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

A Appendix (Revisiting Topology)

In which we continue the thread from chapter 4, introducing the reader to certain connections between (modal) logic and general topology, using this to tell one possible story that can address the three questions raised at the end of chapter 4, and finally where we briefly consider how this story might relate to the bigger story we have been telling, throughout the book, about the rethinking of space.

Returning to the realm of general (point-set) topology: recall how, as we ascended in generality, leaving the notion of distance behind, the resulting notions of the open (closed) sets of a topology really just capture requirements on the interrelations of the subsets of a given set, requirements essentially governing how new members of a collection of subsets can be built from old ones (through intersection and union). In the usual treatment of these matters in general topology, the defining axioms—stability under finite intersection and stability under arbitrary union—are often left somewhat inscrutable. The aim of this appendix is to begin to lift the veil on such matters by introducing the reader to some notable connections with logic, and using this to tell a story that helps us get a better handle on the "meaning" of the axioms.

To get a better sense of what we are doing here, briefly consider the history of aviation. Consider how, for centuries, early aeronautical engineers attempted to build airplanes by looking to birds and bats, mimicking their flapping wings, general shape, and way of flying. By closely observing such creatures, they came to believe that the secret to manned flight was going to be found in devising complex machines that closely mimicked the flapping-wing techniques of birds. On the whole, these attempts failed, and airplanes achieving sustained flight were long deemed impossible. As an improved understanding of the general principles of aerodynamics was gradually attained, it would occur to people that the apparently essential feature of a bird's wings—the flapping of its wings—could be abandoned. Of course, by dropping this feature, and constructing fixed-wing aircraft, sustained long-range load-carrying flight became possible. The rest is history.

Similarly, it was realized by early "explorers" in analysis that we can drop the apparently essential notion of distance and retain a workable notion of open sets—in fact, achieving an even more powerful notion, capable of accommodating new constructions and broadening our treatment of continuity. But just as a fixed wing is still a wing despite not flapping, the introduction of open sets that no longer deploy a notion of distance—and the attendant general definition of continuity—raises a number of questions concerning such a construction and its relation to the old approach. Following this idea, the three questions we introduced in chapter 4 present themselves. To continue the analogy, such questions remind us of how, despite all the advances in aerodynamics, there apparently remains little agreement on what generates the aerodynamic force known as lift, that is, what keeps things in the air![258] Likewise, there is a surprising lack of explicit awareness or agreement among mathematicians regarding why the defining features of a (general) topology are what they are. Most mathematicians will not hesitate to point to the metric space setting as motivation for the more general notion of a

258. See, for instance, Regis (2021).

topology; and they will be able to point to instances exhibiting the power of the more general notion and show off its utility in capturing nonmetrizable spaces; but it is rare to find any general story from topologists about what is so special about the particular axioms used to define a topology: namely, that the members of the collection should be stable under arbitrary union and under finite intersection.

The aim of what follows is to explore some aspects of an account, and connections to matters beyond topology, that might help clarify why the decisive defining properties of a topology are what they are. Exploration of this question will get at a number of related issues, such as why there is this size asymmetry, what this may reveal about what general topology is about, and why open sets typically seem to be preferred to closed sets (and whether this is even conceptually justifiable or just a historical accident). Answers to these questions appear to have something to do with characterizations and assumptions about the nature of continuity, but one can also argue that they extend to deep connections with logic, models of observations, and the structure of verification, and what these connections mean for the treatment of points and boundaries.

The appendix includes five sections. Section A.1 informally motivates the interpretation of general topological notions in terms of certain logical notions and the nature of verification. Section A.2 makes precise the ideas of the preceding section, opening onto a presentation of modal logic and grounding a deep connection between a certain modal logic and topology. Section A.3 steps back and discusses the ideas behind such a connection, fleshing out some of the philosophical implications both for how we see topology and how we think about certain modalities. In this section, we will also unearth a potential response to the first of the three questions raised in the present chapter. Section A.4 takes up the question of "why opens?" and offers some observations and conjectures concerning the precedence given to them in the modern treatment of topology; we also take the opportunity to discuss some potential (generally ignored) advantages of working with closed sets, especially in relation to sheaves. The appendix concludes, in section A.5, with a return to the third question first raised in the Philosophical Pass of chapter 4—*What is general topology really about?*—and briefly relates such ideas to the broader conceptual advances (regarding space and "topologies") we have seen throughout this book.

A.1 Conceptual Motivation: Topology as Logic of Finite Observations

While deceptively obvious at first glance, it turns out to be a surprisingly deep observation that

> topological reasoning is intimately bound up with reasoning with approximations (or under conditions of error tolerance).

Let us spell this out a bit, using a common illustration.[259] Suppose you are a traffic officer charged with observing cases of speeding vehicles, and that you have radar guns—always with some fixed error tolerance—for precisely this purpose. The radars clock speeds, so as long as a car is moving, it is observable by a radar. In other words, we are imagining that there is some overall field of observation,

the nonnegative real line, while various subsets of this represent possible, more specific observations, which can be captured by our measuring devices. Your department has reserved a special name for the portion of this field of observation involving speeds greater than 50 mph—these have the observable property of *speeding*.

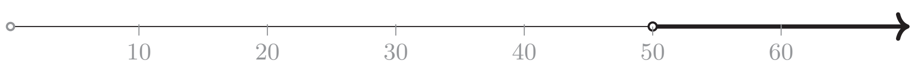

In other words, the property *speeding* can be thought of as being identified with $(50, \infty)$.

Now suppose a car is going 51 mph.

259. For instance, Moss and Parikh (1992) make use of this same metaphor of radars.

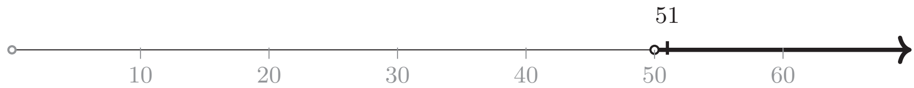

"Clearly," the car going 51 mph is speeding, as 51 "clearly" is in the interval $(50, \infty)$. But, as we already insisted in the dialogue in chapter 4, everything depends on our account of what this "clearly" actually involves. Let us first make a relevant observation. Surely, if 51 is indeed a case of speeding, since such events are observable, we should be able to produce a *verification* (or *piece of evidence*) that tells us unequivocally that what we are observing is indeed a case of speeding. As our radars supply us with our windows of observation, perhaps certain of these can give us the verification we are looking for. Let us consider various cases of such a possible verification.

Case 1: Suppose we witness 51 mph with a measuring device or radar with an error of ± 2, which is equivalent to adopting a somewhat restricted observational window of $(49, 53)$.

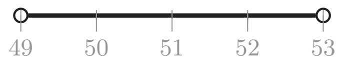

Using such a radar as our measuring device meant to serve as verification of the property in question, we get back only an equivocal answer about whether or not 51 ought to count as having that property: it can really only tell you that there is both speeding and not speeding going on. As potential evidence, the radar just tells us that the reading of the radar is 51. But we want to assert something stronger, namely that "there is speeding." After all, it may be that this radar shows 51 but the car is really going 49.5 mph. Such a measurement device cannot serve as evidence verifying that 51 is a case of speeding.

Case 2: Suppose our department gets more resources and allocates some of those to acquiring more refined radars, with an error of just ± 0.5 mph. Now, as we use such a radar and it shows 51, we have that for the resulting observational window $(50.5, 51.5)$,

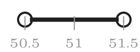

there will only be speeding *as far as the radar can see*! Giving us a read-out of 51, such a radar indeed serves as verification of speeding: for, directed at 51, it does not witness anything but cases of speeding. After all, unlike in the previous case, when pointed at 51, there is nothing that this radar can even see that is not speeding, so the entire radar cannot but verify that 51 really is speeding. In a moment, we will more closely consider the general features of this device; but before doing so, let us consider one last case.

Case 3: But why deal with such approximate regions? Why not just zoom in until we were *at the point itself*, on the nose? Provided we allow that this is even something we could really do (instead of an impossibility that would require our department had infinite resources), our resulting observational windows would then collapse down to the point itself, and by construction, operating with this degenerate window of observation, there would be nothing around, nothing *in the vicinity*, that could inform about how it related to its environs—in particular, whether or not 51 was in $(50, \infty)$. Informally, to verify that you are in (or not in) some interval, such as $(50, \infty)$, you need to be able to "look around"—that is, you need the presence of others (in or not in the interval in question) to help inform you about where you are. If we really could find such an exact measuring device, it could not tell us that 51 was speeding (or not), since there is nothing for it to see except the point itself—and, remember, we are asking about evidence for 51 *being in some interval* (or having some observable property). The basic idea here, in short, is that

> *verifiable knowledge* of an actual location/speed belonging in a region of observation is determined by observing the relationship of the region to environments/neighborhoods of the actual location/speed.

Moreover, as our discussion of the previous cases already suggested, if an observation made by a measuring device of some s is to serve as evidence, or verify a property, then anything that can be said of s must hold not just of s, but also of any other s' within the observational window of that device. The present point is: a hypothetical exact radar would not constitute a verification of anything—for it cannot speak to the relationship of what it purports to observe (the actual speed/point) with any other observable.

Of the three cases presented above, only case 2 describes something that could verify that 51 constituted speeding (i.e., was *in* $(50, \infty)$). Let us thus look at case 2 more closely. First, observe that, while the radar with an error of ± 0.5 mph worked just fine—in a way that the radar from case 1 (with an error of ± 2 mph) did not—there is in fact nothing special about the number 0.5 mph. A number of other radars would have served the same purpose: for instance, any radar of error smaller than ± 0.5 mph; in fact, so too would any radar with error $\pm 1 - \epsilon$ for any $0 \leq \epsilon < 1$.

Now suppose you have an old radar R_1 that can approximate the speed up to an error of ± 1, so that we can think of this radar as providing evidence of the actual case of 51 mph being a case of speeding, on account of its observation of $R_1 := (50, 52)$. But suppose you have another newer more expensive radar R_2 that observes cars with an error of only ± 0.5 mph, so that you can also supply evidence of speeding in $R_2 := (50.5, 51.5)$. This refined instrument R_2 will be able to prove its value in being able to observe instances of speeding that the first radar R_1 could not, for example, when a car is clocked at 50.75 mph. In other words, the two radars do not verify exactly the same things. However, there is an important asymmetrical relationship between what they can verify:

if R_1 can verify an instance of speeding, then R_2 can as well;

but

if R_2 can verify an instance of speeding, R_1 need not be capable of verifying it.

The latter fact describes how R_2 is a *better approximation* (and so represents a possible increase in knowledge). On the general way of looking at things that we have started to motivate, subsets are regarded as the *possible observations*, so that "open sets" containing a point are effectively just pieces of evidence *qua* observable properties concerning actual states (represented by a point or definite speed). We can further regard the key relation between the various subsets of a space as follows:

$x \in V \subseteq U$ says that V is a better approximation to x.

If you think of the open sets in terms of generalized rulers or measuring devices, the idea is that smaller rulers $V \subseteq U$ give you more refined measurements. One might also begin to think of such better approximation sets in terms of *effort*—the more refined the set, the more effort. The interpretation in terms of effort is meant to represent an action—such as a measurement, computation, or approximation—that in general may result in an increase in knowledge or what is knowable. If one thinks of the open sets as involving possible observations, the idea is that for $V \subseteq U$, V is a more sophisticated or refined means of observing properties. As such, it will in principle "cost more effort" compared to any cruder observation. If a radar is attempting to determine whether or not a car is speeding in a 50 mph speed zone, but the accuracy of the radar is ± 2, then a car going 51 mph cannot be verified as speeding; on the other hand, supposing a more refined traffic radar has an accuracy of ± 0.5 mph, then a car going 51 mph can be verified as speeding. The second, more refined radar—and so the one that surely costs more effort to make—can detect a property, namely that of "speeding," which the first radar cannot. That it can detect properties that "fly under the radar" (pun intended!) of cruder radars justifies our thinking of smaller sets in terms of "better approximations."

However, note that whenever a given radar detects a property, any more refined or costly radar must also detect that property. This last sentence gives a way of thinking about the key feature of "openness," and appeals to the key idea of the

Paradigm of Truth Continuity: an observable property will be true of a point whenever a verification (approximation) can be supplied, such that this continues to be true in all better approximations.

Such truth would of course break down provided there was at least one more refined approximation for which the property ceased to be true. On this interpretation that we are starting to suggest, there is a very natural way of understanding the full force of the finite intersection and arbitrary union conditions that define a topology. For reasons that will become apparent in a moment, we will describe the paradigm above by saying

$\Box\phi$ is true at $p \in U$ provided for all refinements $V \subseteq U$ of the observation U, ϕ is still true at $p \in V$.

For now, you might think of $\Box$ in terms of supplying a sort of radar for some observable property—the old radar guns used to have rectangular windows! The reader might want to contemplate on their own some other sensible conditions one might expect to obtain for this $\Box$.

A.2 Explicit Connections to Modal Logic

Recall the discussion of modal logic from chapter 7.2. The first modern formal analysis of modalities is typically attributed to the work of C. I. Lewis, beginning in 1912 and culminating in a paper in 1932.[260] Lewis constructed a series of axiomatic systems of modal logic—called $S1$ through $S5$—based on which of the axioms were held to legislate over the modalities. Lewis was first motivated by concerns over what he took to be an eliding of an important difference in how implication was understood by the algebraists and logicians of his time, on the one hand, and the ordinary meanings of implication, on the other. He thought that these concerns could be met by using an *impossibility* ($\neg\Diamond$) operator to define a more appropriate notion of implication, and accordingly he came to introduce systems taking as primitive the connectives negation, conjunction, and possibility ($\Diamond$).

Gödel would later[261] advance this by initiating a simplified presentation of modal logics: here, modal logic emerges as just an extension of propositional logic—the modal operator $\Box$, or $\Diamond$, is added as an additional connective to propositional logic, together with additional axioms and rules governing its behavior. In that very brief one-page article, Gödel happened to be concerned with formalizing assertions of *provability*, which he sought to capture by means of a propositional connective B (for *beweisbar*), so that $B\alpha$ was to be read as "α is provable" or "it is provable that α." In that paper, he revealed how to define a system given by the axioms and rules of the usual propositional logic together with the following axioms governing his new operator B:

1. $B\alpha \to \alpha$
2. $B(\alpha \to \beta) \to (B\alpha \to B\beta)$
3. $B\alpha \to BB\alpha$

together with the inference rule

from α, infer $B\alpha$.

Gödel then stated that this system is none other than what Lewis had described as the system $S4$, after translating B to $\Box$. He also stated that the appropriate translation of the law of the excluded middle is not derivable, making the further suggestion that the translation of any theorem of the intuitionistic propositional calculus of Heyting is derivable in his system, and conjectured that the converse was true as well. This is of some significance, as we shall see.

These days, it is common to follow Gödel's general approach of presenting modal logic as an extension of propositional logic. While the particular axioms and inference rules adopted do depend on the context of use, or on which properties we expect our modal connectives to have, the main axioms used in modal logic, built on top of a complete axiomatization of propositional logic together with the rules of inference modus ponens, start by adjoining the basic necessitation inference rule

(Necessitation rule **N**) $\frac{\phi}{\Box\phi}$,

260. See, for instance, Lewis (1912).
261. Gödel (1933).

or $\phi \vdash \Box\phi$, where the turnstile symbol $\vdash$ here just tells us that whatever follows it represents either an axiom or a formula that can be derived as a theorem. Thus, this just says that if ϕ can be inferred (is provable) from no assumptions (beyond the axioms of logic), then ϕ is necessary in the modal logic—that is, if ϕ is a theorem (just derivable from nonmodal good old-fashioned logic), then $\Box\phi$ is also a theorem. This can also be construed as an *axiom*, in which case it reads

(**N** as an axiom) $\vdash \Box\top$.

The weakest (or most minimal) modal logic that has proven to be useful is called K, which involves an augmentation of propositional logic with $\Box$ and the rule **N**, together with the further axiom

(Axiom **K**) $\Box(\phi \to \psi) \to (\Box\phi \to \Box\psi)$.

The axiom systems of many other commonly used modal logics are extensions of K: by adding further axioms, we get other well-known modal systems. K on its own assumes very little, and so within K we cannot even prove that "if $\Box\phi$ is true, then ϕ is true." To address this, we could add the further axiom

(Axiom **T**) $\Box\phi \to \phi$,

an axiom that indeed holds in most modal logics.[262] Other elementary axioms are

(Axiom **4**) $\Box\phi \to \Box\Box\phi$,

sometimes called the "Positive Introspection" axiom, and

(Axiom **B**) $\phi \to \Box\Diamond\phi$

(Axiom **D**) $\Box\phi \to \Diamond\phi$

(Axiom **5**) $\Diamond\phi \to \Box\Diamond\phi$.

Various axioms, when conjoined, allow us to formulate other axioms, which are occasionally useful enough to be given their own name. For instance, **K** with the rule **N** can be shown to entail

(Axiom **R**) $\Box(\phi \wedge \psi) \leftrightarrow (\Box\phi \wedge \Box\psi)$.

Variously combining certain of these axioms yields distinct axiom systems corresponding to some notable modal logics (the usual names are given on the left):

- $K := \mathbf{N} + \mathbf{K}$
- $T := \mathbf{N} + \mathbf{K} + \mathbf{T}$
- $S4 := \mathbf{N} + \mathbf{K} + \mathbf{T} + \mathbf{4}$
- $S5 := \mathbf{N} + \mathbf{K} + \mathbf{T} + \mathbf{5}$ (or $\mathbf{N} + \mathbf{K} + \mathbf{T} + \mathbf{4} + \mathbf{B}$).

Observe that K through $S5$ form a nested hierarchy of systems, ultimately built on top of the minimal necessitation rule. Modal logics for which the necessitation rule holds—or which assume axiom **N**—are called *normal* modal logics. This nested hierarchy thus describes how the main normal modal logics relate. There are further, less commonly used normal modal logics, intermediate to those above, such as

- $D := \mathbf{N} + \mathbf{K} + \mathbf{D}$
- $K4 := \mathbf{N} + \mathbf{K} + \mathbf{4}$
- $B := \mathbf{N} + \mathbf{K} + \mathbf{T} + \mathbf{B}$
- $D4 := \mathbf{N} + \mathbf{K} + \mathbf{D} + \mathbf{4}$.

Calling this a "hierarchy" makes sense, as some of these logics are sublogics of others; for instance, $K4$ is a sublogic of $S4$ in the sense that all formulas valid in $K4$ are also valid in $S4$.

262. Anticipating, one could think of this as the "truthfulness of verification" principle.

Both for reasons internal to modal logic and for reasons we shall appreciate in a moment, the logic $S4$ is of special interest. An equivalent axiomatization—one that will prove to be even more revealing—of $S4$ can be given as follows:

Definition 340 The modal logic $S4$ is defined as the logic for which the following axioms hold:

- $\Box\top$ (Axiom **N**)
- $\Box\phi \to \phi$ (Axiom **T**)
- $\Box\phi \to \Box\Box\phi$ (Axiom **4**)
- $\Box(\phi \wedge \psi) \leftrightarrow (\Box\phi \wedge \Box\psi)$ (Axiom **R**)

and where modus ponens

$$\frac{\phi \to \psi \qquad \phi}{\psi}$$

and monotonicity

$$\frac{\phi \to \psi}{\Box\phi \to \Box\psi}$$

are the only rules of inference.

There is also the following derivable theorem, which is often of use:

$$\Box\phi \vee \Box\psi \leftrightarrow \Box(\Box\phi \vee \Box\psi) \qquad \textbf{(or)}.$$

While $S4$ is typically presented via **K, N, T** and **4**, it is easy to show that **Monotonicity** and **R** entail **K**; likewise, **N** and **K** can be shown to entail **Monotonicity** and **R**. Observe also that **4** together with **T** in fact gives us $\Box\phi \leftrightarrow \Box\Box\phi$.

While these axioms are of some interest in their own right, we can already anticipate where we are going with all this by making a purely formal observation, involving nothing more than pattern matching. Recall that we can express a topology, and its open sets, in terms of an interior operator. Have another look at the axioms of $S4$ given in definition 340. Now recall the Kuratowski axioms governing the interior operator, reproduced here for convenience:

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

This striking, apparently purely formal, similarity suggests that the logical connective $\Box$ might be some sort of interior operator, and that there may be some close connections between the modal logic $S4$ and topology. Indeed, $S4$, while having evolved for rather distinct purposes and in different contexts, turns out to be ordinary topology in disguise! Viewed from the other direction, this intimate connection with modal logic and features of certain logics and their interpretations can help us shed light on the defining axioms of a topology.

S4 as the Logic of Topological Spaces

The modern study of propositional logic largely got its start in the nineteenth century as algebra—in the tradition of Boole, who initiated a revolutionary conceptual shift, by combining algebra with logic. Boole had revolutionized logic, and began the project of the mathematization of logic, by applying notions and methods from symbolic algebra to the treatment of logical arguments, while simultaneously revolutionizing algebra by freeing it from its narrow application to arithmetic. By the time Lewis's work on modal logic systems appeared, many advances in the study of algebras had been made, and it wasn't long before modal systems were seen in an algebraic light. One of the overriding morals of the work of the logician Tarski—who proved a major result concerning the relation between $S4$ and topology—involves the insight that the conditions defining a topology (in terms of conditions on open sets, or in terms of an interior operator) are at bottom *algebraic*.

A prominent class of examples of Boolean algebras can be obtained, in general, as follows: let X be any nonempty set and $\mathbb{P}(X)$ the collection of all subsets of X. Then, for subsets A and B of X, if we define $\neg A$ as the set-theoretic complement of A (in X), $A \vee B$ as the union of A and B, and $A \wedge B$ as the intersection of A and B, the set $\mathbb{P}(X)$ becomes a Boolean algebra—an algebra of sets. In this spirit, after some maneuvering involving taking equivalence classes (identifying expressions that are logically equivalent), the set of all propositions of the propositional logic forms a nontrivial Boolean algebra.

Actually, a topological space itself can be represented as the Boolean algebra of all subsets of the space—in this setting, the interior operator is construed as an operator on the algebra, where it takes elements (subsets) of the algebra to other elements of the algebra. As such, a topological space can be represented as a particular algebra equipped with a particular operator—and, in this light, the Kuratowski axioms emerge as nothing other than algebraic equations. More explicitly, the equations will stipulate that the interior of the top element of the algebra is equal to the top element; the interior of any element is less than or equal to that element; the interior of the meet of two elements is equal to the meet of the interiors; and the interior of the interior of any element is equal to the interior of that element. In symbols,

- (**i1***) $I(\top) = \top$;
- (**i2***) $I(a) \leq a$;
- (**i3***) $I(I(a)) = I(a)$;
- (**i4***) $I(a \wedge b) = I(a) \wedge I(b)$.

If we substitute $\Box$ for I in the above equations, we can immediately see that we will just recover the logic $S4$; alternatively, substituting I for $\Box$ in the $S4$ axioms, we should recover the algebraic version of the topological interior axioms governing a topological space. It is intriguing that axioms first formulated to model reasoning with propositions under situations of governed by dual modalities like "possibility" and "necessity," after a very natural translation, end up recapitulating the very same axioms that mathematicians had been using to describe a space.

This connection further motivates the development of a topological semantics for modal logic. Here is one way of starting to see how this goes. Ordinary propositional calculus can be regarded as basically set theory in disguise. If we assume we start with some set X, we might attempt to translate a logical proposition of our propositional logic into a statement about sets by assigning atomic (noncompound) propositions to subsets of X, regarding $\wedge$ as $\cap$, $\vee$ as $\cup$, $\top$ as X, $\bot$ as $\emptyset$, and $\neg$ as the complement (in X). A proposition would then be regarded as true precisely when its translation always equals X, the entire set. Typically, when we think of a way of assigning truth to the various propositions of our logic, we think of interpretation (or valuation) of a formula in terms of extensions of a function from the propositional values into a "truth value" set $\{0, 1\}$. But really, an interpretation could just be seen as an assignment of meaning to the symbols that make up a formula, and the truth-values need not be confined to 0 (false) and 1 (true), but can just be some sets, where these convey information concerning the proposition. Following this lead, we can instead interpret propositions as being assigned to subsets of $\mathbb{P}(X)$, where these subsets effectively act to inform us about *where* the proposition is true or holds. We then make the appropriate substitutions for Boolean connectives as indicated, from which we can also express the implication $\rightarrow$ as $\subseteq$, and $\leftrightarrow$ as $=$.

Suppose we have a mapping $v : \Phi \rightarrow \mathbb{P}(X)$ that sends propositional variables to subsets of X, where we think of this as a general "valuation" in the algebra of subsets of X, as indicated above. In this way, for ϕ an arbitrary formula, we can let

$$v(\phi) = \{x \in X \mid \phi \text{ is true at } x\}.$$

Then, if ϕ and ψ are propositional variables, with $v(\phi) = A$ and $f(\psi) = B$, where A, B are subsets of X, by interpreting logical connectives as set operations, as indicated above, we can extend the function v just as you might expect, so that every formula is mapped to a subset of X:

- $v(\neg\phi) = X \setminus v(\phi)$;
- $v(\phi \vee \psi) = v(\phi) \cup v(\psi)$;

- $v(\phi \wedge \psi) = v(\phi) \cap v(\psi)$;
- $v(\phi \to \psi) = (X \setminus v(\phi)) \cup v(\psi)$.

Observe also that we must have $v(\top) = X$. If a formula ϕ is mapped to the entire set X by all mappings v, then we say that ϕ is *valid* in X. Moreover, observe that $\phi \to \psi$ will be valid with respect to any interpretation in X if and only if $v(\phi) \subseteq v(\psi)$ for any mapping v that sends propositional variables to subsets of X. This can be seen by noting that $\phi \to \psi$ will be valid with respect to any interpretation v in X if and only if $v(\phi \to \psi) = v(\neg\phi \vee \psi) = v(\neg\phi) \cup v(\psi) = (v(\phi))^c \cup v(\psi) = X$—and this will be true precisely when $v(\phi) \subseteq v(\psi)$, for any v.

So far, none of this is that interesting. The important bit comes with the treatment of the modal operator, which is where the topological semantics really comes into force (and derives its name). We have already anticipated how the logical connective $\Box$ might plausibly be interpreted as follows: for any mapping v that maps propositional variables to subsets of X and for any propositional variable ϕ, we interpret $\Box$ as stipulating

$$v(\Box\phi) = \mathbf{int}(v(\phi)).$$

Supposing we have a space, then, the topological semantics proceeds as follows. We define a *topological model* on a topological space X. This is a pair (X, v), with $v : \Phi \to \mathbb{P}(X)$ a valuation that takes values in the set of all subsets of X, as before. Here, we can define the truth of any formula ϕ at a point x of a topological model $\mathcal{M} = (X, v)$ by induction—where, we use the notation $\mathcal{M}, x \Vdash \phi$ to mean "x semantically entails ϕ with respect to v" (or "ϕ is satisfied by v at $x \in X$"), and this is the case provided $x \in v(\phi)$. Where $\mathcal{O}$ is the set of open neighborhoods of x, we can then show, just as you might expect, that:

- $\mathcal{M}, x \nVdash \bot$;
- $\mathcal{M}, x \Vdash \phi \wedge \psi$ iff $x \Vdash \phi$ and $x \Vdash \psi$;
- $\mathcal{M}, x \Vdash \phi \vee \psi$ iff $x \Vdash \phi$ or $x \Vdash \psi$;
- $\mathcal{M}, x \Vdash (\phi \to \psi)$ iff $\mathcal{M} \nVdash \phi$ or $\mathcal{M}, x \Vdash \psi$;
- $\mathcal{M}, x \Vdash \neg\phi$ iff $\exists U \in \mathcal{O}$ s.t. $\forall y \in U\ y \nVdash \phi$;
- $\mathcal{M}, x \Vdash \Box\phi$ iff $\exists U \in \mathcal{O}$ s.t. $(x \in U$ and $\forall y \in U(\mathcal{M}, y \Vdash \phi))$.

In this way, we can say how

$$v(\Box\phi) = \mathbf{int}(v(\phi))$$

by noting that

$$\begin{aligned} x \in v(\Box\phi) &\text{ iff} \\ \mathcal{M}, x \Vdash \Box\phi &\text{ iff} \\ \exists U \in \tau \text{ s.t.}(x \in U \text{ and } \forall y \in U(\mathcal{M}, y \Vdash \phi)) &\text{ iff} \\ \exists U \in \tau \text{ s.t.}(x \in U \text{ and } \forall y \in U(y \in v(\phi))) &\text{ iff} \\ \exists U \in \tau \text{ s.t. } (x \in U \text{ and } U \subseteq v(\phi)) &\text{ iff} \\ & x \in \mathbf{int}(v(\phi)). \end{aligned}$$

In other words, the formula $\Box\phi$ is true throughout the interior of the set of points where ϕ is true. In the presence of such a model and associated semantics, we then say that a formula ϕ is *satisfied* by (or in) the model provided it is true throughout the entire space (i.e., $v(\phi) = X$), and that ϕ is *valid* in the space X provided it is satisfied in *every* model defined over X.

Moreover, given the modal translation, we can in fact say that

- $v(\neg\phi) = \mathbf{int}(v(\phi)^c)$;
- $v(\phi \to \psi) = \mathbf{int}(v(\phi)^c \cup v(\psi))$.

Building on this, the axioms of $S4$, once interpreted topologically, simply restate the Kuratowski conditions that a topological interior is expected to satisfy. Of course, the rule **N** just corresponds to the condition

$$\mathbf{int}(X) = X.$$

And recalling that combining axiom **K** with the rule **N** gives us our alternative axiom **R**, it is easy to see how this corresponds to the condition

$$\mathbf{int}(A) \cap \mathbf{int}(B) = \mathbf{int}(A \cap B).$$

Finally, axiom **T** just amounts to asserting that the interior of a region is a subset of that region,

$$\mathbf{int}(A) \subseteq A,$$

while axiom **4** (added to axiom **T**) states that the interior of the interior of a region is just the interior of that region,

$$\mathbf{int}(A) \subseteq \mathbf{int}(\mathbf{int}(A)).$$

That **int** is monotone in the sense that

$$A \subseteq B \Longrightarrow \mathbf{int}(A) \subseteq \mathbf{int}(B)$$

is represented by the fundamental inference rule

$$\frac{\phi \rightarrow \psi}{\Box\phi \rightarrow \Box\psi}.$$

Exploiting the striking formal similarity between Gödel's axioms for provability logic and Kuratowski's axioms for a topological space, McKinsey and Tarski proved that, under the above interior-based interpretation, the topologically valid statements are exactly those provable in the the modal system $S4$—that is, that $S4$ is the logic of topological spaces.[263] More explicitly, we can not only show that $S4$ is sound with respect to topological spaces in general, via the topological semantics, but we can also show that it is complete. To begin to appreciate this, recall how both propositional modal logic and topological spaces can be seen as algebras. In describing a valuation, then, we could have described these valuations without assuming anything about being given a topological space. Instead, given a set X, v is just a map from propositions to $\mathbb{P}(X)$. Now, we can consider some mapping

$$i : \mathbb{P}(X) \rightarrow C,$$

where $C = i(\mathbb{P}(X)) \subseteq \mathbb{P}(X)$. In other words, with the map i we are considering a choice of some collection of subsets of X.

In relation to our valuation and what to do about the valuation of $\Box\phi$, we can then say that $v(\Box\phi) = i(v(\phi))$. But what should i be? Here is the fascinating result. Given any set X, we can show that *any* interpretation of $\Box$ in X that happens to satisfy precisely the axioms of $S4$ will be such that the image of this interpretation forms a topology on X. $S4$ is sound with respect to any interpretation in any topological space; moreover, $S4$ can be shown to be complete over all topological spaces—in other words, if a formula is valid in every topological space, then it will be derivable from the logic $S4$. Supposing the map i makes $S4$ sound, we can show that the set $C = i(\mathbb{P}(X))$ is a topology on X and that it is one only on the condition that it satisfies the axioms and rules of $S4$!

In short, let us now ask the following question: Is the set $C = i(\mathbb{P}(X))$ a topology for X? What conditions—axioms and rules—must i satisfy to guarantee that $i(\mathbb{P}(X))$ is a topology on X? As it turns out, the set $C = i(\mathbb{P}(X))$ is indeed a topology, and is one precisely when it satisfies all the axioms of the modal logic $S4$ (where no proper sublogic of $S4$ can guarantee that C is a topology). The following table shows how the particular modal axioms transfer over to topological conditions – in particular, a checkmark means that the collection C does have the property in question, while an $\times$ means that C does not have that property, when governed by the particular modal axioms of that row.

263. See McKinsey and Tarski (1944).

If we drop any of the axioms of $S4$, we can see that C will not necessarily be a topology—in other

Modal axioms	Topological axioms			
	$X \in C$	$\emptyset \in C$	$X_1, X_2 \in C \Rightarrow (X_1 \cap X_2) \in C$ (finite intersection)	$\{X_i\}_{i \in I} \subseteq C \Rightarrow \bigcup_{i \in I} X_i \in C$ (arbitrary union)
$\mathscr{L}_\Box + \mathbf{N}$	✓	×	×	×
$\mathscr{L}_\Box + \mathbf{N}, \mathbf{T}, \mathbf{4}$	✓	✓	×	×
$\mathscr{L}_\Box + \mathbf{N}, \mathbf{K}$	✓	×	✓	×
$\mathscr{L}_\Box + \mathbf{N}, \mathbf{K}, \mathbf{4}$	✓	×	✓	×
$\mathscr{L}_\Box + \mathbf{N}, \mathbf{K}, \mathbf{T}$	✓	✓	✓	×
$\mathscr{L}_\Box + \mathbf{N}, \mathbf{K}, \mathbf{4}, \mathbf{T}$	✓	✓	✓	✓

words, no proper sublogic of **S4** can guarantee that C is a topology.[264] By observing the table, we can even decipher which axioms of modal logic must correspond to, or are responsible for, which of the axioms of topological spaces.

- **N** just stipulates that the whole space is open, that is, that $X \in C$;
- **T** with **N** gives that $\emptyset \in C$;
- **R** (equivalently, **N** and **K**) is just the stability under finite intersection condition; and
- **or** (derived from **N**, **K**, **4**, and **T** together) is the condition that open sets are stable under finite unions (which can be extended to arbitrary unions by using an infinitary extension of the modal language).

McKinsey and Tarski further showed that $S4$ is sound and complete with respect to particular spaces—namely, the logic of any dense-in-itself metrizable space.[265] This result implies, in particular, that $S4$ is the logic of the real line $\mathbb{R}$ (endowed with the usual topology), the rationals $\mathbb{Q}$, the Cantor space, or any Euclidean space. This basically shows how $S4$ fundamentally characterizes any dense-in-itself metric space, in particular the spaces we are most familiar with—thereby solidifying the spatial core of modal logic.

Altogether, *soundness* of $S4$ for a given topological space is more or less had for free, where soundness just means that if something is provable, then it is valid, that is,

$$\vdash_{S4} \phi \text{ implies } \Vdash_X \phi.$$

For its part, the completeness of $S4$ for a given topological space X is trickier. Completeness of $S4$ for a space X asserts that every validity in X is provable within $S4$, that is,

$$\Vdash_X \phi \text{ implies } \vdash_{S4} \phi.$$

The equivalent contrapositive of this claim of completeness can be more helpful:

$$\nvdash_{S_4} \phi \text{ implies } \nVdash_X \phi;$$

that is, if ϕ is *not* a theorem of $S4$, then it is *not* valid in X—which formulation effectively informs us that checking for completeness, in this context, amounts to checking that space for a class of refuting models. McKinsey and Tarski's result concerning how $S4$ characterizes any dense-in-itself metric space basically informs us that in any dense-in-itself metric space there will always be the resources to refute all nontheorems of $S4$.

264. Xu (2016) has a nice discussion of these matters; the idea for the above table came from that work.

265. A space is *dense-in-itself* when every point is the limit of other points in the space. The most familiar spaces of all—including any finite-dimensional Euclidean space—are such a space.

Stepping back, observe that we can also define

$$v(\Diamond\phi) = v(\neg\Box\neg\phi) = (v(\Box\neg\phi))^c = (\mathbf{int}(v(\neg\phi)))^c = (\mathbf{int}((v(\phi))^c))^c,$$

where this is just the closure of the set $v(\phi)$. Moreover, in general, we can construe the boundary as

$$\partial(\phi) = \Diamond\phi \wedge \Diamond\neg\phi.$$

In the context of modal logics, the notion of boundary generally captures something like the *contingency* of ϕ.

In connection with the topological semantics, we can also observe that for the law of excluded middle $\phi \vee \neg\phi$ to be valid in a topological space X, each closed set must also be open. A concrete counterexample to the excluded middle being valid can be given by the Sierpiński space described earlier, where $\{a\}$ is closed and $\{b\}$ is open, and where we define $v(\phi) = \{b\}$. Then, $a \nVdash_v \phi$. But since the only open neighborhood of a is the entire space $\{a, b\}$, which does indeed contain a point that satisfies ϕ, we must have $a \nVdash_v \neg\phi$. Thus, $a \nVdash_v \phi \vee \neg\phi$. This fundamentally gets at the deep connections with Heyting algebras as models of intuitionistic propositional logic.

Altogether, by interpreting each ϕ as a subset A of a topological space X (so that each propositional variable represents a region of the space, and so too does any formula), and making the appropriate substitutions, we can see that the modal logic $S4$ axioms governing $\Box$ reduplicate those of the topological interior operator. Going the other way, with the axioms of $S4$, we appear to have rediscovered the algebraic version of Kuratowski's axioms for a topological interior. This connection is not merely formal, but lets us attempt a justification of the defining features of a topology.

One of the merits for the topologist of seeing $S4$ as the logic of topological spaces—or topology in disguise—is to help ground an important verificationist interpretation, in which setting the axioms of a topology appear to achieve some clarification. This is explored in the next section.

One of the merits of the topological semantics for the modal logician, for their part, is that it does not validate

$$\neg\Box\phi \to \Box\neg\Box\phi,$$

known by philosophers under the name "negative introspection" principle, which seems to require that if one cannot verify something as true, then one must be able to verify that one cannot verify it as true. Some have argued that a good notion of *belief* is captured by *not knowing that you do not know*—so that one can define a belief operator

$$B\phi := \Diamond\Box\phi.$$

In those terms, the principle of negative introspection effectively makes it impossible to have incorrect beliefs, for one always knows whatever one believes that one knows. The principle is generally regarded as undesirable—especially applied in epistemic contexts, involving verifications and knowledge—since it seems to say that one cannot believe one knows things that one does not actually know. Rational agents are constantly believing they know things that they do not in fact know, so there are reasons for not wanting to enforce this principle. That the topological semantics does *not* validate it is encouraging.

Finally, there are merits to this connection that go beyond any benefit to one side or the other. For instance, in example 198 (chapter 7), we mentioned Kuratowski's 14-set theorem, which says that in a topological space, there can be no more than fourteen distinct sets that can be generated from a fixed set by taking closures and complements (or using the interior, closure, and complements). Why this should be true seems never to be explained—and one is left with the feeling that such a result is extremely mysterious, if not entirely arbitrary. In modal logic textbooks, one finds sequences of boxes and diamonds formed to get finite modalities—but, on account of certain logical equivalences, sequences of modal operators can often be reduced. One is sometimes asked to show that in $S4$ in particular, up to equivalence, there are only fourteen modalities. Again, there is not usually an explanation of why this should be the case. Authors in either camp rarely seem to acknowledge that these results are two ways of looking at the same result.

A.2.1 Models of Information

The results of McKinsey and Tarski laid the groundwork for further work into spatial logics. Moreover, the completeness results in particular drew the attention of epistemic logicians, leading certain logicians to use the deep connection between modal logic and topology to reflect features of $S4$ as an epistemic system back onto topology—where this led to an epistemic reevaluation of the interior semantics, in the context of which topologies came to be seen as *models for information*. The idea of regarding topological spaces as codifying information structures had itself been considered since the 1930s, leading to the use of topological spaces as models for intuitionistic languages.[266] Building on this use of topological models for broad classes of intuitionistic languages, the interior semantics sketched above has been seen as providing insight into evidential approaches to knowledge.

In the topological models for intuitionistic languages, as used in epistemic settings, elements of a given open basis are taken to be "pieces of observable evidence," while open sets of the topology are generally treated as "observable properties" that can be verified based on the observable evidence. Evidence as open sets can be grounded in thinking of a collection of observable evidence *directly* obtainable by an agent—through things such as computation, approximation, measurement—as a subbasis, so that the collection of observable evidence forms a basis for the topology. Furthermore, taking open sets as pieces of evidence is further supported by how measuring devices used to compute things like the speed of a car or height of a person are always *approximating devices*—in the sense that they do not give exact values, but are defined by a definite range of error tolerance.[267]

On this view of things, points in a topological space X are seen as (partial) states of information (about how things actually are).[268] Taking increasingly accurate measurements with less error-tolerant approximating devices gives rise to better and better approximations to the actual state. Building on this, the meaning of the fundamental logical connective as a kind of interior can be developed in terms of the notion of *verifiability*.[269] Open sets correspond to properties that are *in principle verifiable* by the agent—in short, open sets are being construed in terms of meaningful propositions, where these are just propositions whose truth is equivalent to their verifiability.

The general idea is that there are properties that are observable (directly) by an agent. These are regarded as forming an open basis for a topology. Because open sets can be expressed as those sets equal to their interior, we can see **int** of those properties as the set of states in which the properties are verifiable, suggesting that we read $\Box$ as follows:

$\Box\phi$: = "it is verifiably true that ϕ."

Another way to see $\Box$ is as expressesing the notion of "continuous truth."

Fundamentally, we are understanding "verification" here in such a way that

> to verify something is to observe something that entails it,

where this is entirely in keeping with the earlier view of a neighborhood as a kind of proof or evidence. Open neighborhoods U of an "actual state" x act as sound or truthful evidence. An actual state x is then in the interior of ϕ if and only if there exists a sound piece of evidence U that justifies ϕ. If an open U is included in a set representing a proposition ϕ, then we can say the proposition ϕ is entailed or justified by the evidence U. That agent knows ϕ (to be true) if they have a correct justification for it, where this is based on a sound piece of evidence justifying it.

As Bezhanishvili and Holliday (2019) puts it, given any set U of states of information,

> U is *verifiable*, relative to one's current information state, iff it is possible to achieve such verification of U after a finite amount of time, starting from the current information state.

266. See, for instance, Troelstra and Dalen (1988) for discussion.

267. See Baltag et al. (2019) for further elaboration on this.

268. See Dalen (2002) and Scott (1968).

269. This way of looking at things is nicely summarized in Bezhanishvili and Holliday (2019). One can find further variations of this overall idea in developments by computer scientists, for instance in domain theory (see Abramsky 1987, Vickers 1996) and work in formal epistemology.

Seen in this light, the features of the logic of verifiability, or the process of finite observation, can begin to shed some light on why the conditions of a topology are what they are. By taking the family of observable propositions as a basis, the features of the modal operator generate what we can think of as a "structure of verifiability"—and these just recapitulate, in another guise, the core defining properties of a topology. Moreover, this interpretation comes with a natural dual notion whereby closed sets can be expressed in terms of *falsifiable properties*: whenever a property is false, it is falsified by a sound piece of evidence. Finally, on this picture, the set of *boundary points* will correspond to those properties that are neither verifiable nor falsifiable.[270]

A.3 The Idea of All This

Making use of the approach sketched above, and taking advantage of the intimate connection between the interior and the modal operator $\Box$, we can attempt something like a *justification* of the axioms of the interior operator, that is, of the defining features of a topology.[271]

First of all, by taking $\mathbf{int}(U)$ to be the set of states in which U is verifiable, it is evident that $\mathbf{int}(U) \subseteq U$. Next, if it is possible to perform any finite sequence of possible verifications in a finite amount of time, then the **int** operator ought to distribute over finite intersections—where this effectively just means that observations can be refined in a finitary way. This is something like a continuity condition on our verifications. By contrast, we are *not* assuming that it is possible to perform an infinite sequence of verifications in a finite amount of time. Put otherwise, while observations can indeed be refined in a finitary manner—such that *each* observation can be refined by another—we do not assume that there generally exists a single and universal observation that could embrace *all* the observations in one go (which corresponds to not assuming that **int** distributes over arbitrary intersections). The "stability under *finite* intersections" condition then captures the idea of that agent's ability to conjoin only finitely many pieces of evidence into a single piece of evidence. An infinite conjunction of such properties would require confirming or verifying all of them, something it is not reasonable to assume is available to an agent capable of producing verifications. For the justification of the last condition on the interior operator, we can assume that we are working with a model of verification

> according to which by verifying U, we are also verifying that one has verified U, which implies that if it is possible to verify U, then it is also possible to verify that it is possible to verify U. (Bezhanishvili and Holliday 2019)

270. In section A.1, we sometimes used the language of "effort." Technically, modeling effort involves some extension of the semantics just sketched. If, following Moss and Parikh (1992), we let $(X, \mathcal{O})$ be a "subset space," where X is a nonempty set of states and $\mathcal{O}$ is a collection of subsets of X (not necessarily forming a topology, though topologies provide a particular case of this construction), then again elements of $\mathcal{O}$ are = *possible observations* or *possible observation sets*. Formulas are interpreted not just with respect to the actual state, but with respect to pairs of the form (x, U), where $x \in U \in \mathcal{O}$, and x represents the way the actual state of affairs happens to be. The *neighborhood* U with $x \in U \in \mathcal{O}$ is interpreted as a truthful observation that can be made about the actual state x. Moss and Parikh give a subset space semantics as follows: given a pair (x, U), the modality $\Box$ quantifies over all subsets of U in $\mathcal{O}(X)$ that include the actual state x, while another modality K (acting as a "knowledge" modality) quantifies over the elements of U. In other words, we have a logic that formalizes reasoning about sets and points: $\Box$ quantifies *over* the sets, while K quantifies *in* the sets. In this setup, $(x, U))$ is called a *neighborhood situation* if U is a neighborhood of x, that is, $x \in U \in \mathcal{O}$. If at (x, U) ϕ is known, this is to be interpreted as saying that we can move from the given reference point x to any other point y in the given neighborhood situation (x, U). Similarly, using the $\Box$ modality, we can shrink the neighborhood around a given reference point. In this manner, knowledge (via K) is interpreted locally in a given truthful observation set U. Effort (via $\Box$), for its part, is interpreted as neighborhood-refinement, where "more effort" corresponds to a smaller neighborhood, and a smaller neighborhood corresponds to a more refined truthful observation—and so, a possible increase in knowledge. The smaller the observation set is, the more informative whatever information we have, and the more effort we will have spent to obtain this. See Moss and Parikh (1992) for more details.

271. The main ideas of the next paragraph are covered in Bezhanishvili and Holliday (2019), which has a valuable discussion of these, and related, matters.

In other words, $\mathbf{int}(U) \subseteq \mathbf{int}(\mathbf{int}(U))$. These conditions taken together further imply things about disjunctions. Observe that an arbitrary disjunction (a union) of opens-as-properties gets verified by confirming *any one of them* and thus also requires only a finite amount of evidence.

Regarding the asymmetry in the finiteness conditions on how verification distributes over union and intersection, there is a closely related tradition of computer scientists who have taken to describing the key properties of a topology in terms of a key computational or logical idea, namely that

> open sets are analogous to *semidecidable* properties.[272]

The idea here is that the notion of *(semi)decidability* may help understand the axiomatic stability properties that characterize a topology—where, computationally speaking, an observable property of a data type is taken to correspond to a semidecision procedure.

In its native setting, one attributes the property of semidecidability to a theory or logical system. A theory is said to be semidecidable provided there is an effective method which, given an arbitrary formula, will always tell correctly when the formula is in the theory, but acts differently when the formula is not in the theory: in that case, it may give either a negative answer or no answer at all. A logical system, for its part, is regarded as semidecidable if there is an effective method for generating theorems (and only theorems) such that every theorem will eventually be generated—yet, in a semidecidable system there may be no effective procedure for checking that a formula is *not* a theorem. Applied to our present setting: this concept would help describe how, if something is observably true of a set, this can be *decided* by some evidence or verification; but if it is false, the same procedure may "run forever" and not have the means of verifying that it is not true of the set.

To better see what is going on here, referring back to our informal discussion in A.1, consider the following proposition or assertion: "The car is speeding." Affirmations and refutations of such propositions can be carried out by means of what can actually be observed, and an observation must be made in finite time (i.e., after a finite number of steps), in particular by supplying a piece of evidence. Whenever the car is unmistakably speeding, so that it is definitely true, such a proposition can be affirmed. Whenever the car is unmistakably not speeding, so that it is definitely false, it can be refuted. But in an important way, the real crux of the matter boils down to how all the other borderline cases are handled. The two extremes can help reveal what is going on here. Vickers (1996) describes how

> An assertion is *affirmative* (or *affirmably true*) iff it is true precisely in the circumstances when it can be affirmed.

Then,

> if we declare that for all border cases the assertion is *false*, then by "true" we will mean "affirmably true."

In other words—on the affirmative interpretation—any sets corresponding to true are treated as *open sets*.

We can also speak of refutative assertions as follows:

> An assertion is *refutative* iff it is false precisely in the circumstances when it can be refuted.

Then,

> if we declare that for all border cases the assertion is *true*, then by "true" we will mean "irrefutable" (in the sense of being not falsifiable).

In other words – on the refutative interpretation – sets corresponding to true are treated as *closed sets*.

272. This suggestion apparently first appeared in Smyth (1983); it is discussed in Vickers (1996), which inspired the next few pages.

Let us declare the sets for which we can make affirmations or refute the assertion *open* – since they do not include any boundary. Asserting "The car is speeding" means the car is definitely speeding, while refuting it means it is definitely not speeding.

Suppose, for instance, we can agree that the car really is speeding, where speeding means any speed strictly greater than 50 mph. As suggested in A.1, this affirmation might be tested by using finer and finer "radar guns." Then we can *affirm* the assertion—such an observation can be made—since if the speed is indeed greater than 50 mph, you will eventually discover this, provided only that you use a radar or measuring device with a granularity finer than ± 1.

As Vickers stresses, working with the affirmative interpretation, we can surely affirm $\phi \vee \psi$ either by affirming ϕ or by affirming ψ. This extends to further disjunctions, even to the point where we have an entire family of infinitely many assertions $\{\phi_i\}$—the infinite disjunction $\bigvee_i \phi_i$ can be affirmed by just affirming *any* individual ϕ_i. In short, any disjunction, even an infinite one, of affirmative assertions is still affirmative—we need affirm only one!

We can affirm $\phi \wedge \psi$ by affirming both ϕ and ψ. This extends to any finite conjunction. But to affirm an infinite conjunction $\bigwedge_{i \in I} \phi_i$, we must affirm every single one and that will require an infinite amount of work. Thus, we can only say: any *finite* conjunction of affirmative assertions is still affirmative.

Under different circumstances, it may be that the car is decidedly not speeding – in that instance, the assertion is definitely false, and so can be refuted. To affirm "not ϕ" ($\neg\phi$), we will have to make a finite observation evidencing that ϕ is definitely false (i.e., we refute ϕ).

Under the refutative interpretation, for its part, assertions correspond to falsifiable ones. And for these, the following holds: For conjunctions, if ϕ is falsifiable and ψ is any statement (falsifiable or not), then ϕ and ψ is falsifiable—it is falsifiable by *any* observation that falsifies ϕ. For disjunctions, if you have two falsifiable statements ϕ, ψ, then $\phi \vee \psi$ is falsifiable since you can demonstrate the disjunction to be false merely by supplying two observations, the first falsifying ϕ and the second falsifying ψ. But this does not extend to arbitrary disjunctions. This gives another way of seeing the conditions on closed sets.

The logic of refutative assertions will be such that it is stable under arbitrary conjunctions and finite disjunctions, while the logic of affirmative assertions will be stable under finite conjunctions and arbitrary disjunctions. Whether we take up the logic of affirmative assertions or the logic of refutative assertions, we capture the essence of a topology.

Finally—and this is the point!—suppose "The car is speeding" (ϕ) is false and, as it turns out, the car is going *exactly* 50 mph. Then, on the affirmative interpretation (i.e., using opens), you will never discover whether or not ϕ is true, regardless of how refined the radar used. We will never be able to affirm that it is false. This borderline case helps reveal the key idea of thinking of all this in terms of semidecidability: there is some test that you can perform such that, if ϕ is true, you will eventually discover this, but if it is false, you may never discover this. On the refutative side of things (i.e., using closed sets): if ϕ is false, you will eventually be able to falsify it, but if it is true, you may never discover this.

Altogether, this accords with the reading

> $\Diamond\phi :=$ "It is not verifiably false that ϕ" (or, equivalently, "The hypothesis that ϕ cannot be falsified"),

according to which $\Diamond$ is a closure operator, and $S \subseteq X$ is falsifiable if and only if $\mathbf{cl}(S) \subseteq S$—another way of seeing that we can falsify a falsifiable proposition unless it happens to be true. Furthermore, a boundary of a set $S \subseteq X$ is just $\partial(S) = \mathbf{cl}(S) \setminus \mathbf{int}(S)$, so that x is a boundary point of S precisely when there is no true piece of evidence that supports neither S nor $\neg S$.

A.4 Why Opens?

In chapter 4, I mentioned that I would attempt to address the question of why, in spite of the historical precedence given to closed sets, present-day topologists often seem to regard open sets as more primitive. Building on previous discussions, here are some plausible reasons one could give for such a preference:

1. The epistemic-topological framework sketched above helps reveal how key axioms of a topology appear to encode the notion of semidecidability, that is, the notion of a condition whose truth can be verified in finite time (but whose falsehood cannot necessarily be verified in finite time). It is sometimes suggested that this interpretation encodes the intuition from metric space settings where verifying that a point is in an open ball can be done with finite-precision computations by just producing an even smaller open ball in which it lies, yet where it is not necessarily possible to do the same sort of thing for a closed ball (if the point lies on the boundary, that we cannot make arbitrarily precise measurements in finite time becomes an issue).

   In this same epistemic-topological framework, another more general reason could have to do with the following. We can technically freely consider things either in terms of affirmative propositions or refutative propositions and convert between them. Affirmative assertions, though, for their part, do seem more primitive—or less derivative or parasitic—than their refutative counterparts. And affirmative assertions are the ones that correspond—in both intuitive and precise ways—to open sets. This priority of the affirmably true might be further developed into something like a proper defense of the priority of open sets.
2. Another reason I have heard for preferring open sets over closed sets is the apparent ease with which certain formulations of continuity, and related results, can be given using open sets and potentially infinite unions thereof—compared to using closed sets, where one must make restrictions to the finite case. Indeed, this can be related to sheaves as well. Over and above the important historical connection between sheaves and local homeomorphisms, there is a viable explanation for the general preference for working with open sets in that what has been called the *open pasting lemma*—or the "local formulation of continuity" (as in Munkres 2015)—stipulates that if $X = \cup_i U_i$ is a possibly infinite union of open sets with continuous sections $s_i : U_i \to E$ that agree on their overlaps, then it follows that the (set-theoretically defined) section $s : X \to E$ will be continuous as well. However, if closed sets are used here, then the gluing argument works only for those covers that have *finitely* many closed sets.[273] More generally, if it is really the case that having arbitrary unions of open sets at one's disposal (such that this is open) is somehow more useful than having arbitrary intersections of closed sets (such that this is closed), then I think that this sort of justification has some merits and deserves to be probed further. On the other hand, to the extent that the results and formulations adduced as defense in fact really only highlight a lack of human ingenuity, or some feature of the current practice of mathematics, preventing us from finding the right formulation and obtaining the same results while using closed sets, then this reason is less compelling.
3. A last reason I will list—which I have not seen developed before—is a little more complicated and difficult to parse. Fundamentally, the idea is that the preference may be influenced by differences in the underlying algebraic structures of open sets versus closed sets.

   The lattice of open subsets of a topological space supplies us with an important and canonical example of a Heyting algebra, while its dual—the lattice of closed subsets—supplies a canonical example of a co-Heyting algebra. For a co-Heyting algebra (like that of closed subsets of a space), we have seen that there is an induced operator $\sim$ that could be called *paraconsistent negation*, some of whose properties were explored in chapter 7. The impact of this alternative negation operator can be initially appreciated by considering that $\mathbf{cl}(U) \cap \mathbf{cl}(U^c)$, where $\mathbf{cl}(U^c) := \sim U$, will in general be nonempty. By defining such an alternative negation with $\sim$ ("non"), as the closure of the complement, we get a co-Heyting algebra of closed sets. In the general setting of co-Heyting algebras, which model paraconsistent logics, inconsistent theories are modeled as the ones that include the formulas that are true at the boundaries. As such, the logic of closed set topologies can be seen as being *paraconsistent* or forming a *paraconsistent algebra*, as a formula ϕ and its paraconsistent negation $\sim \phi$ may intersect at the boundary of their extensions.

273. Curry (2014) makes this observation.

Recall that paraconsistency is the blanket term for a logical system where the principle of noncontradiction may fail, that is, we can have $\alpha \wedge$ not-α. In general, paraconsistent structures are those for which theories that are true at boundary points include formulas and their negation; a boundary of a theory T can be construed as those sequents that neither "fully follow from" T nor "fully contradict" T. Note that, with a paraconsistent logic, this does not mean that *all* contradictions are true—rather, that there are some contradictions that do not entail a trivial theory. In general, the idea here is that the extension of the conjunction of some formulas and their negations may not be the empty set. As such, viewed in semantic terms, there may exist some states for which a formula and its negation are true. There are many different logical structures that are used to represent paraconsistent logics, but co-Heyting algebras have been seen as an especially natural choice.

The reason why closed set topologies form a paraconsistent structure has to do with the fact that theories that are true at boundaries include formulas and their negation—in other words, a formula α and its paraconsistent negation $\sim\alpha$ intersect at the boundary of their extensions. In such settings, the boundary operator also plays a decisive and active role. In general, boundaries play an important role in supplying topological semantics for paraconsistent logics.[274] As Lawvere writes,

> That the notion of boundary is just that of "logical contradiction" (within the realm of closed sets) follows at once from the intuitive notion of motion: indeed, since the unit interval is connected, any continuous path which is in A at time 0 and in $\sim A$ at time 1 must at some intermediate time be in both A and $\sim A$, i.e., must pass through the boundary of A. (Lawvere 1986, 10)

We have seen that sheaves fundamentally involve assignments of data to a base topology. In the usual presentations of sheaves on topological spaces, it is standard practice to almost exclusively work with open set topologies. But in the sheaf construction, the base topology is of special importance. For one thing, the algebra of the base space topology can be seen in terms of the algebra of the sheaf section structure. And historically at least, sheaves were defined over closed sets before they were defined over open sets.[275] By considering sheaves with respect to closed set topologies, then, the sheaf morphisms ought to reflect this algebra—so, using closed sets, we could have morphism algebras equipped with paraconsistent negations. Moreover, by taking closed sets for our base topology, we could introduce into the sheaf construction the valuable notion of boundary—something that will not exist for the corresponding open set sheaf construction. As William James has suggested, building on suggestions of Lawvere:[276]

> One area in which this may work for us is the mathematics of physics where the boundaries of a body are as important as the parts of a body inasmuch as physics concerns itself with the interactions of bodies in a system. Lawvere in the introduction to *Categories in Continuum Physics*, mentions the speculation that there is a role for a closed set sheaf in thermodynamics as a functor from a category of parts of a body to a category of "abstract thermodynamical state-and-process systems" (Lawvere 1986, 9). Lawvere recognises the particular properties of closed set topologies that make them interesting to us, namely that as algebras they provide us with a formalisation of what we call a paraconsistent negation. Sheaves are then of interest to us in our project of developing paraconsistent logic in categories for the way in which they transport algebras of a topology into the structure of a category of sheaves over that topology. (James 1996, 127–128)

274. See Başkent (2013), Goodman (1981), and Mortensen (2000).

275. Historically, in 1946 Jean Leray first defined a sheaf as a way of assigning modules to closed sets in an inclusion-reversing manner. For a historical account of the early development of the sheaf concept, the reader can consult Fasanelli (1981).

276. The interested reader should also consult James (1992, 1995).

There may indeed be considerable virtues to be had by giving more attention to closed sets, especially in relation to sheaves, and the reader is urged to explore these further. For instance, by looking at sheaves over closed sets, ultimately paraconsistency could be integrated into topos theory. Moreover, we can mention that in any presheaf topos, the lattice of all subobjects of any given object is an example of a co-Heyting algebra (as well as being a Heyting algebra).

All this is to make the following point: The co-Heyting operations are in general not preserved by substitution (inverse image) along maps—and this is as opposed to the Heyting "not" (¬) and even the "possibility"-like operators given by Grothendieck topologies. I believe that this—and subtleties having to do with how boundaries and continuity interact—may contribute (albeit in the background, or beneath awareness) to the preference for working with open set topologies. By encoding the notion of "having no boundary," open sets are able to avoid some of the subtleties associated with explicitly working with boundaries and the induced paraconsistent structure supporting some inconsistency tolerance (which generally seems to require more careful treatment). By comparison, there are close connections between Heyting algebras and intuitionistic logic (which logic is particularly useful for representing a number of situations relevant to the mathematical treatment of continuity). On account of these features of the algebraic structures of open versus closed sets, we might propose the following:

> topology in terms of open sets ≈ algebra of observables where the law of excluded middle can fail but the law of noncontradiction holds in general;

> topology in terms of closed sets ≈ algebra of observables where the law of noncontradiction can fail but the law of excluded middle holds in general.

Finally, related to this issue of boundaries: a great deal of mathematics is concerned with *continuity* in its many different forms. Inquiring into the continuity of something, say some function f defined on some region, at a point p, is to invite comparison of the behavior of f *all around* p—yet, at boundaries, we cannot necessarily go around the point without leaving our area of concern. This merely suggests one of the related ways in which one might expect to find, on account of the general features of boundaries, that there is a more natural alignment of continuity with open, rather than closed, sets—as the former exclude all boundaries, while the latter include them all.

A.5 What Is Topology Really About?

At least as far as general (point-set) topology goes, the story just told suggests that the essence of topology can be seen as being closely connected to certain features of logic and the logic of verification, specifically as this relates to the implications of how we formulate notions of *negation* and our dealings with *boundaries*. Guided by such a story, a think a reasonable case can be made for wanting to view

> topology as a formal framework for studying reasoning about systems of parts of a whole in such a way that the nondegeneracy of *boundaries* is respected.

In particular,

> open set topology ≈ the structure of reasoning in the absence of boundaries (or with approximations),

while

> closed set topology ≈ the structure of reasoning in the presence of boundaries.

Depending on whether we choose to admit boundaries or not, certain prices will have to be paid—involving the tolerance of violations of a generalized "law of noncontradiction," and how negations must be treated. And it is arguably these prices that are truly determinative of the structure we call a topology.

However, what does all this have to do with the wider notions we have explored in this book—from the familiar topologies to Grothendieck topologies to Lawvere-Tierney topologies? Already in proposition 132 (chapter 5) we saw that we can effectively drop the open sets of a topological space X and just look at the category of sheaves on X—this later category can then be used, by looking at the subsheaves of the terminal sheaf, to recover the original data of the open sets. This suggests that, in some sense, sheaves are the more essential of the two, even if at first it may appear as if the space on which sheaves are built needs to be assumed by the sheaf construction. If that is indeed the case, how can we reconcile everything we said about sheaves—their power, how to think about them, and how to use them to think more deeply about space—with what this particular logical story seems to suggest about the nature of topology and space? I leave that to the reader to ponder!

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