscr{C} \subseteq \mathbb{P}(X)$ of subsets of X such that $X = \bigcup \mathscr{C}$. The set being covered can of course be the entire space itself X, in which one speaks of a cover of the space. But more generally, we can just define a cover for any subset of a space.

Definition 111 Let X be a topological space and $S \subseteq X$. Then an (*open*) *cover* of S is a collection $\{U_i\}_{i \in I}$ of (open) subsets U_i whose union contains S, that is, $S \subseteq \bigcup_{i \in I} U_i$. Note that if $S = X$, then this is just to require that

$$\bigcup_{i \in I} U_i = X.$$

If we are working with open sets, given an open set topology, we will often just speak of an *open cover*, where this is exactly the same thing as a cover, except we specify that all the members of our collection of sets doing the covering are themselves open sets. But observe that this definition would work just as well for a *closed cover*, where this would just be a collection $\{C_i\}_{i \in I}$ of closed subsets of X such that $\bigcup_{i \in I} C_i = X$. Closed covers can be got from open covers by taking the closure of each of the open subsets, which incidentally makes it so that every point $x \in X$ is in the interior of one of the closed subsets C_i (a further stipulation sometimes placed on closed covers).

Example 112 Consider $\mathbb{R}$ with the usual topology. Let

$$U_n = (n, n+2), \text{ where } n \in \mathbb{Z}.$$

Then the collection $\{U_n\}_{n\in\mathbb{Z}}$ forms an open cover of $\mathbb{R}$—one where, incidentally, we will have many overlaps.

Example 113 Again consider $\mathbb{R}$. Let

$$U_n = (-n, n), \text{ where } n \in \mathbb{Z}.$$

Then the collection $\{U_n\}_{n\in\mathbb{Z}}$ forms an open cover of $\mathbb{R}$—a nested cover where we will have many more overlaps than in the previous example, since each set contains all the preceding ones.

Example 114 We can have a cover that has considerable redundancy. For instance, take the set $(0, 1)$ in $\mathbb{R}$. This set has an open cover with the collection

$$\left\{\left(0, \frac{3}{4}\right), \left(\frac{1}{4}, \frac{1}{2}\right), \left(\frac{1}{4}, 1\right)\right\}.$$

Observe, though, that we do not need all these subsets. Already with the first and last members of the cover we will have a cover of $(0, 1)$, that is, with

$$\left\{\left(0, \frac{3}{4}\right), \left(\frac{1}{4}, 1\right)\right\}.$$

In such case, we speak of the latter cover as a *subcover*.

As such examples can help one see, be sure to observe how, in general, a subcover consists of *fewer* open sets, not *smaller* subsets.

Example 115 Take $S = \mathbb{R}$. The family $\{U_n\}$ of open intervals

$$U_n = (n-1, n+1), \text{ where } n \in \mathbb{Z}$$

forms an open cover of $\mathbb{R}$ that contains no nontrivial subcovers.

Example 116 Suppose our space is $\mathbb{R}$ and $S = (0, 1)$. Then, the collection

$$\left\{\left(-1, \frac{1}{2}\right), \left(0, \frac{3}{2}\right)\right\}$$

is an open cover of S. This is an example of a *finite cover*, as it consists of only a finite number of sets (in this case, just two).

Example 117 (Nonexample) A *Cantor set* is constructed by iteratively deleting the open middle third from a set of line segments. For instance, one can first delete the open middle third $\left(\frac{1}{3}, \frac{2}{3}\right)$ from the interval $[0, 1]$, giving us the two line segments $\left[0, \frac{1}{3}\right] \cup \left[\frac{2}{3}, 1\right]$. One then deletes the open middle third of each of those remaining segments, leaving us with

$$\left[0, \frac{1}{9}\right] \cup \left[\frac{2}{9}, \frac{1}{3}\right] \cup \left[\frac{2}{3}, \frac{7}{9}\right] \cup \left[\frac{8}{9}, 1\right].$$

One can continue indefinitely with this process. The Cantor ternary set then contains all points of the interval $[0, 1]$ that are *not* deleted at any step in this infinite process.

Such a set cannot cover $[0, 1]$, and Cantor sets in general will fail to cover $\mathbb{R}$.

Exercise 14 Given X a topological space, is $\{X\}$ an open cover of X?

Solution Yes, a rather trivial one, but it is indeed a cover—specifically, it is the coarsest cover of the space. Obviously, this fact would apply to *any* open subset as well: namely, for an open subset U of a space X, the coarsest cover of U is just the cover given by one element $\{U\}$.

There are generally going to be many different ways to cover a given set, that is, there may exist many open covers for subsets of a topological space. For instance, still working with $S = (0, 1)$, consider the following.

Example 118 Let $U_n = \left(-\frac{n}{3}, \frac{n}{3}\right)$. Then

$$\bigcup_{n\in\mathbb{N}} U_n$$

clearly contains $(0, 1)$, making it an open cover of S. But observe that the subcover that contains only the first three sets U_1, U_2, U_3 already covers $(0, 1)$—and so too would any larger subcover of the original cover! This means that the original open cover has a finite subcover that consists of only three open sets. As such, this makes the original cover a rather inefficient one.

As it turns out, not every cover has a finite subcover, as one can see in a number of ways. For instance, consider the set

$$S = (-\infty, -1] \cup [1, \infty).$$

Then S is a closed subset of $\mathbb{R}$, as the complement of S is $(-1, 1)$, which is an open set of $\mathbb{R}$. The collection

$$\{(n-1, n+2) \mid n \in \mathbb{Z}\}$$

is an open cover of S. Yet it has no finite subcover.

Exercise 15 Take $U_n = \left(\frac{1}{n}, 1\right)$, an open cover of $(0, 1)$. Does this cover have a finite subcover?

Solution It is straightforward to show that

$$\bigcup_{n\in\mathbb{N}} \left(\frac{1}{n}, 1\right) = (0, 1).$$

However, we can also show that no finite collection of the $\{U_n\}_{n\in\mathbb{N}}$ can act as a cover of $(0, 1)$.

Those sets for which every cover has a finite subcover are rather special, and so are given a special name, namely *compact*. Compactness is a very important topological property, one that plays a central role in much of topology. Open covers are not terribly interesting on their own, without further properties or without engaging more advanced constructions or properties. It is generally fairly simple to find an open cover of a set. Properties like compactness, by contrast, are less trivial because they involve saying something about *every* open cover of a set. Metric spaces, for their part, are rather significant spaces, in part because of the fact that they are *paracompact*—where this is a property that can be defined in terms of open covers, and appears to account for certain of the "nicer" properties of such spaces.

One further notion concerning covers that will be especially useful to us in the subsequent story we tell of sheaves is that of *refinement of covers*. Subcovers are an example of this, as every subcover of a cover will be a refinement of that cover. As we advance in the development of ideas, and especially as we arrive at the sheaf notion, it will often be natural to ask the question: If a certain property holds for some cover, for what other covers can we expect it to hold?

As we already began to appreciate, there are potentially many different covers of a set. When we formally introduce sheaves on topological spaces in the next chapter, we will see that this sheaf notion effectively just combines the data of a functor (presheaf) with the notion of a cover. In the definition of a sheaf, the sheaf axioms will be required to hold for all covers. But we will also see that if F is a sheaf for some cover $\mathcal{V}$, then it is a sheaf for every cover that $\mathcal{V}$ refines. Since we will have to verify the sheaf axioms on spaces with covers, whenever such spaces have a finest cover, verifying that the sheaf axioms hold on such a space will amount to just checking it on the finest cover, since this will guarantee it for all covers by the result we just mentioned. The following notion will thus be rather useful in the formulation of such things.

Definition 119 Suppose we have two covers $\mathcal{U}$ and $\mathcal{V}$ of a subset S. We say that $\mathcal{V}$ *refines* (or is a *refinement* of) the cover $\mathcal{U}$ if for each $V \in \mathcal{V}$ there is a $U \in \mathcal{U}$ such that $V \subseteq U$.

Exercise 16 Which covers, if any, are refinements of the trivial cover $\{U\}$ of an open subset U?

Solution This one should be easy. Without knowing anything else, you should have been able to say that *every cover* of U will refine the trivial cover $\{U\}$.

We can note that the collection of all covers of a set U in fact forms a category $\mathbf{Cov}(U)$: its objects are covers and its morphisms are supplied by the refinement relation, that is, there is a unique morphism $\mathcal{V} \to \mathcal{U}$ whenever $\mathcal{V}$ refines $\mathcal{U}$.

There is a final notion related to covers and their refinements that we will occasionally make use of: that of the *nerve* of a cover. The nerve of a cover gives us something like a combinatorially friendly representation of the data of a cover, one where we can effectively forget about points in the space and instead get a simplified "approximation" of the space by working instead with a structure that represents the abstract relations between elements of a cover. While the notion of a nerve can be useful for formulating the sheaf definition, we postpone its formal definition, and further discussion, until chapter 9.

4.3.3 Category of Topological Spaces

So much of topology is driven by the need to study continuous functions. From an early age, we become accustomed to the idea of a function and functional dependence, where change in one variable quantity relates to change in another variable quantity, thus establishing a functional dependence between one variable (input) and another (output). Such dependence can unfold in a number of general ways, of course, but we typically learn to think of the property of continuity of functions in terms of

> a big change in the output implies that there must have been a big change in the input.

The standard ϵ-δ definition of continuity one learns in calculus is one way to get a better handle on the implied question of "how big?" There, of course, we say that

> a change in the output greater than ϵ implies a change in the input greater than δ,

or,

> if the change in input is bounded by δ, then the change in output is bounded by ϵ.

More explicitly, a function $f:\mathbb{R}\to\mathbb{R}$ is defined as continuous provided for every $x\in\mathbb{R}$ and every real number $\epsilon>0$, another $\delta>0$ can be found such that

$$|f(x)-f(y)|<\epsilon \text{ for every } x\in\mathbb{R} \text{ with } |x-y|<\delta.$$

Regardless of which maximal error $\epsilon>0$ is required, the idea is that there is always an interval around x—which is $(x-\delta, x+\delta)$, with size $\delta>0$—where all approximated function values $f(y)$ deviate by less than ϵ from the function value $f(x)$ being approximated.

In the context of the real line or plane, where we first learn these notions, we think about it in terms of an implied metric, so that continuity involves something like a control problem: we can control the error $f(x)$ to be lower than ϵ by keeping the error in the argument sufficiently small, that is, smaller than $\delta>0$. If you are measuring some x, using it to compute $f(x)$ where f is a continuous function, this ϵ-δ criterion allows you to find the maximal error δ in x (i.e., $|y-x|<\delta$), which guarantees that the final error $|f(y)-f(x)|$ will be smaller than ϵ. Such a δ may be found only if small changes around the argument x also determine small changes around the function value $f(x)$. Thus, for functions continuous at x, we must have

$$y\approx x \implies f(y)\approx f(x),$$

which just says that whenever y is sufficiently close to our point of interest x, then $f(y)$ will be approximately $f(x)$. Such an idea can of course be described using the notion of an ϵ-neighborhood: for every ϵ-neighborhood $(f(x)-\epsilon, f(x)+\epsilon)$ around $f(x)$, there is always a δ-neighborhood $(x-\delta, x+\delta)$ around x, whose function values are all mapped into the ϵ-neighborhood.

As we ascend beyond the narrower context of Euclidean spaces, and even ultimately beyond metric spaces, we get a more general treatment of the notion of continuity as well.

Definition 120 A function $f:X\to Y$ between two topological spaces X, Y, is *continuous* iff $f^{-1}(U)$ is open in X for every subset U open in Y.

This shows how we can determine whether or not a function is continuous without using any information about a metric. We need only know which subsets of X and Y are declared open. Notice that we could just as well have defined continuity using closed sets, as $f^{-1}(Y\setminus A)=\{x\in X \mid f(x)\notin A\}=X\setminus f^{-1}(A)$.

Now, one can easily verify that for any topological space X, the identity map $\mathrm{id}_X: X\to X$ is continuous; that for any topological spaces X, Y, Z, and any continuous functions $f:X\to Y$ and $g:Y\to Z$, the composition

$$g\circ f: X\to Z$$

is itself continuous; and this composition will be associative.

Altogether, this informs us that topological spaces, together with continuous functions, form a category, one we call **Top**. More explicitly, the category **Top** has topological spaces for objects, and continuous functions for morphisms. In more detail: we have seen that a *topological space* is a pair $(X, \mathscr{O}(X))$—usually abbreviated by the carrier set X—where X is a set and $\mathscr{O}(X)$ are the open sets of the topology on X. Described in a more categorical way, the morphisms here are just functions $f : X \to Y$ such that for every $V \in \mathscr{O}(Y)$, the preimage $f^{-1}(V)$ in the order of all subsets of X is in $\mathscr{O}(X)$, that is, so that there exists an arrow along the top of the following square making the diagram commute (where the vertical arrows are just inclusions):

$$\begin{array}{ccc} \mathscr{O}(Y) & \longrightarrow & \mathscr{O}(X) \\ \downarrow & & \downarrow \\ \mathbb{P}(Y) & \xrightarrow[f^{-1}]{} & \mathbb{P}(X). \end{array}$$

In other words, a continuous map $f : X \to Y$ gives rise to a function $f^{-1} : \mathscr{O}(Y) \to \mathscr{O}(X)$ that carries an open subset $U \subseteq Y$ to its preimage $f^{-1}(Y)$, open in X—recapitulating the usual notion of continuity, as defined in general topology.

Notice how in the setup above, morphisms $(X, \mathscr{O}(X)) \to (Y, \mathscr{O}(Y))$ in **Top** already include a morphism between orders, just one that goes in the opposite direction $\mathscr{O}(Y) \to \mathscr{O}(X)$. In other words, there is a functor $O : \mathbf{Top}^{op} \to \mathbf{Pos}$ that takes a space X to its underlying poset $\mathscr{O}(X)$ of open subsets.

Given a topological space X, the open sets $\mathscr{O}(X)$, ordered among themselves by inclusions, forms a poset. As such, we can describe the following:

Definition 121 For a topological space X, the *category of open subsets* $\mathscr{O}(X)$ (or, if you prefer, **Open**(X)) of X is the category that has

- for objects: open subsets $U \hookrightarrow X$ of X; and
- for morphisms: inclusions $V \hookrightarrow U$ of open subsets $V, U \subseteq X$.

Given a topological space, this category of open subsets of that space will prove to be a category of great interest to us, as we can examine data assigned to a space in terms of what it does to the open subsets.

When we think about covers in terms of $\mathscr{O}(X)$, the poset of open subsets of X, ordered by inclusion, we are saying that an I-indexed family of open subsets $V_i \hookrightarrow U$ *covers* U provided the full diagram consisting of the sets V_i together with the inclusions of all their pairwise intersections

$$V_i \hookleftarrow V_i \cap V_j \hookrightarrow V_j$$

has U for its colimit. Roughly, one can think of a covering of a given object U as some sort of *decomposition* of that object into simpler ones, the resulting simpler pieces of which, when taken altogether, can be used to recompose all of U. At the outset, it is reasonable to just think of this in terms of specifying a collection of subregions that can be laid over a given region in such a way that the entire region is thereby covered, where an entirely obvious but still decisive observation is that such subregions making up the cover can *overlap* one another. The naive image to keep in mind is that we have a region U that we want to cover with some collection of pieces into which it may be regarded as being decomposed. Suppose we have some $V_1 \subseteq U$ and $V_2 \subseteq U$:

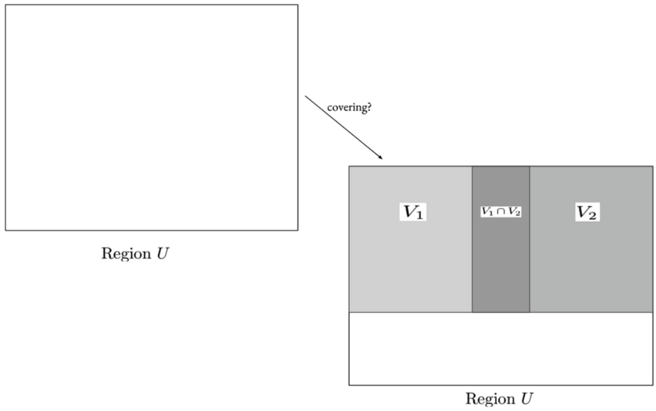

Clearly, V_1 and V_2 collectively fail to cover U, yet we can observe that there is a subregion where V_1 and V_2 overlap, which we call $V_1 \cap V_2$ and regard as specifying another "piece." Since V_1 and V_2 collectively cover more of U than either does individually, we should also consider the larger region (the entire northern half of U) that results from joining V_1 and V_2. We might continue in this manner, working our way up to a collection of subregions of U that actually cover all of U. For instance, we might have another V_3, laid on top of the entire southern half of the region (and partly overlapping with each of V_1 and V_2), such that the entire region U is now covered by the collection $\{V_1, V_2, V_3\}$. Altogether, the data of such a system of open sets, ordered by subset inclusion, will have the structure of a poset (this means, in particular, that we can regard $\mathcal{O}(X)$ as a category). In our particular case, this could be displayed by the diagram:

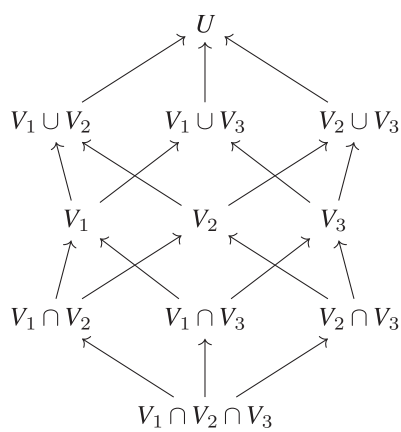

$\mathcal{O}(X)$

revealing the components of the space, together with their relevant inclusion relationships as members of a cover of the entire space.

Sheaves on a topological space can be described as particular presheaves on the open subsets $\mathcal{O}(X)$, presheaves that satisfy a further property. What ultimately will distinguish presheaves and sheaves is that sheaves are a special kind of presheaf—one that is “sensitive to” the information or structure of a cover.

In some sense, the sheaf notion is one that unfolds in an analogous fashion to something we see in analysis and in the context of metric spaces. In such settings, one first learns how to think about the continuity of a function f in terms of one that commutes with limits—where the limits come from sequences of points $\{a_n\}_{n=1}^{\infty}$ converging to a point a—in the sense that $\lim_{n\to\infty} f(a_n) = f\left(\lim_{n\to\infty} a_n\right) = f(a)$, and where $f(a)$ is ultimately independent of the sequence chosen to approximate a. At a much higher level of abstraction, the notion of a sheaf almost seems to amount to a retelling of this story. We saw in the last chapter how functors that commute with (the categorical version of) limits are said to be continuous. Letting the relevant limits now come from open covers, on this general categorical version a sheaf is fundamentally just a functor that commutes with limits—and, as such, we appear to have described something like a purified version of the continuity we first came to know in studying continuous functions on the real line and plane. The ingredients we will need to develop this notion of a sheaf will just be that of a functor (presheaf) and the notion of a cover of a space, so we have everything we need to jump right in.

4.4 Philosophical Pass: Open Questions

Box 4.1
Questions Concerning Elements of Topology

In the usual presentations of general topology, one notable oversight is that the characteristic conditions constituting a topology are almost never justified or properly clarified. One is left wondering why these conditions are so important, why they are what they are. Moreover, one is told that open and closed sets are basically formally dual to one another—yet nearly all treatments of topology go on to work almost exclusively with open sets, treating them as somehow primitive or indirectly alluding to their specialness. Finally, in the standard accounts, the true scope or activity of topology seems to be artificially limited, or misunderstood. In order to help rectify these things, the appendix takes up and addresses the following questions:

1. **Why are the axioms of a general topology what they are?** The standard account informs us that notions of general topology and the axioms of a topology itself arise from an abstraction from Euclidean space to metric spaces, followed by a further abstraction of certain properties of the open sets as definitive for a topology. While the Euclidean space motivation is easy enough to understand on its own, this often masks *why* the properties abstracted in the second step are the ones that survive the abstraction and come to constitute a general topology, arguably one of the most powerful notions in all of mathematics. One would like a clarification of these axioms and some account of why they are what they are. If the notion of a topological space really is some sort of generalization of features found in metric spaces, simply referring back to the example or instance of particular metric spaces, and the appearance of such features therein, cannot do anything to clarify why the axioms of a topology are what they are.

   Suppose some concept X, defined by a collection of properties T, is a generalization of concept Y, for which the same properties can be, and were first, observed to hold as a matter of fact. As a generalization, this means in particular that Y is an instance of X, yet concept X captures or describes phenomena or instantiations that are not Y. As such, an *explanation or justification* of why X is not only governed, but even characterized, by properties T instead of others, cannot be given by appealing to the simple *fact* that such properties can also be found governing instance Y.

2. **Why opens?** Nearly every modern text on general topology appears to give some sort of precedence to open sets, treating them as the primitives, even while assuring us that we could equally well have used closed sets, simply by appropriately dualizing the open set account. The status quo perspective on such matters seems to have little problem equivocating between taking open sets as primitive (presumably for some sensible, though never forthcoming, reason), on the one hand, and assuring us that "it's all the same, purely a matter of convention whether we use opens or closed sets" (which seems at times to involve a fundamental misunderstanding of what is and what is not entailed by such formal duality).

   This equivocation has historical roots. Historically, the first of the topological ideas to arise, deeply rooted as they were in problems in analysis, was that of the limit point of a set, used in formulating the notion of a closed set, which arose shortly thereafter. The very first formulation of what is now universally called a topological space, given by Kuratowski, is given in terms of the primitive operation of the closure of a set. Surprisingly, at least given its current status, the idea of an open set emerged last, and mathematicians appear to have been quite fine without it for longer than you might assume. (Moore (2008) has some useful history on these matters—especially on how many prominent mathematicians of the nineteenth century seem to have had little use for "open sets," and instead were initially, and for some time, apparently far more interested in defining and working with boundary points, limit points, and closed sets. The Sierpińksi quote below is taken from this paper.)

   It was not until the 1920s that definitions of a topology in terms of open sets even appeared. In his 1928 book *Introduction to General Topology*, Sierpiński paved the way for taking the notion of "open set" to be primitive, offering an axiomatic definition in terms of open sets that was close to the standard definition now given for a topological space. In the preface he wrote:

   > The axiomatic development based on the concept of an open set (as a basic concept) seemed to us simpler and more intuitive than other axiomatic treatments which will be mentioned. (Sierpiński 1934, iii)

   Later textbooks followed suit, usually with little more of substance than the fairly unconvincing justification given above.

   In short, we would like to know if there really are any mathematical—or even just more persuasive philosophical—reasons for treating open sets as primitive, as somehow more desirable to work with than their closed set counterparts. Or perhaps this is all confused and we need to do things differently.

3. **What is topology really about?** The standard story lends itself to a not unreasonable story about what we are doing when we are doing topology. While plausible on its own terms, it does seem that topology is about something slightly different, and more general, than the usual account would have us believe.

This is a section of doi:10.7551/mitpress/12581.001.0001

Sheaf Theory through Examples

By: Daniel Rosiak

Citation:

Sheaf Theory through Examples
By: Daniel Rosiak
DOI: 10.7551/mitpress/12581.001.0001
ISBN (electronic): 9780262370424
Publisher: The MIT Press
Published: 2022

The open access edition of this book was made possible by generous funding and support from Arcadia – a charitable fund of Lisbet Rausing and Peter Baldwin, and MIT Press Direct to Open

The MIT Press

The open access edition of this book was made possible by generous funding from Arcadia—a charitable fund of Lisbet Rausing and Peter Baldwin.

ARCADIA
A charitable fund of Lisbet Rausing and Peter Baldwin

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in LaTeX by the author.

Library of Congress Cataloging-in-Publication Data

Names: Rosiak, Daniel, author.
Title: Sheaf theory through examples / Daniel Rosiak.
Description: Cambridge, Massachusetts : The MIT Press, [2022] | Includes bibliographical references and index.
Identifiers: LCCN 2021058949 | ISBN 9780262542159 (paperback)
Subjects: LCSH: Sheaf theory.
Classification: LCC QA612.36 .R67 2022 | DDC 514/.224–dc23/eng20220521
LC record available at https://lccn.loc.gov/2021058949