

# Sheaf Theory Through Examples

Daniel Rosiak

MIT Press  
2020



# Contents

0.1	Introduction . . . . .	5
0.2	A First Pass at the Idea of a Sheaf . . . . .	9
0.3	Outline of Contents . . . . .	18
<b>1</b>	<b>Categorical Fundamentals for Sheaves</b>	<b>21</b>
1.1	Categorical Preliminaries . . . . .	21
1.1.1	Aside on “No Objects” . . . . .	29
1.1.2	A Few More Examples . . . . .	30
1.1.3	Some New Categories From Old . . . . .	32
1.2	Prelude to Sheaves: Presheaves . . . . .	35
1.2.1	Functors . . . . .	35
1.2.2	Examples of Functors . . . . .	37
1.2.3	Natural Transformations . . . . .	50
1.3	Yoneda: The Most Important Idea in Category Theory . . . . .	58
1.3.1	First, Enrichment! . . . . .	58
1.3.2	Downsets and Yoneda in the Miniature . . . . .	63
1.3.3	Representability Simplified . . . . .	67
1.3.4	More on Representability, Fixed points, and a Paradox . . . . .	72
1.3.5	Yoneda in the General . . . . .	77
1.3.6	Philosophical Pass: Yoneda and Relationality . . . . .	80
1.4	Adjunctions . . . . .	84
1.4.1	Adjunctions through Morphology . . . . .	84
1.4.2	Adjunctions through Modalities . . . . .	102
1.5	Final Thoughts on Fundamentals . . . . .	123
<b>2</b>	<b>Presheaves Revisited</b>	<b>125</b>
2.1	Seeing Structures as Presheaves . . . . .	125
2.2	The Presheaf Action . . . . .	132
2.2.1	Right Action Terminology . . . . .	132
2.2.2	Four Ways of Acting as a Presheaf . . . . .	134
2.2.3	Philosophical Pass: The Four Action Perspectives . . . . .	148

<b>3</b>	<b>First Look at Sheaves</b>	<b>151</b>
3.1	Sheaves: The Topological Definition . . . . .	153
3.1.1	A Sheaf as Restriction-Collation . . . . .	156
3.2	Examples . . . . .	158
3.2.1	Philosophical Pass: Sheaf as Local-Global Passage . . . . .	173
3.2.2	Three Historically Significant Examples . . . . .	175
3.2.3	Bundles to (Pre)Sheaves . . . . .	192
3.2.4	(Pre)Sheaves to Bundles . . . . .	196
3.2.5	The Bundle-Presheaf Adjunction . . . . .	200
3.2.6	Take-Aways . . . . .	203
3.2.7	What is <i>Not</i> a Sheaf . . . . .	204
3.2.8	Presheaves and Sheaves in Order Theory . . . . .	209
<b>4</b>	<b>Sheaf Cohomology through Examples</b>	<b>215</b>
4.1	Simplices and their Sheaves . . . . .	215
4.1.1	Sheaf Morphisms and Some Operations on Sheaves . . . . .	228
4.2	Sheaf Cohomology . . . . .	231
4.2.1	Primer on (Co)Homology . . . . .	232
4.2.2	Cohomology with Sheaves . . . . .	238
4.2.3	Philosophical Pass: Sheaf Cohomology . . . . .	242
4.2.4	A Glimpse into Cosheaves . . . . .	244
<b>5</b>	<b>Toposes</b>	<b>251</b>
5.1	Revisiting Covers: Towards General Sheaves . . . . .	251
5.2	Grothendieck Topologies: First Pass . . . . .	253
5.2.1	Sieves . . . . .	255
5.2.2	Grothendieck Topology Defined . . . . .	257
5.2.3	Sheaves on a Site . . . . .	262
5.2.4	2-enriched Simpler Version . . . . .	264
5.2.5	Grothendieck toposes . . . . .	272
5.2.6	The Plus Construction . . . . .	273
5.3	Further Examples . . . . .	276
5.3.1	Philosophical Pass: The Idea of Toposes . . . . .	290
5.4	Elementary Toposes . . . . .	292
5.4.1	The Subobject Classifier . . . . .	292
5.4.2	Lawvere-Tierney Topologies and their Sheaves . . . . .	299
5.4.3	Topos Logic . . . . .	312
5.5	Morphisms of Toposes . . . . .	315
5.5.1	Geometric Morphisms Defined . . . . .	315
5.6	Towards Cohesive Toposes . . . . .	320
5.6.1	Foray into Cohesive Toposes . . . . .	322
5.6.2	Philosophical Pass: A New Dialectical Science? . . . . .	330

# Introduction to Sheaves Through Examples

## 0.1 Introduction

In many cases, events and objects are given to observation as extended through time and space, and so the resulting data is local and distributed in some fashion. For now, we can think of this situation in terms of data being indexed by, or attached to (“sitting over”), given regions or domains of some sensors. In saying that the data is *local*, we just mean that it holds only within, or is only defined over, a certain region, i.e., its validity is restricted to a prescribed region or domain or reference context, and we expect that whenever a property holds at a point of its extended domain, then it also holds at “nearby” points. In other words, even if the measurements or data concerning some system of regions are not all related in a “nice” way, a given measurement will be closely related to its “neighbors,” or to nearby measurements. We collect temperature and pressure readings and thus form a notion of ranges of possible temperatures and pressures over certain geographical regions; we record the fluctuating stockpile of products in a factory over certain business cycles; we accumulate observations or images of certain patches of the sky or the earth; we gather testimonies or accounts about particular events understood to have unfolded over a certain region of space-time; we build up a collection of test results concerning various parts of the human body; we amass collections of memories or recordings of our distinct interpretations of a certain score of music; we develop observations about which ethical and legal principles or laws are respected throughout a given region or network of human actors; we form a concept of our kitchen table via various observations and encounters, assigning certain attributes to those regions of space-time delimiting our various encounters with the table, where we expect that the ascribed properties or attributes are present throughout the entirety of a region of their extension. Even if certain phenomena are not intrinsically local, frequently its measurement or the method employed in data collection may still be local.

But even the least scrupulous person does not merely accumulate or amass local

or partial data points. From an early age, we try to understand the various modes of *connections* and *cooperations* between the data, to patch these partial pieces together into a larger whole whenever possible, to resolve inconsistencies among the various pieces, to go on to build coherent and more global visions out of what may have only been given to us in pieces. As informed citizens or as scientists, we look at the data given to us on arctic sea-ice melting rates, on temperature changes in certain regions, on concentrations of greenhouse gases at various latitudes and various ocean depths, etc., and we build a more global vision of the changes to our entire planet on the basis of the connections and feedbacks between these various data. As investigators of a crime, we must “piece together” a complete and consistent account of the events from the partial accounts of various witnesses. As doctors, we must infer a diagnosis and a plan of action from the various individual test results concerning the parts of a patient’s body. We take our many observations concerning the behavior of certain networks of human actors and try to form global ethical guidelines or principles to guide us in further encounters.

Yet sometimes information is simply not local in nature. Roughly, one might think of such non-locality in terms of how, as perceivers, certain attributes of a space may appear to us in a particular way but then cease to manifest themselves in such a way over subparts of that space, in which case one cannot really think of the perception as being built up from local pieces. For a different example: in the game of Scrabble<sup>TM</sup>, one considers assignments of letters, one by one, to the individual squares in a lattice of squares, with the aim of building words out of such assignments. One might thus suspect that we have something like a “local assignment” of data (letters in the alphabet) to an underlying space ( $15 \times 15$  grid of squares). Yet this assignment of letters to squares in order to form words is not really local in nature, since, while we do assign letters one by one to the grid of squares, the smallest unit of the game is really a *legal word*, but not all sub-words or parts of words are themselves words, and so a given word (data assignment) over some larger region of the board may cease to be a word (possible data assignment) when we restrict attention to a subregion.

Even when information is local, there are many instances where we cannot synthesize our partial perspectives into a more global perspective or conclusion. As investigators, we might fail to form a coherent version of events because the testimonies of the witnesses cannot be made to agree with what other data or evidence tells us regarding certain key events. As musicians, we might fail to produce a compelling performance of a score because we have yet to figure out how to take what is best in each of our “trial” interpretations of certain sections or parts of the entire score and splice them together into a coherent single performance or recording of the entire score. A doctor who receives conflicting information from certain test results, or testimony from the patient that conflicts with the test results, will have difficulty making a diagnosis. In explaining the game of rock-paper-scissors to children, we tell them that rock beats scissors, scissors beats paper, and paper

beats rock, but we cannot tell the child how to win *all the time*, i.e., we cannot answer their pleas to provide them with a global recipe for winning this game.

For distinct reasons, differing in the gravity of the obstacle they represent, we cannot always “lift” what is local or partial up to a global value assignment or solution. A problem may have a number of viable and interesting local solutions but still fail to have even a single global solution. When we do not have the “full story,” we might make faulty inferences. Ethicists might struggle with the fact that it is not always obvious how to pass from the instantiations or particular variations of a seemingly locally valid prescription, valid or binding for (say) a subset of a network of agents, to a more global principle, valid for a larger network. In the case of the doctor attempting to make a diagnosis out of conflicting data, it may simply be a matter of either collecting more data, or perhaps resolving certain inconsistencies in the given test results by ignoring certain data in deference to other data. Other times, as in the case of rock-paper-scissors, there is simply nothing to be done to overcome the failed passage from the given local ranking functions to a global ranking function, for the latter simply does not exist. The intellectually honest person will eventually want to know if their failure to lift the local to the global is due to the inherent particularity or contextuality of the phenomena being observed or whether it is simply a matter of their own inability to reconcile inconsistencies or repair discrepancies in data-collecting methods so as to patch together a more global vision out of these parts.

*Sheaf theory* is the roughly 70-year old collection of concepts and tools designed by mathematicians to tame and precisely comprehend problems with a structure exactly like the sorts of situations introduced above. The reader will have hopefully noticed a pattern in the various situations just described. We produce or collect assignments of data indexed to certain regions, where whenever data is assigned to a particular region, we expect it to be applicable throughout the entirety of that region. In most cases, these observations or data assignments come already distributed in some way over the given network formed by the various regions; but if not, they may become so over time, as we accumulate and compare more local or partial observations. In certain cases, together with the given value assignments and a natural way of decomposing the underlying space, revealing the relations between the regions themselves, there may emerge correspondingly natural ways of restricting assignments of data along the subregions of given regions. In such cases, in this movement of decomposition and restriction, the glue or system of translations binding the various data together, permitting some sort of transit between the partial data items, becomes explicit; in this way, an internal consistency among the parts may emerge, enabling the controlled gluing or binding together of the local data into an integrated whole that now specifies a solution or system of assignments over a larger region embracing all of those subregions. Such structures of coherence emerging among the partial patches of local data, once explicitly acknowledged and developed, may enable a unique *global* observation or solution, i.e., an observation

that no longer refers merely to yet another local region but now extends over and embraces all of the regions at once; as such, it may even enable predictions concerning missing data or at least enable principled comparisons between the various given groups of data. Sheaves provide us with a powerful tool for precisely modeling and working with the sort of local-global passages indicated above. Whenever such a local-global passage is possible, the resulting global observations make transparent the forces of coherence between the local data points by exhibiting to us the principled connections and translation formulas between the partial information, making explicit the glue by which such partial and distinct clumps of data can be “fused” together, and highlighting the qualities of the distribution of data. And once in this framework, we may even go on to consider systematic passages or translations between distinct such systems of local-to-global data.

On the other hand, when faced with *obstructions* to such a local-global passage, we typically revise our basic assumptions, or perhaps the entire structure of our data, or maybe just our manner of assigning the data to our regions. We are usually motivated to do this in order to allow precisely such a global passage to come into view. When we can satisfy ourselves that nothing can be done to overcome these obstructions, we examine what the failure to pass from such local observations to the global in this instance can tell us about the phenomena at hand. *Sheaf cohomology* is a tool used for capturing and revealing precisely obstructions of this sort.

The purpose of this book is to provide an inviting and (hopefully) gentle introduction to sheaf theory, where the emphasis is on explicit constructions, applications, and a wealth of examples from many different contexts. Sheaf theory is typically presented as a highly specialized and advanced tool, belonging mostly to algebraic topology and algebraic geometry (the historical “homes” of sheaves), and sheaves accordingly have acquired a somewhat intimidating reputation. And even when the presentation is uncharacteristically accessible, emphasis is typically placed on abstract results, and it is left to the reader’s imagination (or “exercises”) to consider some of the things they might be used for or some of the places where they can be found. This book’s primary aim is to dispel some of this fear, to demonstrate that sheaves can be found all over (and not just in highly specialized areas of advanced math), and to give a wider audience of readers a more inviting tour of sheaves. Especially over the last few years, the interest in sheaves among less and less specialized groups of people appears to be growing immensely; but, whenever I spoke to newcomers to sheaves, I invariably heard that the existing literature was either too specialized or too forbidding. This book accordingly also aims to fill a gap in the existing literature, which for the most part tends to either focus exclusively on a particular use of sheaves or assumes a formidable pre-existing background and high tolerance for abstraction. I do not share the view that applications or concrete constructions are mere corollaries of theorems, or that examples are mere illustrations with no power to inform “deeper” conceptual advances. I am not sure if I would go as far as to endorse Vladimir Arnold’s idea that “The content of a



mathematical theory is never larger than the set of examples that are thoroughly understood,” but I do believe that one barrier to the wider recognition of the immense power of sheaf theory lies in the tendency to present much of sheaf theory as if it were a forbiddingly abstruse or specialized tool, or as belonging mainly to one area of math. One thing this book aims to show is that it is no such thing. Moreover, well-chosen examples are not only useful, both pedagogically and “psychologically,” in helping newcomers get a better handle on the abstract concepts and advance forwards with more confidence, but can even jostle experts out of the rut of the ‘same old examples’ and present interesting challenges both to our fundamental intuitions of the underlying concepts and to preconceptions we might have about the true scope of applicability of those concepts.

Before outlining the contents of the book, the next section offers a more detailed, but still “naive,” glimpse into the *idea* of a sheaf via a toy construction, with the aim of better establishing intuitions about the underlying sheaf idea.

## 0.2 A First Pass at the Idea of a Sheaf

Suppose we have some ‘region’, which, for the moment, we can represent very naively and abstractly as



We are less interested in the “space itself” and more in how the space serves as a site where various things *take place*. In other words, we think of this region as really just an abstract domain supporting various *happenings*, where such happenings carry information for appropriate sensors or “measuring instruments” (in a very generalized sense), so that interrogating the space becomes a matter of asking the sensors about what is happening on the space.<sup>1</sup> For instance, the region might

---

<sup>1</sup>The description of sheaves as “measuring instruments” or the “meter sticks” on a space that we are invoking—so that the set of all sheaves on a given space supply one with an arsenal of all the meter sticks measuring it, yielding “a kind of ‘superstructure of measurement’”—ultimately comes from Grothendieck, who was largely responsible for many of the key ideas and results in the early development of sheaf theory. In speaking of (another early sheaf theorist) Jean Leray’s work in the 40s, Grothendieck said this:

The essential novelty in his ideas was that of the (Abelian) sheaf over a space, to which Leray associated a corresponding collection of cohomology groups (called “sheaf coefficients”). It is as if the good old standard “cohomological metric” which had been used up to then to “measure” a space, had suddenly multiplied into an

be the site of some happenings that supply *visual information*, so that as a sensor monitors the happenings over a region (or some part of it), it collects specifically visual information about whatever is going on in the area of its purview:




There might then be another sensor, taking in visual information about another region or part of some overall ‘space’, offering another “point of view” or “perspective” on another part of the space; and it may be that the underlying regions monitored by the two sensors overlap in part:

---


unimaginably large number of new “meter sticks” of every shape, size and form imaginable, each intimately adapted to the space in question, each supplying us with very precise information which it alone can provide. This was the dominant concept involved in the profound transformation of our approach to spaces of every sort, and unquestionably one of the most important mathematical ideas of the 20th century. (*Récoltes et Semailles*, Promenade 12)

Then the sheaves on a given space will incorporate “all that is most essential about that space...in all respects a lawful procedure [replacing consideration of the space by consideration of the sheaves on the space], because it turns out that one can “reconstitute” in all respects, the topological space by means of the associated “category of sheaves” (or “arsenal” of measuring instruments)...[H]enceforth one can drop the initial space...[W]hat really counts in a topological space is neither its “points” nor its subsets of points, nor the proximity relations between them; rather, it is the *sheaves on that space, and the category that they produce*” (Promenade 13). The reader for whom this is overwhelming should press on and rest assured that we will have a lot more to say about all this later on in the book, and the notions and results alluded to in the above will be motivated and discussed in detail.




Since we are ultimately interested in the informative happenings on the space, we want to see how the distinct “perspectives” on what is happening throughout the space are themselves related; to this end, a very natural thing to do is ask how the data collected by such neighboring sensors are related. Specifically, a very natural thing to ask is whether and how the perspectives are *compatible* on such overlapping sub-regions, whenever there are such overlaps between the underlying regions over which they, individually, collect data.

A little more explicitly: if we assume the first sensor collects visual data about its region (call it  $U_1$ ), we may imagine, for concreteness, that the particular sort of data available to the sensor consists of sketches, say, of characters or letters (so that the underlying region acts as some sort of generalized sketchpad or drawing board)




While not really necessary, the sensor might even be supposed to be equipped to “process” the information it collects, translating such visual inputs into reasonable guesses about which possible capital letter or character the partial sketch is supposed to represent. In any event, attempting to relate the two “points of views” by


considering their compatibility on the region where their two surveyed regions overlap, we are really thinking about first making a selection from each of the collections of data assigned to the individual sensors:




Corresponding to how the underlying regions are naturally related by an “inclusion” relation, the compatibility question, undertaken at the level of the selections (highlighted in gray above) from the collections of all informative happenings on the respective regions, will involve looking at whether those data items “match” (or can otherwise be made “compatible”) when we restrict attention to that region where the individual regions monitored by the separate sensors overlap:



If the given selection from what they individually “see” does match on the overlap, then, corresponding to how the regions  $U_1$  and  $U_2$  may be joined together to form a larger region,



at the level of the data on the happenings over the regions, we can pull this data back into an item of data given now over the entire space  $U_1 \cup U_2$ , with the condition that we expect that restricting this new, more comprehensive, perspective back down to the original individual regions  $U_1$  and  $U_2$  will give us back whatever the two individual sensors originally “saw” for themselves:




In other words, given some selection from what sensor 1 “sees” as happening in its region  $U_1$  and from what sensor 2 “sees” as happening in its region  $U_2$ , provided

their “story” agrees about what is happening on the overlapping region  $U_1 \cap U_2$ , then we can paste their individual visions into a single and more global vision or story about what is happening on the overall region  $U_1 \cup U_2$  (and we expect that this story ultimately “comes from” the individual stories of each sensor, in the sense that restricting the “global story” down to region  $U_1$ , for instance, will recover exactly what sensor 1 already saw on its own).


Another way to look at this is as follows: while the sensor on the left, when left to its own devices, will believe that it may be seeing a part of any of the letters  $\{B, E, F, P, R\}$ , checking this assignment’s compatibility with the sensor on the right amounts to constraining what the left sensor believes by what the sensor on the right “knows,” in particular that it cannot be seeing an  $E$  or an  $F$ . Symmetrically, the sensor on the right will have its own “beliefs” that might, in the matching with the left sensor, be constrained by whatever the left sensor “knows.” In matching the two sensors along their overlap, and patching their perspectives together into a single, more collective, perspective now given over a larger region (the union of their two regions), we are letting what each sensor individually “knows” constrain and be constrained by what the other “knows.”

In this way, as we cover more and more of a ‘space’ (or, alternatively, as we decompose a given ‘space’ into more and more pieces), we can perform such compatibility checks at the level of the data on the happenings on the ‘site’ (our collection of regions covering a given space), and then “glue together,” piece by piece, the partial perspectives represented by each sensor’s local data collection into more and more embracing or “global” perspectives. More concretely, continuing with our present example, suppose there are two additional regions, covering now some southwest and southeast regions, respectively, so that, altogether, the four regions cover some region (represented by the main square):




where we have left implicit the obvious intersections ( $U_1 \cap U_2$ ,  $U_3 \cap U_4$ ,  $U_1 \cap U_3$ , etc.). With the four regions  $U_1, U_2, U_3$ , and  $U_4$ , to each of which there corresponds a particular sensor, we have the entire central region  $U = U_1 \cup U_2 \cup U_3 \cup U_4$  ‘covered’. Part of what this means is that, were you to invite *another* sensor to observe the happenings on some further portion of the space, in an important sense, this extra sensor would be superfluous—since, together, the four regions monitored by the four individual sensors already have the overall region ‘covered’.

For concreteness, suppose we have the following further selections of data from the data collected by each of these new (southwest and southeast) sensors, so that altogether, having performed the various compatibility checks (left implicit), the resulting system of “points of view” on our site can be represented as follows:




This system of mutually compatible local data assignments or “measurements” of the happenings on the space—where the various data assignments are, piece

by piece, constrained by one another, and thereby patched together to supply an assignment over the *entire* space covered by the individual regions—is, in essence, what constitutes our *sheaf*. The idea is that the data assignments are being “tied together” in a natural way



where this last picture is meant to serve as motivation or clarification regarding the agricultural terminology of ‘sheaf’:





Here one thinks of various ‘regions’ as the parcels of an overall ‘space’ covered by those pieces, the collection of which then serves as a ‘site’ where certain happenings are held to take place, and the abstract sensors capturing local snapshots or measurements of all that is going on in each parcel are then regarded as being collected together into ‘stalks’ of data, regarded as sitting over (or growing out of) the various parts of the ground space to which they are attached. A selection of a particular snapshot made from each of the individual stalks (collections of snapshots) then amounts to a cross section and the process of restriction (along intersecting regions) and collation (along unions of regions) of these sections then captures how the various stalks of data are “bound together.”

To sum up, then: the first characteristic feature of this construction is that some information is received or assigned *locally*, so that the records or observations made by each of the individual sensors are understood as being “about,” or indexed to, the entirety of some limited region, so that whenever something holds or applies at a “point” of that region, it will hold nearby as well. Next, since together the collection of regions monitored by the individual sensors may be seen as *collectively covering* some overall region, we can check that the individual sensors that cover regions that have some overlap can “communicate” their observations to one another, and a natural expectation is that, however different their records are on the non-overlapping region, there should be some sort of *compatibility* or *agreement* or *mutual constraining* of the data recorded by the sensors over their shared, overlapping region; accordingly, we ask that each such pair of sensors covering overlapping regions “check in” with one another. Finally, whenever such compatibility can be established, we expect that we can bind the information supplied by each sensor together, and regard them as patching together into a *single sensor supplying data over the union* of the underlying (and partially overlapping) individual regions, in such a way that were we to “restrict” that single sensor back down to one of the original regions, we would recover exactly the partial data reported by the original sensor assigned to that individual region.

While most of the more fascinating and conspicuous examples of such a construction come from pure and applied math, something very much like the sheaf construction appears to be operative in so many areas of “everyday life.” For instance, related to the toy example discussed above, even the way our binocular vision systems work appears to involve something like the collation of images into a single image along overlapping regions whenever there is agreement (from the input to each separate eye).<sup>2</sup> More generally, image and face recognition appears to operate, in a single brain (where clusters of neurons play the role of individ-

---

<sup>2</sup>That visual information processing itself appears to fundamentally involve some sort of sheaf-like process appears even more acutely in other species, such as certain insects like the dragonfly, whose compound eyes contain up to 30,000 facets, each facet within the eye pointing in a slightly different direction and taking in light emanating from only one particular direction, resulting in a mosaic of partially overlapping images that are then integrated in the insect brain.

ual sensors), in something like the patchwork “sum of parts” way described above. Moving beyond the individual, collective knowledge itself appears to operate in a fundamentally very similar way: a society’s store of knowledge consists of a vast patchwork built up of partial records and data items referring to particular (possibly overlapping) regions, each of which data items can be (and often are!) checked for compatibility whenever they involve data that both refer to, or make claims about, the same underlying domain.

The very simple and naive presentation given to it above runs the risk of downplaying the power and scope of this construction; it would be difficult to overstate just how powerful the underlying idea of a sheaf is. An upshot of the previous illustration, though, is that while sheaves are often regarded as highly abstract and specialized constructions, whose power derives from their sophistication, the truth is that the underlying idea is so ubiquitous, so “right before our eyes,” that one might even be impressed that it was finally named explicitly so that substantial efforts could then be made to refine our ideas of it. In this context, one is reminded of the old joke about the fish, where an older fish swims up to two younger fish, and greets them “morning, how’s the water?” After swimming along for some time, one of the younger fishes turns to the other and says

“What the hell is water?”

In this same spirit, Grothendieck would highlight precisely this “simplicity” of the fundamental idea behind sheaves (and, more generally, toposes):

As even with the idea of sheaves (due to Leray), or that of schemes, as with all grand ideas that overthrow the established vision of things, the idea of the topos had everything one could hope to cause a disturbance, primarily through its “self-evident” naturalness, through its simplicity (at the limit naive, simple-minded, “infantile”) – through that special quality which so often makes us cry out: “Oh, that’s all there is to it!”, in a tone mixing betrayal with envy, that innuendo of the “extravagant”, the “frivolous”, that one reserves for all things that are unsettling by their unforeseen simplicity, causing us to recall, perhaps, the long buried days of our infancy.... (*Récoltes et Semailles*, Promenade 13)

### 0.3 Outline of Contents

The rest of the book is structured as follows. The first chapter is dedicated to exposition of the most important category-theoretic concepts, tools, and results needed for the subsequent development of sheaves. Category theory is indispensable to the presentation and understanding of the notions of sheaf theory. While in the last decade there have appeared a number of accessible introductions to category the-

ory,<sup>3</sup> feedback from readers of earlier drafts of this book convinced me that the best approach to an introduction to sheaves that aims to reach a much wider audience than usual would need to be as self-contained as possible. In this first chapter, all the necessary categorical fundamentals are motivated and developed. The emphasis here, as elsewhere in the book, is on explicit constructions and examples. For instance, the concept of an *adjunction*, and key abstract properties of such things, is introduced and developed first through an example involving “dilating” and “eroding” an image, then again through a development of “possibility” and “necessity” modalities applied to both modeling consideration of attributes of a person applied to them *qua* the different “hats” they wear in life, and then applied to graphs of traveling routes. While the reader already perfectly comfortable with category theory is free to skip this chapter or just skim through it, or refer back to later cited examples as needed, there are a few novel examples and philosophical discussions of important results such as the Yoneda lemma that may interest the expert as well.

Chapter 2 returns to presheaves (introduced in Chapter 1) to consider them in more depth. It discusses four main perspectives on *presheaves*, develops a few notable examples of each of these, and develops some useful ways of understanding such constructions more generally. This is done both for its own sake and in order to build up to the following chapter dedicated to the initial development of the sheaf concept.

Chapter 3 introduces sheaves (specifically on topological spaces) and some key sheaf concepts and results—as always, through a diverse collection of examples. Throughout this chapter, some of the vital conceptual aspects of sheaves in the context of topological spaces are motivated, teased out, and illustrated through the various examples, and sometimes the same aspect is revisited from new perspectives as the level of complexity of the examples increases.

Chapter 4 is dedicated to a “hands on” introduction to sheaf cohomology. The centerpiece of this chapter is an explicit construction, with worked-out computations, involving sheaves on complexes. There is also a brief look at *cosheaves* and an interesting example relating sheaves and cosheaves.

Chapter 5 revisits and revises a number of earlier concepts, and develops sheaves from the more general perspective of *toposes*. The important notions in topos theory (especially as this relates to sheaves) are motivated and developed through a variety of examples. We move through various layers of abstraction, from sheaves on a site (with a Grothendieck “topology”) or Grothendieck toposes to elementary toposes. The last few sections are devoted to illustrations, through concrete examples, of

---

<sup>3</sup>The general reader without much, or any, background in category theory is especially encouraged to have a look at the engaging and highly accessible [Spi14]. Readers with more prior mathematical experience may find [Rie16], displaying the ubiquity of categorical constructions throughout many areas of mathematics, a compelling introduction. [LR03] is also highly recommended, especially for those readers content to be challenged to work many things out for themselves through thought-provoking exercises, often giving one the feeling of “re-discovering” things for oneself.

some slightly more advanced topos-theoretical notions and examples. The book concludes with an abridged presentation of some special topics, including a brief introduction to *cohesive toposes*. There are many other directions the book could have taken at this point, and more advanced sheaf-theoretical topics that might have been considered, but in the interest of space, attention has been confined to this short final section on the special topic of cohesive toposes.

Throughout each chapter, I occasionally pause for a few pages to highlight, in a more “philosophical” fashion (in what I call “Philosophical Passes”), some of the important conceptual features to have emerged from the preceding technical developments. The overall aim of the “Philosophical Pass” sections is to periodically step back from the technical details and examine the contributions of sheaf theory and category theory to the broader development of ideas. These sections may provide some needed rest for the reader, letting the brain productively “switch modes” for some time, and giving one something to think about “beyond the formal details.” A lot of category theory, and the sheaf theory built on it, is deeply philosophical, in the sense that it speaks to, and further probes, questions and ideas that have fascinated human beings for millenia, going to the heart of some of the most lasting and knotty questions concerning, for instance, what an individual object is, the nature of the concept of ‘space’, and the dialectics of continuity and discreteness. I hope it is not entirely due to my bias as someone who doubles as a professional philosopher that I believe that this sort of “behind the scenes” reflection is an indispensable part not just of doing good mathematics but also of advancing our inquiry, as human beings, into some of these fundamental questions.