Introduction

An Invitation

In many cases, events and objects are given to observation as extended through time and space, and so the data about these is local and distributed in some fashion. For now, we can think of this situation in terms of the data being indexed by, or attached to, given delimited regions or domains of some sensors. In a very general and rough way, by *local* we typically understand that something is being compared to what is around or nearby it; this is as opposed to the *global*, generally understood to mean compared to everything or across an entire domain of interest. Satisfying a property at a local level does not necessarily entail that the same will obtain at the global level. In saying that the data is local, we just mean that it holds throughout, or is defined for, a certain limited region; that is, its validity is restricted to a prescribed region or partial domain or reference context. We also use this language of locality to describe a way of evaluating a property or data ascribed to a part or point of its extended domain in terms of what that property or data looks like viewed from its immediate surroundings—that is, whenever it holds somewhere, it should also hold *nearby*.

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 piece 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 their 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 been given to us only 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, and so on, 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 nonlocality 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™, one considers the assignment 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 (15×15 grid of squares). Yet this assignment of letters to squares 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 subwords 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 piece 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 piece. 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*, that is, 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 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 inabilities 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 seventy-year-old collection of concepts and tools designed by mathematicians to tame and precisely comprehend problems with a structure like the sorts of situations introduced above. I hope the reader will have 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 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 of the space and restriction of the data assigned to the space, 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, that is, an observation that no longer refers merely to yet another local region but now extends over and embraces all the regions at once. As such, it may even enable predictions concerning missing data or at least enable principled comparisons between 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 in this instance to pass from such local observations to the global can tell us about the phenomena at hand. *Sheaf cohomology* is a tool used for capturing and revealing precisely obstructions of this sort.

The very natural distinction between local and global, hinted at above, in fact posits a large class of problems involving relations between the local and global. For instance, given an overall domain of interest, or space, if we consider some part of that, when is it possible, just through knowledge of such portions, to deduce knowledge about the whole domain of interest? Perhaps unsurprisingly, the antagonism between the local and the global found its initial articulation within the frameworks of geometry and topology (the study of space), where there is a very natural account of locality or what it means for something to hold locally. One of the virtues of sheaves and associated techniques (like sheaf cohomology) is to have allowed for an appreciation of how this local-global dialectic is still more universal and reaches beyond its initial appearance in the context of topology and geometry.

The main purpose of this book is to provide an inviting and (I hope) gentle introduction to sheaf theory, where the emphasis is on explicit constructions and applications, using 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. 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 specialized groups of people appears to be growing immensely, but whenever I spoke to newcomers about sheaves, they often expressed 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 preexisting 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 it 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 forward with more confidence but they 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 and discussing some of its unique features, the next section offers a more explicit, 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.

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

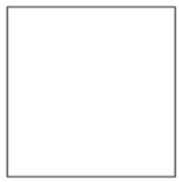

We are less interested in the "space itself" and more in how the space serves as a site where various things *take place*. In other words, we think of this region as really just an abstract domain supporting various *happenings*, where such happenings carry information for appropriate sensors or measuring 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.[1] For instance, the region might 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:

1. The description of sheaves as "measuring instruments" or the "meter sticks" associated to 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 1940s, 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 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. (Grothendieck 1986, 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*. (Grothendieck 1986, 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 in the book, and the notions and results alluded to in the above will be motivated and discussed in detail.

The related "sensor" perspective has been developed more recently, to great effect, in the work of Robert Ghrist, Michael Robinson, and Justin Curry, for example, Curry (2014, chap. 10).

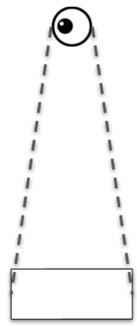

There might then be another sensor, taking in visual information about another region or part of some overall space, offering another "point of view" on another part of the space; and it may be that the underlying regions monitored by the two sensors overlap in part:

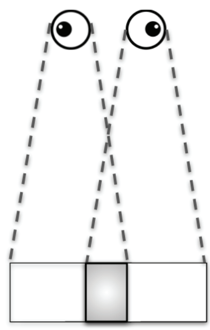

Since we are ultimately interested in the informative happenings attached to the space, we want to see how the distinct perspectives on what is happening throughout the space are themselves related; to this end, a very natural thing to do is to ask how the data collected by such neighboring sensors are related. Specifically, it is very natural to ask whether and how the perspectives are *compatible* on such overlapping subregions, 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):

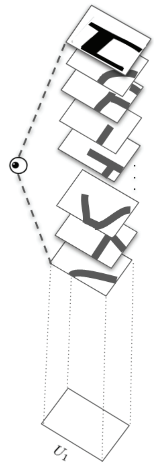

While not really necessary, the sensor might even be supposed to be equipped to process the information it collects, translating such visual inputs into reasonable guesses about which possible capital letter or character the partial sketch is supposed to represent. In any event, attempting to relate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

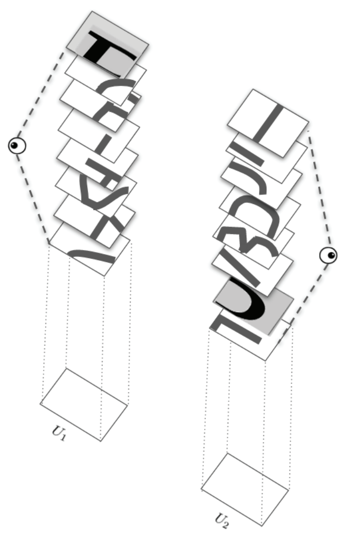

Corresponding to how the underlying regions are naturally related by a relation of inclusion, the compatibility question, undertaken at the level of the selections (highlighted in gray above) from the collections of all informative happenings on the respective regions, will involve looking at whether 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:

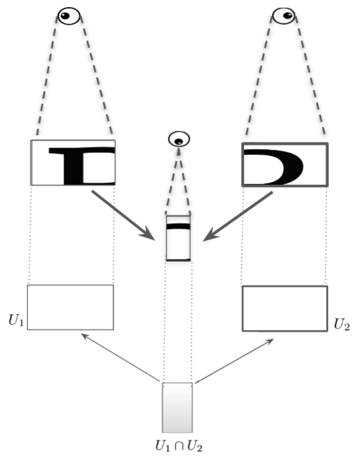

If the given selection from what they individually "see" does match on the overlap, then, corresponding to how the regions U_1 and U_2 may be joined together to form a larger region,

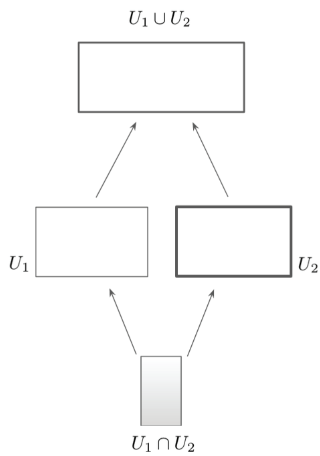

at the level of the data on the happenings over the regions, we can pull this data back into an item of data given now over the entire space $U_1 \cup U_2$, with the condition that we expect that restricting this new, more comprehensive, perspective back down to the original individual regions U_1 and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

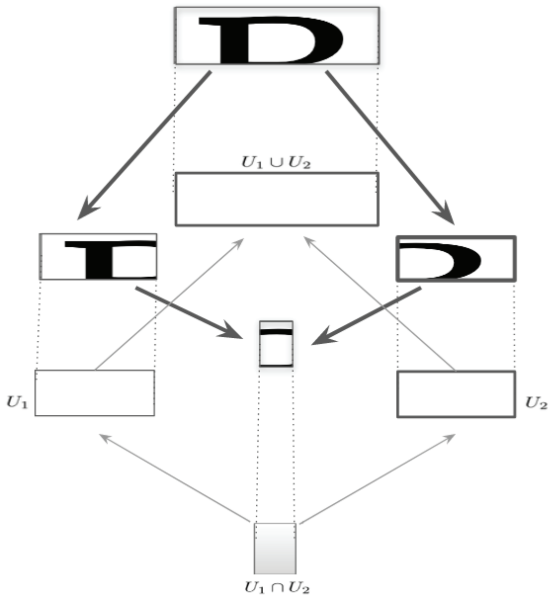

In other words, given some selection from what sensor 1 sees as happening in its region U_1 and from what sensor 2 sees as happening in its region U_2, provided their "story" agrees about what is happening on the overlapping region $U_1 \cap U_2$, then we can paste their individual visions into a single 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" 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 of 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.):

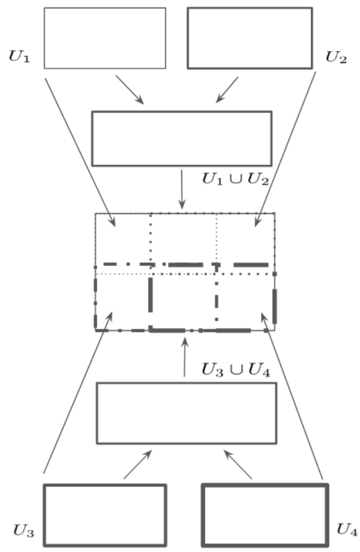

With the four regions U_1, U_2, U_3, and U_4, to each of which there corresponds a particular sensor, we have got the entire region $U = U_1 \cup U_2 \cup U_3 \cup U_4$ "covered." Part of what this means is that, were you to invite *another* sensor to observe the happenings on some further portion of 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:

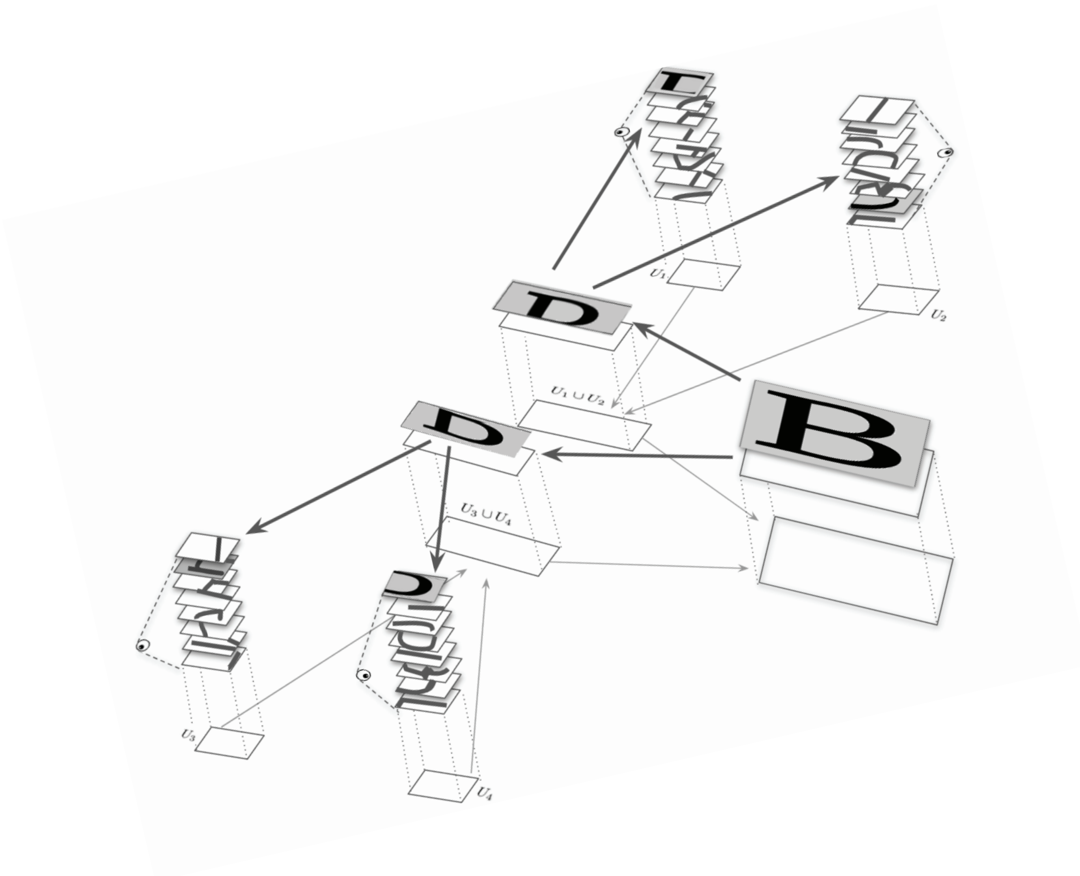

This system of mutually compatible local data assignments or "measurements" of the happenings on the space—where the various data assignments are, piece by piece, constrained by one another, and thereby patched together to supply an assignment over the *entire* space covered by the individual regions—is, in essence, what constitutes our *sheaf*. The idea is that the data assignments are being "tied together" in a natural way

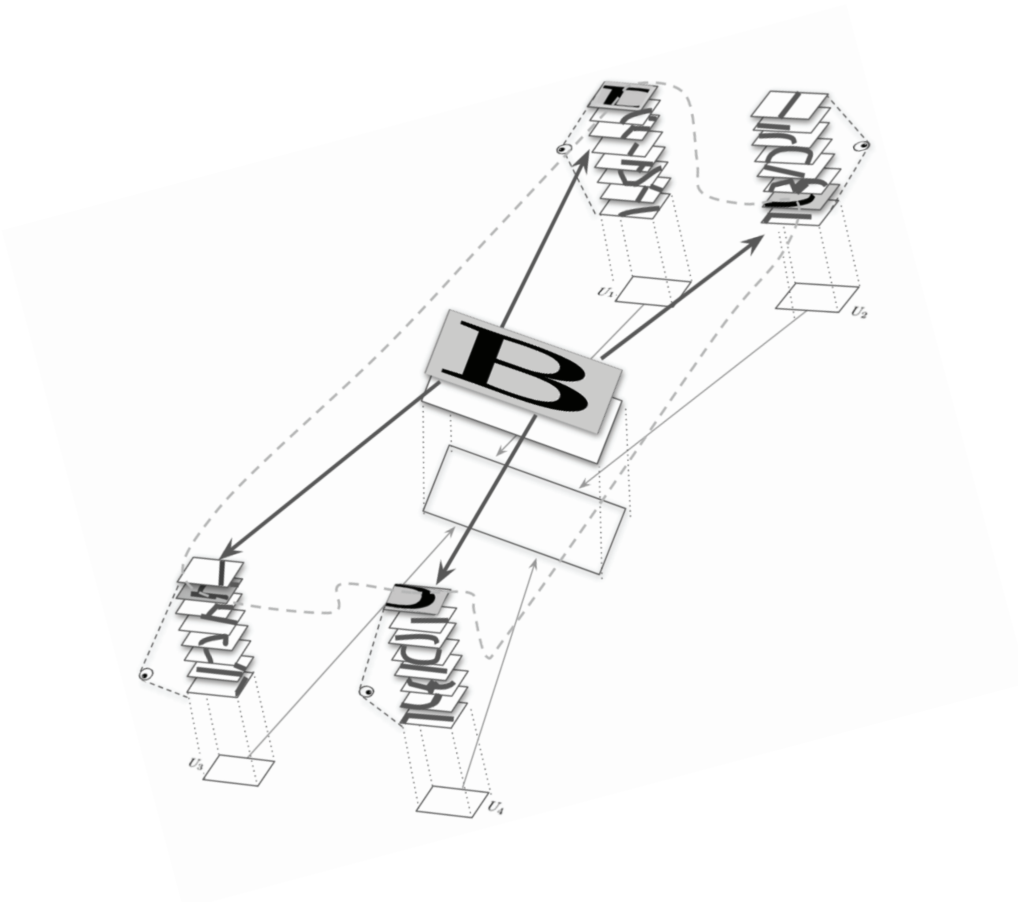

where this last picture is meant to serve as motivation or clarification regarding the agricultural terminology of “sheaf”:

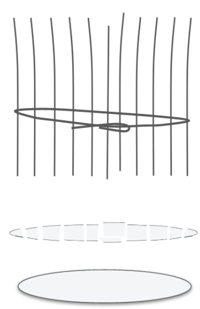

Here one thinks of various regions as the parcels of an overall space covered by those pieces, the collection of which then serves as a site where certain happenings are held to take place, and the abstract sensors capturing local snapshots or measurements of all that is going on in each parcel are then 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) amounts to a cross-section and the process of restriction (along intersecting regions) and collation (along unions of regions) of these sections 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 nonoverlapping 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 together the information supplied by each sensor, 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 system works 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).[2] More generally, image and face recognition appears to operate, in a single brain (where clusters of neurons play the role of individual 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 similar way: a society’s store of knowledge consists of a vast patchwork built up of partial records and data items referring to delimited (possibly overlapping) domains of interest, each of which data items can be (and often are!) checked for compatibility whenever they involve data that refers to, or makes claims about, the same underlying domain.

The very simple and naive presentation given to it above admittedly 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

2. That visual information processing itself seems 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.

was finally named explicitly so that substantial efforts could 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. (Grothendieck 1986, promenade 13)

Outline of Contents

The rest of the book is structured as follows. The first three chapters, together with the sixth and seventh chapters, are dedicated to exposition of the most important category theoretic concepts, tools, and results needed for the 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 theory,[3] 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 these chapters, all the necessary categorical fundamentals are accordingly motivated and developed. The emphasis here, as elsewhere in the book, is on explicit constructions and creative examples. For instance, the concept of an *adjunction*, and key abstract properties of such things, is introduced and developed first through an extended example involving “dilating” and “eroding” an image, then through the development of “possibility” and “necessity” modalities applied to modeling the 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.

Chapter 1 introduces categories, some important examples of categories, and some of what one can do with categories.

Chapter 2 develops functors and presheaves in considerable depth. It discusses four main perspectives on presheaves, works through some notable examples of each of them, 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 chapters,

3. The general reader without much, or any, background in category theory is especially encouraged to have a look at the engaging and highly accessible Spivak (2014). Readers with more prior mathematical experience may find Riehl (2016) a compelling introduction, displaying as it does the ubiquity of categorical constructions throughout many areas of mathematics. Lawvere and Rosebrugh (2003) are 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 rediscovering things for oneself.

especially chapter 5, dedicated to the initial development of the sheaf concept. Natural transformations are also introduced in this chapter.

Chapter 3 covers universal properties and some important universal constructions.

All that is needed to offer a definition of a sheaf is the notion of a presheaf (covered in chapter 2) and some basic notions from topology, such as that of a cover. With the aim of exposing the reader to the sheaf notion sooner rather than later, chapter 4 covers the requisite notions from general topology, and raises some more philosophical questions that are taken up in later parts of the book (including the appendix).

Chapter 5 introduces sheaves (on topological spaces) and some key sheaf concepts and results through some initial 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.

Chapter 6 is dedicated to the Yoneda results—perhaps the most important idea in category theory—and the associated Yoneda philosophy.

Chapter 7 returns to, and completes, the treatment of categorical foundations for sheaves, by covering adjunctions. As usual, the key features of this construction are teased out through a variety of examples and worked-out constructions.

Chapter 8 returns to sheaves and covers some more involved results, rooted in historically significant examples. This chapter also includes a section on what is *not* a sheaf, or when and how the sheaf construction fails, as well as an important case where the notions of sheaf and presheaf coincide.

Chapter 9 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 10 revisits and revises a number of earlier concepts, and develops sheaves from the more general perspective of Grothendieck toposes. The important notions 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") to elementary toposes, the topic of chapter 11. The later sections of chapter 11 are devoted to illustrations, through concrete examples, of some slightly more advanced topos-theoretical notions. The book concludes with an abridged presentation of some special topics, including a brief glimpse into *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 in that final section to the special topic of cohesive toposes.

Finally, there is an appendix, dedicated to exploring in greater depth the open philosophical questions raised in chapter 4 on general topology and the concept of space, doing so by building on some of the constructions introduced in chapter 7's treatment of adjunctions.

Remarks on Distinct Features of This Book

This book has three notable features that may deserve brief discussion:

1. an emphasis on pictures;
2. an emphasis on detailed worked-out examples from different areas of application; and

3. an emphasis on ideas.

Regarding the first of these: a colleague once told me that they had read an entire book on sheaf theory, but it was not until years later, after they saw a simple and evocative picture drawn of a certain sheaf, that they finally felt like they understood what sheaves were about. I suspect that this person is by no means alone in their experience. If this is really a fair description of the experience of some newcomers to sheaves, you could say that, at least as far as sheaves are concerned, a picture is worth not a thousand words but many thousands of words! Inspired by this experience, I have tried to include, throughout the book, a great many pictures.

The second feature of the book is that it takes part in the burgeoning area of *applied category theory*, and as such aims to expand the repertoire of examples of sheaves, beyond those that have already had great impact within mathematics. As in any area of life, there can be a kind of "groupthink" that takes over an academic niche, and examples are usually the first things to suffer the negative consequences of this common phenomenon—for instance, many standard texts on sheaves start with the constant sheaf and then are satisfied to mention a handful of other standard examples and well-established uses within mathematics, before pressing on with abstract results. Especially in recent years, there has been something of a push against this, with a number of exciting new applications of sheaves to topological data analysis,[4] to sensor networks,[5] to opinion dynamics (including selective opinion modulation and lying) on social networks,[6] to target tracking,[7] to dynamical systems and behavior types,[8] to name just a few. This book has been greatly inspired by such efforts.

Regarding the third feature of this book: 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 them 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 millennia, going to the heart of some of the most lasting and knotty questions concerning, for instance:

- What is an object (and can we give an entirely relational account of objects, that is, display an object in terms of all its relations)?
- What is universality?
- What is negation?
- What fundamental notions are codified by our concept of space?

4. As in Curry (2014).
5. As in Robinson (2016b).
6. As in Hansen and Ghrist (2020).
7. As in the work of Robert Ghrist.
8. As in Schultz and Spivak (2017) and Schultz, Spivak, and Vasilakopoulou (2016).

While a number of other issues will be discussed, some of the main philosophical issues that will be explored in the book engage a few decisive dialectics, notably that of the

- local-global,
- continuous-discrete,
- particular-universal, and
- object-relation.

The struggle to articulate the peculiar relations and antagonisms between each of the members of such pairs has been ongoing for centuries, and while mathematics has advanced inquiry into these matters more than any other discipline, it remains the case that there is a great history to investigating such dialectics, and they are not the sole property of mathematics. Occasionally stepping back to ground specialized treatments of these matters in the broader discussion is useful not only for reminding us of some of the stakes of our formalism, but also for connecting the activity of mathematics back to the longer history and future of inquiry, as human beings, into such fundamental questions.

A word about philosophy.[9] Specialization has manifold benefits, and even if it didn't, it seems to be the price we must pay, as beings with very limited resources, for doing something well. At times though—especially during times like ours, an age of increased specialization—the incentive structures for engaging with something outside one's specialization and subspecializations can deteriorate. Whether the thick boundaries of the adult's specialized world have barely been felt by them, or because they still have the luxury of not being overly concerned with the pursuit of excellence, children are good at refreshingly disrespecting the adult's divided world. As we grow out of being a child, those boundaries become more and more real for us, yet most of us (even the hardened specialist) do not really entirely outgrow or utterly forget that state of the child, nor do we ever come to fully believe in the reality of those boundaries. And even if we tell ourselves that we do, the child seems to return, however faintly or mischievously, in unexpected ways. We find ourselves wondering if such a thing as humor can be defined within music in a purely musical way, or if certain growth patterns found throughout nature can tell us something about the impulse movements of financial markets. Through a mixture of curiosity, a drive to unify and organize, or sometimes just a stupid whim, we retain something of this impulse to *take concepts beyond where we are told they belong*. Such inquiries can only be vague and tentative at first, and there is always a risk they will not lead anywhere. Over time, certain inquiries mature and start to appear a little differently to us: we find ourselves seriously considering if there is life on other planets, or how we can get machines to learn complicated behaviors purely using reinforcements built into the environment, the way so many animals do. If we look carefully, even in those more established questions we can still recognize that same childlike impulse to disregard the myriad cues that exert pressure on the questioner to leave concepts where they belong: "'Life' is a *here* thing!" "Learning is something only carbon-based beings can do!"

9. No part of this book rests on the remarks made in the next three paragraphs. They are provided for context, and were prompted by questions I received from separate mathematician colleagues curious about how I, as a professional philosopher, understood "philosophy" and its relation to category theory.

When it works, taking concepts beyond the confines of their native setting can have the effect of attaining greater generality. This impulse to attain greater generality—which is born out of taking concepts beyond where they belong—is the minimal working sense of "philosophical" that I intend in the present context. In this sense, philosophy is something that we all do and that does not at all belong to "the philosopher." And while it is perhaps one of the greatest beneficiaries of the advantages of specialization, mathematics fundamentally shares this same strong drive toward the general—which may in part account for why, throughout the centuries, there has been a great deal of interaction between the disciplines of philosophy and mathematics, even to the point that for much of history it would have been difficult to draw a sharp line between the two. This intimate bond becomes especially evident with category theory. One could argue that, at least in large part, philosophy (in a more traditional sense) has evolved as the informal study of universality (and universal phenomena). One could argue that category theory is the formal study of formal universality. As such, it is no surprise that there appear to be a number of especially strong connections between the matters pursued by category theorists and those of philosophers.

I happen to believe that many of the staple questions that were originally the provenance of the philosopher will eventually be handled with the care they deserve once they are adequately framed as problems within category theory, and that in the near future every major philosophical dialectic—universal-particular, continuous-discrete, global-local, quality-quantity—and even less obvious problems, such as those of "personal identity," will be handed over to, and considerably enriched by, the category theorist. In the other direction, a variety of basic elements of category theory appear to raise philosophical questions of their own, and certain more advanced developments (such as with cohesive toposes, discussed in chapter 11) seem almost inherently philosophical, and poised to attack a number of the traditional philosophical problems. But we are probably at least 100 years away from a world in which one can adequately realize that category theory is everything philosophy ever strove to be, and let that long, rich, and frustrating tradition take on a new form. In the meantime, one of the aims of this book is to encourage those from each camp to engage with the other—and the "Philosophical Pass" sections are opportunities to step back from the formal details, gather our thoughts, relate the mathematical concepts to broader or tangential conceptual developments, and occasionally engage in a little pushing of the formal concepts beyond where they belong.

I would encourage all readers to pursue the philosophical sections of this book—though they are set off in boxes to mark them off from the rest of the text so that the more narrowly focused reader can easily find their way around them should they insist on reading only the mathematics. To encourage the more strictly mathematical reader to engage with those sections, though, I will just add that it seems that nearly all great mathematicians of the past have let themselves be provoked by, and at times have even engaged with, the philosophical dimensions of their work.

Finally, while emphasis on concrete examples from unexpected areas beyond the confines of pure mathematics is already unusual enough for a text on sheaf theory, and while engagement with philosophical dimensions of the mathematics is itself atypical for a primarily mathematical text, the reader might be even more surprised to find these two things paired together. In response to this reaction, let me bastardize the philosopher Kant and say this: knowledge of examples and applications without a sense of the general ideas these exemplify and are powered by is *blind*, while knowledge of general ideas without a familiarity with all sorts of examples and applications is *empty*. Various philosophers of the past, like Aristotle and Spinoza, have set as an ideal for the most demanding and adequate kind of knowledge one that can look past the apparent immediacy of the universal and the particular, taken on their own, and instead achieve a more unified understanding of the subtle mediations between our knowledge of the universal and of the particular. Moreover, it has been my experience that often the only way to really grasp the most general, and to appreciate the various needs to keep pushing things in the direction of the more general, is to sink as deeply as possible into certain particular problems. In a peculiar slogan: often what is furthest (most general) can be most readily approached through closer consideration of what is nearest (least general). In this connection, I believe that the ideal mathematician would represent some sort of fusion between the Grothendieckian impulse toward extreme abstraction and general ideas, on the one hand, and the intimate exploration and care for particulars embodied by the likes of a Ramanujan, on the other. While I would not pretend to achieve anything remotely close to this fusion for myself, I do believe that it is a noble ideal to strive for, and the atypical pairing of engagement with general ideas and respect for examples found in this book has been influenced by that belief.

What the Book Is (and Requirements of the Reader)

I should add a word about what this book aims to be and who it is for. One reviewer of an earlier version characterized the book's most significant contribution as

> providing an accessible sheaf theory book filled with fun examples, with a broad philosophical bent.

I think this is a very clear statement of what I have wanted to achieve with this book. There also happens to be a great gap between the few accessible books on the basic category theory (and other prerequisites) needed to develop sheaves and any currently published book on sheaf theory. Anyone who would find a bridge over that chasm useful, or who would be engaged by a sheaf theory book that meets the above description, will likely find this book valuable.

Realistically, though, anyone who would find their way to this book will likely have some prior mathematical training and interests. The primary audience of this book should include open-minded mathematicians, scientists and engineers with some broader mathematical interests, and mathematically inclined philosophers. Because of the distinctness of these three groups, I highlight, at the beginning of each chapter, the mathematical and philosophical goals and topics explored. As for those with interests of the practical sort: there are a number of examples, constructions, and discussions that should be of interest; however, there may be certain sections (appealing to those with more abstract aspirations or those with a philosophical bent) that might be of less interest to such a reader. Such readers might try dipping their toes into those sections and skimming on first reading, focusing most of their attention on the examples.

As for general requirements of the reader, I have tried to make this book as self-contained as possible and minimize the prerequisites in order to extend the reach as far as possible to nonexperts. I thus assume only some basic familiarity with set theory and mathematical reasoning—all other concepts needed for the formulation and understanding of sheaves, including the basics of category theory, topology, and anything else, are motivated and introduced in this book.

In the end, I have tried to write the book I wish I had when I was first learning sheaf theory. There are some outstanding books on sheaf theory—notably Mac Lane and Moerdijk's *Sheaves in Geometry and Logic*—but such texts can be rather demanding on the beginner, assume a great deal of mathematical maturity, and generally appeal to a rather expert and self-sufficient audience. In this book, I have tried to assume a great deal less than such texts, to engage a broader audience, and generally adopt a more gentle approach.

What the Book Is Not

As one might already imagine, given its unique aims and approach, this book is not meant to be a standard textbook for experts learning about sheaf theory as it is usually taught in one of its specialized contexts, such as algebraic geometry. An expert reader who has certain expectations about what this book should be, based on standard specialized references on sheaves, will surely have those expectations violated.

In this connection, this book deliberately minimizes treatment of applications to problems in algebraic geometry, one of the historical homes of sheaf theory. This was intentional—in part since these applications require a level of mathematical maturity which this book tries not to assume of the reader, in part because there are already many references devoted to sheaves in algebraic geometry. Beyond this, the omission is also somewhat philosophical. Tom Leinster wrote, in 2010, a blog post entitled "Sheaves Do Not Belong to Algebraic Geometry":

> They are, of course, very *useful* in algebraic geometry (as is the equals sign). Also, human beings discovered them while developing algebraic geometry, which is why many of them still make the association. But. . . sheaves are an inevitable consequence of general ideas that have nothing to do with algebraic geometry.

This is a perspective I share, and I have accordingly sought to avoid including applications to algebraic geometry, with the aim of redistributing the somewhat disproportionate control algebraic geometers have taken over these (demonstrably more general and far-reaching) ideas.

This book is also not meant to be a complete reference. This is part of a trade-off one must make when attempting to appeal to, and sustain the interest of, a wider audience of nonexperts. There are a number of additional topics I would have loved to cover, and further examples I would have loved to include, yet doing so with the aim of completeness could have easily made this book extend to over 1,000 pages. It seemed to me more desirable to welcome more newcomers to sheaves with a book of a more manageable size.

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.

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