Sheaves on Topological Spaces and their Logic

Leo Lobski

Year 4 Project
School of Mathematics
University of Edinburgh
6 April 2018

Abstract

We introduce presheaves on a topological space as contravariant functors, and sheaves as presheaves satisfying the gluing axiom (Definition 2.3). It is demonstrated how sheaves can be used to distinguish between local and global properties of structures in a topological space. We proceed to define bundles on a topological space together with their sections, and construct an adjunction between the category of presheaves and the category of bundles. Chapter 2 is concluded with an important result, the equivalence of categories of sheaves and étale bundles.

In Chapter 1, we introduce the notion of an elementary topos as a generalisation of the category of sets. It turns out that the category of sheaves is a topos, which is the object of discussion in Chapter 3. The main aim there is to construct the subobject classifier for sheaves, and to show that the logic it gives rise to is in general non-Boolean. To this end, we define Boolean and Heyting algebras, and demonstrate how any topos gives rise to Heyting algebras of subobjects.

This project report is submitted in partial fulfilment of the requirements for the degree of *BSc Mathematics*.

Contents

Introduction

Suppose you are dropped on a random point on a map. By simply looking around, you can verify that you are indeed on a map representing some part of the planet Earth. However, if someone asks you whether the map you are on is contained within the boundaries of a country[1], you are clueless. In order to find this out you need to take a step back, you need to know something about the entire map rather than what is happening around each point. The difference between 'maps' and 'maps that are contained within a country' can be captured by noticing that maps with no extra requirements satisfy a universal property. Namely, given any collection of maps which agree whenever they describe the same geographical area (that is, the maps use the same projection, are from the same historical period, include the same details and so on), we can construct a new map by 'gluing' these pieces together. Moreover, there is exactly one such map which coincides with each of the maps in the original collection, meaning we have no choice in how to glue the maps. This corresponds to the idea that we can verify that we are on some map by 'looking around'; we can identify a small piece of map around each point, then gluing the pieces together (which necessarily coincide on the overlaps) reproduces the original map. Furthermore, this property fails for maps that are contained within a country. Indeed, consider a map of Haparanda (a border town in Sweden) and that of Tornio (a border town in Finland), while both of these are contained within a country, the glued map of the cross-border agglomeration formed by the two towns is not.

The observations made in the introductory example above lead to the distinction between local and global properties. Since maps with no additional requirements are determined by considering some neighbourhood around each point, we have the intuition that the property of being a map is local. In contrast, since determining whether a map is contained in a country requires 'the whole picture', we would like to say that this property is global. One of the aims of this work is to make this distinction precise by introducing the notion of a sheaf.

After defining a sheaf, we will provide a number of examples illustrating that sheaves do indeed distinguish local and global properties of structures living in a topological space. The canonical example is the sheaf of continuous functions (Example 2.6); if a function is continuous at every point, then it is continuous. This is often contrasted with the presheaf of bounded functions (Example 2.8); a function that is bounded in some open neighbourhood of each point is by no means bounded.

[1] Is there at least one country such that the entire geographical area represented by the map lies inside that country?

We will provide two perspectives on sheaves. The first viewpoint is the one introduced thus far, which is to regard a sheaf as a special kind of functor providing a useful tool for distinguishing between local and global properties. In this guise, we will proceed to show that the category of sheaves on a topological space is equivalent to the category of special kind of maps into that topological space; namely, to the category of local homeomorphisms. This is the subject of Chapter 2. The second perspective is to view sheaves as a structure in its own right. We will see that the category of sheaves is in fact a topos, meaning roughly that sheaves have many of the familiar properties of sets. In particular, this will allow us to talk about Heyting algebras of sheaves in Chapter 3, which in turn provide the appropriate language for propositional logic of sheaves.

I will assume no familiarity with sheaf or topos theory, as we will start the discussion by defining an elementary topos in Section 1.2 and a sheaf on a topological space in Section 2.1. However, I will assume that the reader is familiar with elementary category theory, although nothing beyond a typical introductory course or textbook will be needed. In particular, I will freely use such notions as functors, natural transformations, adjoints, slice categories (which are a special case of comma categories), equivalence of categories, initial and terminal objects, limits and colimits. For an introduction to category theory, the reader is referred to Tom Leinster's *Basic Category Theory* [4], or to Emily Riehl's *Category Theory in Context* [8], which both introduce the subject from the beginning. Another possibility is Saunders Mac Lane's *Categories for the Working Mathematician* [5]. Robert Goldblatt introduces category theory and elementary toposes with no assumed categorical background in his *Topoi, the Categorial Analysis of Logic* [2]; the emphasis there, as the title suggests, is on formulating topos logic.

Despite saying that I will assume familiarity with limits and colimits, Chapter 1 begins with a review of equalisers and pullbacks as well as exponentials. These will be used to define toposes and sheaves, hence the definitions will be stated so that we can refer to them later. We then proceed to define a subobject classifier, which is an important part of the definition of an elementary topos, as it will allow us to define a Heyting algebra in an arbitrary topos. We conclude the first chapter by defining an elementary topos, and making sure that it does indeed generalise the category of sets by checking that **Set** is an elementary topos.

In Chapter 2, we define sheaves and presheaves. We emphasize the perspective that sheaves are functors that distinguish between local and global properties. We then move on to defining bundles over a topological space together with their sections. These are just continuous maps into the topological space (bundles) together with a collection of right inverses for these maps (sections). We explicitly construct an adjunction between the category of bundles and the category of presheaves on a topological space. In the same way that we define a sheaf as special kind of presheaf, where 'special' has something to do with locality, we define a special type of bundle, namely, an étale bundle as a local homeomorphism. The main result of the chapter is the fact that the aforementioned adjunction restricts to an equivalence of categories of sheaves and étale bundles. This is proved in Section 2.3.

Chapter 3 begins by defining a lattice and a Heyting algebra as a property of the lattice. The idea behind a Heyting algebra is to provide a model for (intuitionistic) propositional logic. We show that the law of excluded middle is equivalent to the law of double negation in any Heyting algebra. This will allow us to define a Boolean algebra by imposing the additional requirement on a Heyting algebra that either of these two equivalent laws is satisfied. Whereas a Heyting algebra models intuitionistic logic, a Boolean algebra is a suitable description of classical logic. We will show that Heyting algebras are intimately connected to toposes in the following way. For any object in a topos, a collection of its subobjects forms a lattice. Moreover, this lattice has a Heyting algebra structure. We will show this in Section 3.2.

We will conclude this work by connecting the discussion of algebras in a topos to sheaves. It turns out that both sheaves and presheaves form a topos. We will not prove these results in detail despite the fact that the proofs provide interesting and important insights into the essence of both sheaves and presheaves. The guiding rationale behind the proofs is outlined in Section 3.3. Nonetheless, we will be able to construct a subobject classifier for sheaves, which turns out to be the set of all open sets of the underlying topological space, ordered by inclusion. This will allow us to show that 'the logic of sheaves' is in general not Boolean, hence reasoning in the category of sheaves is intuitionistic rather than classical.

Chapter 1

Preliminaries

In the first section, we remind the reader of three constructions that make the category of sets and functions **Set** special as compared to an arbitrary category. In the second section these constructions will be first used to define a subobject classifier, which in turn is a core ingredient in the definition of an elementary topos. If the reader is comfortable with definitions of an equaliser, exponential and pullback, the first section can be safely skipped and used as a reference when needed later in the discussion. Same applies to the second section if the reader has seen the definition of an elementary topos before.

1.1 Categorical structures

We begin by recalling some categorical constructions needed for the subsequent discussion of toposes and sheaves. Namely, these will be two important types of limits, equaliser and pullback, as well as exponential. All of these are motivated by the same constructions in the category of sets, whose universal properties are then taken as definitions in a general category.

Definition 1.1. (Equaliser) Let $f, g : A \rightrightarrows B$ be a pair of morphisms in some category. An *equaliser* of f and g is a morphism $e : E \to A$ such that $ef = eg$, which is universal with this property in the sense that for any morphism $e' : E' \to A$ with $e'f = e'g$, there is a unique map $u : E' \to E$ such that the diagram

$$\begin{tikzcd}
E' \arrow[d, dashed, "u"'] \arrow[dr, "e'"] & & \\
E \arrow[r, "e"] & A \arrow[r, shift left, "f"] \arrow[r, shift right, "g"'] & B
\end{tikzcd}$$

commutes.

In the category **Set**, an equaliser of f and g is simply the subset

$$E = \{a \in A : f(a) = g(a)\}$$

of A together with the inclusion map $e : E \hookrightarrow A$. The unique map u is then just e' itself with the codomain restricted to E, since we have $e'(E') \subseteq E$.

Definition 1.2. (Exponential) Let X and Y be objects of a category with binary products. An *exponential* of Y and X is an object Y^X together with a morphism $ev : X \times Y^X \to Y$ such that for any morphism $f : X \times Z \to Y$ there exists a unique morphism $\bar{f} : Z \to Y^X$ making the diagram

$$\begin{tikzcd}
X \times Z \arrow[d, dashed, "\mathrm{id}_X \times \bar{f}"'] \arrow[dr, "f"] & \\
X \times Y^X \arrow[r, "ev"] & Y
\end{tikzcd}$$

commute. The morphism $ev : X \times Y^X \to Y$ is called the *evaluation*.

In **Set**, the object Y^X is the set of all maps from X to Y and $ev : X \times Y^X \to Y$ is the evaluation of a map $g : X \to Y$ at $x \in X$ (hence the name). The unique map $\bar{f}$ then sends $z \in Z$ to

$$f(\text{-}, z) : X \to Y : x \mapsto f(x, z).$$

The last concept of this section is that of a pullback. This will be used in the next section to define a subobject classifier, and will be very important in Chapter 3 when we will define a logic in a topos. An extremely vague description of a pullback, which is nonetheless useful for intuition and memorisation, is to say that it is 'the universal commutative square'. Here are the details.

Definition 1.3. (Pullback) A commutative square

$$\begin{tikzcd}
A \arrow[r, "f^*"] \arrow[d, "g^*"'] & B \arrow[d, "g"] \\
C \arrow[r, "f"] & D
\end{tikzcd}$$

in some category is a *pullback* if for any pair of maps $h : X \to B$ and $k : X \to C$ such that the outer 'square' in the diagram

$$\begin{tikzcd}
X \arrow[dr, dashed, "u"] \arrow[drr, bend left, "h"] \arrow[ddr, bend right, "k"'] & & \\
& A \arrow[r, "f^*"] \arrow[d, "g^*"'] & B \arrow[d, "g"] \\
& C \arrow[r, "f"] & D
\end{tikzcd} \tag{1.4}$$

commutes (i.e. $gh = fk$), there exists a unique morphism $u : X \to A$ such that the diagram above commutes.

Namely, it is sufficient that $g^*u = k$ and $f^*u = h$ for the diagram to commute, as everything else commutes by assumptions. We often say that f^* is the pullback of f along g, and similarly for g^*.

In **Set**, a pullback of $f : C \to D$ and $g : B \to D$ is the subset

$$A = \{((c, b) \in C \times B : f(c) = g(b)\}$$

of the product $C \times B$. The maps g^* and f^* are then the inclusion $A \hookrightarrow C \times B$ composed with the projection map from $C \times B$ to C and B, respectively. The unique map u is the product map $k \times h : X \to C \times B$ with the codomain restricted to A, since we have

$$\{(c, b) \in C \times B : c \in \text{im } k \text{ and } b \in \text{im } h\} \subseteq A$$

by commutativity of the outer square in (1.4).

In a general category, equalisers, exponentials and pullbacks are not guaranteed to exist unless we make further assumptions about the category in question. However, if they do exist, they are unique up to a unique isomorphism. For equalisers and pullbacks, this can either be shown directly, or it follows from a general result for limits; see, for example Proposition 3.1.7 in Riehl [8, p. 76]. For exponentials, this is shown using universality of the unique map $\bar{f}$; see e.g. Exercise 6.1 in McLarty [7, p. 65].

The next two results will both serve as an illustration of how pullbacks work and will be used later.

Proposition 1.5. *The pullback of a monic is monic.*

Proof. Suppose that the diagram

$$\begin{tikzcd} A \arrow[r, "f^*"] \arrow[d, "g^*"'] & B \arrow[d, "g"] \\ C \arrow[r, "f"] & D \end{tikzcd}$$

is a pullback in an arbitrary category, and further that f is a monic. We need to show that f^* is a monic. Hence suppose that $v, w : E \rightrightarrows A$ are morphisms such that $f^*v = f^*w$. Then $gf^*v = gf^*w$, which implies by commutativity of the square that $fg^*v = fg^*w$. Since f is monic, we have $g^*v = g^*w$. We thus have maps $k \coloneqq g^*v : E \to C$ and $h \coloneqq f^*v : E \to B$ such that the outer square in (1.4) commutes. Hence there is a unique $u : E \to A$ such that $f^*u = h$ and $g^*u = k$. But $h = f^*v = f^*w$ and $k = g^*v = g^*w$, whence $v = u = w$ by uniqueness of u. $\square$

Lemma 1.6. *(The pullback lemma) Suppose that the diagram*

$$\begin{tikzcd} \bullet \arrow[r] \arrow[d] & \bullet \arrow[r] \arrow[d] & \bullet \arrow[d] \\ \bullet \arrow[r] & \bullet \arrow[r] & \bullet \end{tikzcd}$$

in an arbitrary category commutes, and that the square on the right is a pullback. Then the left square is a pullback if and only if the outer rectangle is a pullback.

The proof of the pullback lemma consists of repeated (mechanical) applications of the universal property of pullbacks and is not particularly illuminating. We thus omit the proof here; however, it can be found as Theorem 4.8 in McLarty [7, p. 45].

The last definition of this section is that of a poset. It may seem somewhat out of place, as it has a set-theoretic rather than category-theoretic flavour. We will, however, give a way to interpret a poset as a category, an observation which will be used in both subsequent chapters.

Definition 1.7. (Poset) A *partially ordered set*, or a *poset*, is a set P equipped with a relation '$\leq$' which is reflexive, transitive and satisfies

$$\text{if} \quad x \leq y \quad \text{and} \quad y \leq x, \quad \text{then} \quad x = y,$$

for all $x, y \in P$.

Any poset $(P, \leq)$ can be viewed as a category if we declare that there is a morphism from $x \in P$ to $y \in P$ if and only if $x \leq y$. A typical example is a power set of an arbitrary set, where the subsets (elements of the power set) are ordered by inclusion.

1.2 Elementary toposes

In this work, we will take the perspective that a topos is a generalised category of sets and functions. The definition and all the constructions will thus mimic those in **Set**. We have already started applying this viewpoint in the previous section, by connecting each of the concepts introduced there to the analogous concepts in **Set**. It should be emphasized that this is only one of myriads of perspectives to topos theory. For a brief introduction to some of these, the reader is referred to Leinster [3]. We will next introduce subobject classifier needed for the definition of a topos.

The notion of a subobject generalises the notions of subset, subgroup, subring, subspace etc; in other words, any object, which is in some sense contained in another object of the same type and inherits its structure. Let $\mathcal{C}$ be a category and C an object of $\mathcal{C}$. A *subobject* of C is an isomorphism class of monics $m : D \rightarrowtail C$ in the slice category $\mathcal{C}/C$. In detail, two monics $m : D \rightarrowtail C$ and $n : E \rightarrowtail C$ belong to the same subobject if there exists an isomorphism $f : E \xrightarrow{\sim} D$ (in $\mathcal{C}$) such that the diagram

$$\begin{tikzcd}
E \arrow[rr, "\sim", "f"'] \arrow[dr, tail, "n"'] & & D \arrow[dl, tail, "m"] \\
& C &
\end{tikzcd}$$

commutes. Although strictly speaking a subobject is an isomorphism class of monics, it is common to refer to a particular choice of a monic $m : D \rightarrowtail C$ as 'a subobject of C'. We will adopt this conventional abuse of language.

Definition 1.8. (Subobject classifier) Let $\mathcal{C}$ be a category with a terminal object 1. A *subobject classifier* (or a *truth-value object*) of $\mathcal{C}$ is an object Ω together with a morphism $\mathtt{t} : 1 \to \Omega$ such that for every monic $m : D \rightarrowtail C$ in $\mathcal{C}$ (subobject of C) there is a unique morphism $\chi_m : C \to \Omega$ making the following diagram into a pullback.

$$\begin{tikzcd}
D \arrow[r, tail, "m"] \arrow[d] & C \arrow[d, "\chi_m"] \\
1 \arrow[r, "\mathtt{t}"] & \Omega
\end{tikzcd}$$

The morphism $\mathtt{t}$ is often called $\mathtt{true}$ or $\top$, and the morphism χ_m is called the *characteristic morphism* of m.

If 0 is the initial object of some category, and C is any object in the same category, we will denote the unique map from 0 to C by 0_C. If, in addition to a terminal object and a subobject classifier, a category also has an initial object, we define the map $\mathtt{f} : 1 \to \Omega$, called $\mathtt{false}$, as the unique morphism such that

$$\begin{tikzcd}
0 \arrow[r] \arrow[d] & 1 \arrow[d, "\mathtt{f}"] \\
1 \arrow[r, "\mathtt{t}"] & \Omega
\end{tikzcd}$$

is a pullback. That is, $f = \chi_{0_1}$, the characteristic morphism of the unique map from 0 to 1. The morphism $\mathtt{false}$ is sometimes denoted by $\perp$.

In **Set**, the subobject classifier is the pair $(2, \mathtt{t})$, where $2 = \{0, 1\}$ is the two-element set, and $\mathtt{t}$ maps the one element set to $1 \in 2$. The choice of the image of the $\mathtt{true}$ map as 1 indicates that we think of 1 as being the truth-value standing for 'true', and correspondingly we think of 0 as 'false', the image of the map $\mathtt{false}$. Since the subobjects in **Set** correspond to (isomorphism classes of) subsets of C, the characteristic morphism χ_m is the characteristic function of the image $\operatorname{im} m$.

Definition 1.9. (Cartesian closed category) A category is said to be *Cartesian closed* if it has products of any finite number objects, and exponentials of any two objects.

Note that a Cartesian closed category necessarily has a terminal object, as it is the empty product, that is, the product of no objects.

Definition 1.10. (Elementary topos) An *elementary topos* (or simply *topos*) is a Cartesian closed category with finite limits and a subobject classifier.

Since product is a particular type of a limit, a category having all finite limits already implies that it has all finite products. Hence a more economical list of properties defining a topos without reference to Cartesian closed categories is as follows. A topos is a category which has

i. all finite limits,

ii. exponentials for any two objects, and

iii. a subobject classifier.

We can simplify this even further with the following result.

Proposition 1.11. *A category has all finite limits if and only if it has all finite products and equalisers.*

The 'only if' direction is immediate, since products and equalisers are special cases of limit. The 'if' direction is the Corollary 1 in section V.2 in Mac Lane [5, p. 109], together with the observation that a terminal object is the product of no objects.

Corollary 1.12. *A category is a topos if and only if it has finite products, equalisers, exponentials for any two objects, and a subobject classifier.*

We should check that **Set** has all of these properties. We already saw how to form equalisers of any two functions and exponentials of any two sets, we know how to form a product of a finite number of sets, and we have constructed the subobject classifier of sets. It then follows by the above corollary that **Set** is a topos. Hence topos is indeed a generalisation of **Set**.

It follows from the definition of a topos that any topos has all the constructions defined in this chapter. There is one further, rather non-obvious consequence of the definition which we are going to need. Namely, every topos has all finite colimits [3, p. 6]. This in particular implies that every topos has an initial object, that is, the coproduct of no objects.

Chapter 2

Sheaves on Topological Spaces

Our first perspective on sheaves is to view them as a tool used to differentiate between local and non-local properties of structures in a topological space. The first goal of this chapter is to make the previous sentence precise. The second goal is to express the category of sheaves as something purely topological, namely, we will prove it is equivalent to the category of local homeomorphisms (étale bundles).

The category of sheaves is defined as a subcategory of presheaves, which is itself a functor category (a category whose objects are functors and morphisms natural transformations), while the category of local homeomorphisms is a subcategory of the slice category $\mathbf{Top}/X$, where X is a topological space. Hence the equivalence of these categories shows that the notion of a sheaf on a topological space is natural in the sense that it corresponds to a continuous map into the space, which is a local homeomorphism as defined in Section 2.3. In order to get to this equivalence, we will need to define bundles over X, which is just a shorthand name for a continuous map into X, as well as sections of a bundle, which are just the right inverses of the map.

2.1 Presheaves and sheaves

A topology on a set X is most commonly specified by describing which subsets of X are open. Furthermore, each open subset is a topological space in its own right, inheriting the subspace topology. A map of topological spaces $f : X \to Y$ can be restricted to any subset of X. In this way, f defines a family of maps of topological spaces by restriction $f|_U : U \to Y$ for each open subset U of X. The notion of a presheaf generalises these observations. We further note that this process can be reversed, that is, given a collection of open subsets U_i of X and maps $f_i : U_i \to Y$, there is a unique map having the union of the collection $\bigcup_i U_i$ as its domain such that its restriction to each U_i gives back f_i. This observation leads to the definition of a sheaf. While reading this section, the reader is encouraged to return to the non-mathematical motivation involving we began the Introduction with.

In its full generality, an $\mathcal{S}$-*valued presheaf* on a category $\mathcal{C}$ is a functor,

$$P : \mathcal{C}^{op} \to \mathcal{S}.$$

In this work, however, we will be only considering **Set**-valued sheaves and presheaves over a topological space X. More precisely, given a topological space X, let $\mathcal{O}(X)$ denote the poset category of open subsets of X. A presheaf is then a functor

$$P : \mathcal{O}(X)^{op} \to \mathbf{Set}.$$

An element $s \in P(V)$ is called a *section* of P over $V \subseteq X$. Note that whenever we have $U, V \in \mathcal{O}(X)$ with $U \subseteq V$, we have a map $P(V) \to P(U)$ in **Set**. We will denote the image of a section $s \in P(V)$ under this map $s|_U$ and call it the *restriction* of s to U. The reason for this terminology will become apparent from the following example, which we already referred to when motivating the definition.

Example 2.1. Let X and Y be any topological spaces. For $U \in \mathcal{O}(X)$, let $P(U) = Y^U$ denote the set of all functions from U to Y; and for $U \subseteq V$ and $f \in P(V)$, let $f|_U$ be the actual restriction of $f : V \to Y$ to the subset U. Then P so defined is a presheaf.

To illustrate that presheaf is indeed more general than the motivating example, consider the following.

Example 2.2. (A constant presheaf) Let X be a topological space and let A be a set. The *constant presheaf* with value A sends each open set of X to A. The sections are just the elements of A, and the restriction is necessarily the identity on A.

Note that from functoriality of P it follows that for $U, V, W \in \mathcal{O}(X)$ with $U \subseteq V \subseteq W$ and $s \in P(W)$, we have $(s|_V)|_U = s|_U$, analogously to the case with the actual restriction map. Hence a presheaf takes some structure living in the topological space and isolates it into sets, remembering when it makes sense to restrict the structure to a smaller set.

We are now ready to define a sheaf on a topological space.

Definition 2.3. (Sheaf) A presheaf $P : \mathcal{O}(X)^{op} \to \mathbf{Set}$ is a *sheaf* if for any family $(U_i)_{i \in I}$ of open subsets of X, and for any family $(s_i)_{i \in I}$ such that $s_i \in P(U_i)$ for each $i \in I$ with the property

$$s_i|_{U_i \cap U_j} = s_j|_{U_i \cap U_j} \quad \forall i, j \in I, \tag{2.4}$$

there exists a unique $s \in P\left(\bigcup_i U_i\right)$ such that

$$s|_{U_i} = s_i \quad \forall i \in I.$$

In words, whenever the elements in the family $(s_i)_{i \in I}$ are consistent in the sense that they agree on the intersections $U_i \cap U_j$, we require that they can be 'pasted together' to give an element s in the union $P\left(\bigcup_i U_i\right)$, which restricts back to the original elements s_i. This condition is often referred to as *gluing*. Moreover, we require that there is precisely one way of doing this, that is, if there is some element $t \in P\left(\bigcup_i U_i\right)$ such that $t|_{U_i} = s_i$ for each $i \in I$, then $t = s$. This definition can be reformulated in a very concise way.

Proposition 2.5. *A presheaf P on a topological space X is a sheaf if and only if the following diagram is an equaliser for any family of open sets $(U_i)_{i \in I}$, writing $U = \bigcup_i U_i$.*

$$P(U) \xrightarrow{\ e\ } \prod_i P(U_i) \underset{q}{\overset{p}{\rightrightarrows}} \prod_{i,j} P(U_i \cap U_j)\ ,$$

where for $s \in P(U)$ and for $(s_i)_{i \in I} \in \prod_i P(U_i)$,

$$\begin{aligned} e(s) &= (s|_{U_i})_{i \in I}, \\ p((s_i)_{i \in I}) &= (s_i|_{U_i \cap U_j})_{i,j \in I}, \\ q((s_i)_{i \in I}) &= (s_j|_{U_i \cap U_j})_{i,j \in I}. \end{aligned}$$

Proof. First note that the definitions of e, p and q readily imply that $pe = qe$. We thus need to show that e being universal with this property is equivalent to P being a sheaf as in Definition 2.3. Let Z be a set and let $f : Z \to \prod_i P(U_i)$ be a map such that $pf = qf$. For $z \in Z$, let us write

$$fz = ((fz)_i)_{i \in I} \in \prod_i P(U_i).$$

Note that $pf = qf$ is equivalent to $(fz)_i|_{U_i \cap U_j} = (fz)_j|_{U_i \cap U_j}$ for all $i, j \in I$. Now P is a sheaf if and only if for any such Z and f, and any $z \in Z$, there exists a unique $s \in P(U)$ with $s|_{U_i} = (fz)_i$ for each $i \in I$, or equivalently, $es = fz$ (the 'if' direction follows by taking Z to be the one element set, and noting that f then becomes a choice of an element of $\prod_i P(U_i)$). By uniqueness of such s, this is equivalent to existence of a unique map $\bar{f} : Z \to P(U)$ with $e\bar{f} = f$, and hence to universality of e. □

A structure on a topological space forms a sheaf if the structure is in some sense local (later we will in fact take this as the definition of 'locality'). Since functions are determined by their value at each point, the presheaf of Example 2.1 is in fact a sheaf. What this is saying is something rather trivial, namely, if we have a family of functions defined on some pieces of the space U_i, and the functions agree on all the intersections of these pieces, then there is exactly one way to produce a function on the union of U_i's agreeing with all the functions in the family. To make this slightly less trivial, consider the following examples.

Example 2.6. Define the presheaf P as in Example 2.1, except that require all functions to be continuous. That is, each open $U \subseteq X$ is mapped to $P(U) = C(U, Y)$, the set of continuous functions from U to Y. To see that the sheaf condition is satisfied, let $(f_i : U_i \to Y)_{i \in I}$ be a family of continuous functions, for $U_i \subseteq X$ open, such that $f_i|_{U_i \cap U_j} = f_j|_{U_i \cap U_j}$ for all $i, j \in I$. We have already noted that as a map of sets, there exists a unique function $f : \bigcup_i U_i \to Y$ with $f|_{U_i} = f_i$ by simply defining $f(u) = f_i(u)$ whenever $u \in U_i$. Moreover, f is continuous, as for any open $V \subseteq Y$ we have $f^{-1}(V) = \bigcup_i f_i^{-1}(V)$, which is open.

A similar argument shows that the functions still form a sheaf if we replace 'continuous' with 'differentiable', and the codomain Y with $\mathbb{R}$, or some other space where differentiation makes sense. This works since differentiability is defined locally (i.e. in some neighbourhood of each point), if each f_i is differentiable, then so is f. We can restrict this even further.

Example 2.7. Define the presheaf $P : \mathcal{O}(\mathbb{R})^{op} \to \mathbf{Set}$ by mapping any open $I \subseteq \mathbb{R}$ to the set of all analytic functions from I to $\mathbb{R}$. Then P is a sheaf. This follows, again, since analyticity is defined locally around each point.

Analogously, functions with any property that is defined locally in this sense will give rise to a sheaf. Some examples include n-times differentiable functions for any positive integer n and complex analytic functions, as discussed in sections 1.1 and 2.2 in Tennison [9, p. 2, p. 17].

The previous examples have all been functions with some additional requirement. This prompts the question; is the sheaf condition so weak that functions with any additional constraint will satisfy it? The following examples illustrate that this is not the case.

Example 2.8. Let the presheaf P over $\mathbb{R}$ map an open set I to the set of all bounded functions form I to $\mathbb{R}$. We claim that P is not a sheaf. To prove this, it is sufficient to provide a counterexample. Hence let $I_n = (n-1, n)$ for each $n \in \mathbb{N}$ and define $f_n(x) = n$ for each $x \in (n-1, n)$ (see Figure 2.1). Now each f_n is a constant function and hence bounded. Moreover, the functions f_n (trivially) agree on the intersections of the I_n's. However, if we define a function f from $\bigcup_n I_n = \mathbb{R}_+ \setminus \mathbb{N}$ to $\mathbb{R}$ by letting $f(x) = f_n(x)$ whenever $x \in I_n$, it is not bounded, as given any $R \in \mathbb{R}_+$ we can find an $n \in \mathbb{N}$ so that

$$f(x) = f_n(x) = n > R.$$

Thus $f \notin P\left(\bigcup_n I_n\right)$.

Example 2.9. Let P be the presheaf over $\mathbb{R}$ taking $I \in \mathcal{O}(\mathbb{R})$ to the set of uniformly continuous functions from I to $\mathbb{R}$. Note that the same counterexample as in Example 2.8 works here, each section f_i is uniformly continuous, while f fails to be uniformly continuous. Here is another example. Let $I_n = (n-1, n)$ and $J_n = (n - \frac{1}{2}, n + \frac{1}{2})$ for each $n \in \mathbb{N}$. Define the families of functions $f_n : I_n \to \mathbb{R}$ by

$$f_n(x) = (2n-1)x - n(n-1),$$

and $g_n : J_n \to \mathbb{R}$ by

$$g_n(x) = \begin{cases} n^2 & \text{if } x \in \mathbb{N}, \\ f_n(x) & \text{if } x \in J_n \cap I_n, \\ f_{n+1}(x) & \text{if } x \in J_n \cap I_{n+1}. \end{cases}$$

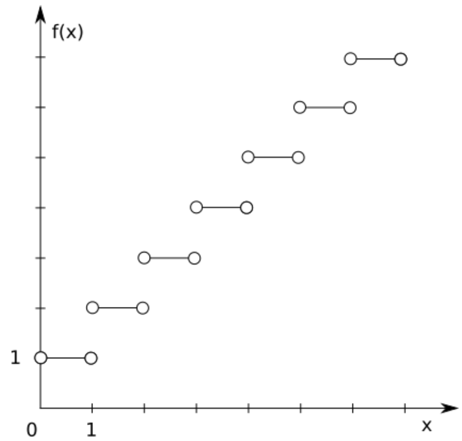

Figure 2.1: The function in example 2.8 is not bounded.

Now $\bigcup_n (I_n \cup J_n) = \mathbb{R}_+$, and (f_n, g_n) agree on the intersections and are all uniformly continuous. The functions f_n and g_n are so defined that the function $f : \mathbb{R}_+ \to \mathbb{R}$ is obtained by connecting each point of the form (n, n^2) to $(n+1, (n+1)^2)$ by a straight line (see Figure 2.2). Such f is not uniformly continuous. To see this, take $\epsilon = 1$, then for any $\delta > 0$, choose an $n \in \mathbb{N}$ such that $n > \frac{1}{\delta} + \frac{1}{2}$. Now letting $x = n$ and $y = n + \frac{\delta}{2}$, we have $|y - x| < \delta$, but

$$\begin{aligned}
|f(y) - f(x)| &= f(y) - g_n(x) \\
&\geq f_n(y) - g_n(x) \qquad \text{(as } f_i \text{ are increasing)} \\
&= (2n - 1)\left(n + \frac{\delta}{2} - n\right) \\
&> \frac{2}{\delta} \cdot \frac{\delta}{2} \\
&= 1.
\end{aligned}$$

Hence for no δ is $|f(y) - f(x)|$ bounded by $\epsilon = 1$, and so f is not uniformly continuous.

Sections of a presheaf do not necessarily need to be functions, as in the following example.

Example 2.10. Let X be a topological space. Let us fix some function $g : X \to \mathbb{R}$. Define a presheaf P by mapping each open $I \subseteq X$ to the set of all sequences of functions on I converging uniformly to $g|_I$. By restriction of sections in this case we mean the restriction of each element (i.e. each function) in the sequence. As one may anticipate, because of the uniformity, P is not a sheaf. The example we provide to demonstrate this is very similar to the one for bounded functions in Example 2.8. Take $X = \mathbb{R}$ and g to be identically equal to zero. Let $I_n = (n-1, n)$ for each $n \in \mathbb{N}$ and define a sequence $(f_{n_i})_{i \in \mathbb{N}}$ by $f_{n_i}(x) = \frac{n}{i}$ for $x \in I_n$. Hence each f_{n_i} is a sequence of constant functions starting at n and approaching zero, that is, $f_{n_i} \to g|_{I_n}$ uniformly for each n. Now the sequence f_i defined by $f_i(x) = f_{n_i}(x)$ for $x \in I_n$ looks like the infinite steps in Figure 2.1, which approach zero as $i \to \infty$. However, the convergence is not uniform, as the steps are unbounded.

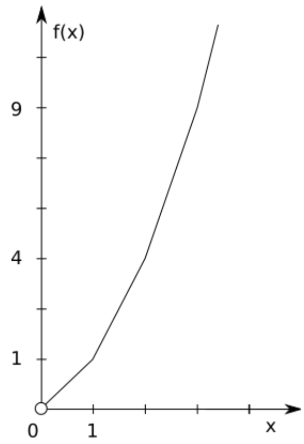

Figure 2.2: The function in example 2.9 is not uniformly continuous.

The previous example is particularly illuminating in highlighting the difference between local and global properties. If we do not require the convergence to be uniform but include the sequences which converge pointwise, P as defined above becomes a sheaf. Hence pointwise convergence is a local property, whereas uniform convergence a global one. The meaning of local and global can thus be made precise using sheaves. We say that some structure is *local* when the presheaf of that structure is in fact a sheaf, otherwise the structure is global. We can thus conclude that continuity, differentiability and analyticity are local properties of functions in this precise sense, while uniform continuity and boundedness are global properties.

2.2 Bundles and sections

It turns out there is a way to map an arbitrary presheaf on X to a continuous function into X. Likewise, an arbitrary continuous function gives rise to a presheaf, which moreover turns out to be a sheaf. Combining these two results, we get a procedure mapping each presheaf to a sheaf. Here we will construct these maps in full detail.

For a topological space X, we write $\mathbf{PSh}(X)$ for the category whose objects are all presheaves on X and morphisms are natural transformations of functors. Similarly, we write $\mathbf{Sh}(X)$ for the full subcategory of $\mathbf{PSh}(X)$ of sheaves on X.

Definition 2.11. (Bundle) A *bundle* over a topological space X is an object in the slice category $\mathbf{Top}/X$.

Explicitly, a bundle over X is nothing but a continuous function $p : Y \to X$, where Y is a topological space. Hence the notion of a bundle is just a change of terminology, instead of specifying a topological space and a continuous map from that space to X, we focus on maps to X as objects. This shift of perspective turns out to be so important that we denote the category of bundles over X by $\mathbf{Bund}(X)$, where the morphisms are simply the morphisms in $\mathbf{Top}/X$, hence $\mathbf{Bund}(X) \coloneqq \mathbf{Top}/X$. Note that id_X is terminal in $\mathbf{Bund}(X)$. A bundle p is sometimes called projection; the reason for this terminology is explained by the following definition.

Definition 2.12. (Section) A *(global) section* of a bundle $p : Y \to X$ is a morphism s from id_X to p in $\mathbf{Bund}(X)$, as in the diagram.

$$\begin{tikzcd}
X \arrow[rr, "s"] \arrow[dr, "\mathrm{id}_X"'] & & Y \arrow[dl, "p"] \\
& X &
\end{tikzcd}$$

That is, a section is a continuous function $s : X \to Y$ with $ps = \mathrm{id}_X$.

If we think of the topological space Y as living 'above' X, then s can be thought of as a choice of a subspace of Y. Take, for example, X to be the unit circle and Y the cylinder of unit radius. Then the image of s is exactly a cross-section of the cylinder, so that p projects it back to the circle, as in Figure 2.3. For this reason, sections of a bundle p are sometimes called *cross-sections* (e.g. in [6, ch. II]).

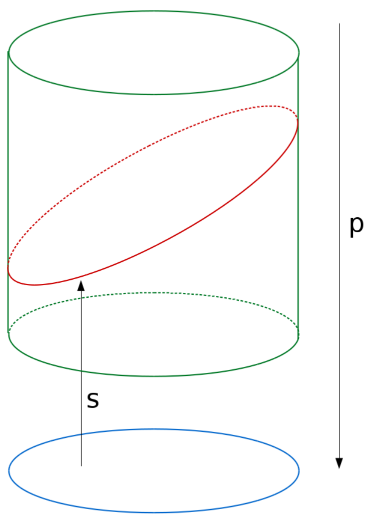

Figure 2.3: The bundle p projects the cylinder onto the the circle; s is one of its sections.

It is useful to allow sections to be defined on some open subset of X only. In detail, if $U \subseteq X$ is open, the bundle $p : Y \to X$ restricts to a bundle

$$p_U : p^{-1}(U) \to U.$$

A section of p_U (or a section of p over U) is then a morphism s from $i : U \hookrightarrow X$ (inclusion) to p in $\mathbf{Bund}(X)$, that is, a continuous function $s : U \to Y$ such that $ps = i$. A section of p_U is called *local*, however, we will not explicitly distinguish between local and global sections unless there is a risk of confusion.

Recall that elements of a presheaf $s \in P(U)$ are also called sections (of P over U). We will next justify this clash of terminology. Given a bundle $p : Y \to X$, we can obtain a sheaf as follows. For an open subset $U \subseteq X$, let $\Gamma_p(U)$ be the set of all sections of p_U. The restriction $\Gamma_p(U) \to \Gamma_p(V)$ for $V \subseteq U$ is given by the usual restriction of functions. Hence Γ_p defines a presheaf over X whose sections over U are exactly the sections of the bundle p over U.

Lemma 2.13. *The presheaf Γ_p is a sheaf for all bundles $p \in \mathbf{Bund}(X)$, called the* sheaf of sections *of p.*

Proof. If U is an open subset of X, the only conditions defining a section s of p over U are continuity and $ps = i_U$ (writing i_U for the inclusion $U \hookrightarrow X$). It follows that for any family $(s_i \in \Gamma_p U_i)_{i \in I}$ where s_i's agree on the intersections of U_i's, the unique function $s : \bigcup_i U_i \to Y$, defined by $x \mapsto s_i(x)$ whenever $x \in U_i$, is continuous and $ps = i_{\bigcup U_i}$. Thus, Γ_p is indeed a sheaf. $\square$

Moreover, the assignment Γ_p is functorial in p.

Lemma 2.14. *The assignment to each bundle of its sheaf of sections is a functor*

$$\Gamma : \mathbf{Bund}(X) \to \mathbf{Sh}(X). \tag{2.15}$$

Proof. Any map $k : p \to p'$ of bundles in $\mathbf{Bund}(X)$ induces a map of sections

$$\Gamma k : \Gamma_p(U) \to \Gamma_{p'}(U)$$

by $\Gamma k(s) = ks$ natural in U, that is, the diagram,

$$\begin{tikzcd}
\Gamma_p(U) \arrow[r, "\Gamma k"] \arrow[d] & \Gamma_{p'}(U) \arrow[d] \\
\Gamma_p(V) \arrow[r, "\Gamma k"] & \Gamma_{p'}(V)
\end{tikzcd}$$

where the vertical maps are the restrictions, commutes whenever $V \subseteq U$. Hence Γk is a morphism of sheaves (i.e. a natural transformation); we have thus indeed defined a functor sending each bundle to its sheaf of sections. $\square$

We immediately ask if we can reverse this process, that is, does there exist a functor sending sheaves to bundles? This indeed turns out to be the case, although the construction required is not quite as straightforward as that for Γ. To this end, we will need to define a germ of a section, which, intuitively, captures the idea of two sections being 'locally the same'.

Let P be a presheaf over X and let x be a point in X. If U and V are some open neighbourhoods of x, we say that the sections $s \in P(U)$ and $t \in P(V)$ *have the same germ* at x if there exists an open neighbourhood W of x such that $W \subseteq U \cap V$ and $s|_W = t|_W \in P(W)$, in which case we write sG_xt. The following example illustrates that, for the presheaf of functions, it need not to be the case that the functions f and g have the same germ at x if they agree on x, the converse is, however, manifestly true.

Example 2.16. Consider the presheaf of functions from $\mathbb{R}$ to $\mathbb{R}$. Let g be the function which is identically zero, and let $f(x) = 0$ for $x \leq 0$ and $f(x) = 1$ for $x > 0$. Then f and g have the same germ at x for any $x < 0$, but not at $x = 0$, as any open neighbourhood of 0 contains points where $f \neq g$.

Lemma 2.17. *The relation $sG_x t$ is an equivalence relation on sections over open neighbourhoods of $x \in X$.*

Proof. Reflexivity and symmetry are clear. To see transitivity, suppose we have $s \in P(U)$ and $t \in P(V)$ and $y \in P(Z)$ such that $sG_x t$ and $tG_x y$. Then there are open neighbourhoods of x, say W and W', with $W \subseteq U \cap V$ and $W' \subseteq V \cap Z$ so that $s|_W = t|_W$ and $t|_{W'} = y|_{W'}$. But then $W \cap W'$ is open, contains x and is a subset of $U \cap Z$. It then follows by functoriality of the restriction map that

$$s|_{W\cap W'} = (s|_W)|_{W\cap W'} = (t|_W)|_{W\cap W'} = (t|_{W'})|_{W\cap W'} = (y|_{W'})|_{W\cap W'} = y|_{W\cap W'},$$

whence $sG_x y$. □

We can thus make the following definitions.

Definition 2.18. (Germ) For a presheaf P over X, let $x \in X$ and U an open neighbourhood of x. For a section $s \in P(U)$, the *germ* of s at x is the equivalence class of s under the relation G_x.

Following the notation in Mac Lane and Moerdijk [6, p. 83], we denote the germ of s at x by $\mathrm{germ}_x s$.

Definition 2.19. (Stalk) Given a presheaf P on X, let the *stalk* of P at x, denoted by P_x, be the set of all germs at x. That is,

$$P_x = \{\mathrm{germ}_x s : s \in P(U), \text{ where } U \subseteq X \text{ open and } x \in U\}.$$

For a presheaf P over X, we can therefore consider the disjoint union of its stalks

$$\Lambda_P := \coprod_{x\in X} P_x = \coprod_{x\in X} \{\mathrm{germ}_x s : s \in P(U), \text{ where } U \subseteq X \text{ open and } x \in U\}.$$

Define a function $p : \Lambda_P \to X$ by $p(\mathrm{germ}_x s) = x$. That is, p returns or projects any germ at x back to x. Next, any section $s \in P(U)$ (over an open subset U) induces a function $\dot{s} : U \to \Lambda_P$ by $\dot{s}(x) = \mathrm{germ}_x s$. Note that we now have $p\dot{s} = i_U$, which makes p look like a bundle over X together with its sections $\dot{s}$. However, presently these are just maps of sets; it thus remains to topologise Λ_P in a suitable way making p and all the $\dot{s}$ continuous. This is achieved by taking the images $\dot{s}(U)$ as the base of open sets. This is to say that $W \subseteq \Lambda_P$ is open if and only if $W = \bigcup_{t\in\Sigma} \dot{t}(V)$, where $\Sigma \subseteq P(V)$ for some open $V \subseteq X$.

To see that $\dot{s}$ is continuous for each $s \in P(U)$, consider the preimage of an open $W \subseteq \Lambda_P$.

$$\dot{s}^{-1}(W) = \dot{s}^{-1}\left(\bigcup_{t\in\Sigma} \dot{t}(V)\right)$$

whence

$$\begin{aligned} x \in \dot{s}^{-1}(W) &\iff \dot{s}(x) = \mathrm{germ}_x s \in \dot{t}(V) \quad \text{for some } V \in \mathcal{O}(X) \text{ and } t \in P(V) \\ &\iff \mathrm{germ}_x s = \mathrm{germ}_y t \quad \text{for some } y \in V \\ &\iff \mathrm{germ}_x s = \mathrm{germ}_x t \\ &\iff x \in \{z \in U \cap V : \mathrm{germ}_z s = \mathrm{germ}_z t\}, \end{aligned}$$

which is open by the definition of a germ.

Similarly, p is continuous, as for any open $U \subseteq X$,

$$p^{-1}(U) = \coprod_{x \in U} P_x = \bigcup_{s \in P(U)} \dot{s}(U),$$

which is open.

Hence for any presheaf P over X, we can assign a bundle $p : \Lambda_P \to X$ as described above. Furthermore, this process is functorial.

Lemma 2.20. *The assignment of the bundle p to a presheaf P is a functor from* $\mathbf{PSh}(X)$ *to* $\mathbf{Bund}(X)$*, denoted by*

$$\Lambda : \mathbf{PSh}(X) \to \mathbf{Bund}(X). \tag{2.21}$$

Proof. Any natural transformation of presheaves $\alpha : P \to Q$ induces maps $\alpha_x : P_x \to Q_x$ by $\alpha_x(\mathrm{germ}_x s) = \mathrm{germ}_x \alpha_U(s)$ for each $s \in P(U)$ and U an open neighbourhood of x.

We need to check that h_x is well-defined, that is, independent of the choice of the open neighbourhood U; this follows from naturality of α. In detail, if $V \subseteq U$ are open, naturality of α amounts to commutativity of

$$\begin{tikzcd} P(U) \arrow[r, "\alpha_U"] \arrow[d] & Q(U) \arrow[d] \\ P(V) \arrow[r, "\alpha_V"] & Q(V) \end{tikzcd}$$

where vertical maps are the restrictions. Now if $\mathrm{germ}_x s = \mathrm{germ}_x t$ for $s \in P(U)$ and $t \in P(V)$, then $s|_W = t|_W$ for some $W \subseteq U \cap V$, consequently, $\alpha_W(s|_W) = \alpha_W(t|_W)$. Naturality then implies that $\alpha_U(s)|_W = \alpha_V(t)|_W$ and so $\mathrm{germ}_x \alpha_U(s) = \mathrm{germ}_x \alpha_V(t)$, whence $\alpha_x(\mathrm{germ}_x s) = \alpha_x(\mathrm{germ}_x t)$.

The disjoint union of maps $\alpha_x : P_x \to Q_x$ therefore gives a map $\Lambda_\alpha : \Lambda_P \to \Lambda_Q$. Continuity of this map follows from the fact that the preimage of $\dot{t}(x) = \mathrm{germ}_x t \in Q_x$ under the map α_x is the set

$$\{\dot{s}(x) = \mathrm{germ}_x s : \alpha_U(s) = t, \text{ where } s \in P(U)\}.$$

Taking the appropriate unions yields that open sets are preserved under the preimage of Λ_α. This thus makes Λ_α into a morphism from $p : \Lambda_P \to X$ to $q : \Lambda_Q \to X$ in $\mathbf{Bund}(X)$. $\square$

Note that if we compose Λ as in (2.21) with Γ as in (2.15), we get

$$\Gamma\Lambda : \mathbf{PSh}(X) \to \mathbf{Sh}(X).$$

This is a functor turning every presheaf into a sheaf, it is known as the *sheafification functor* or the *associated sheaf functor*.

Example 2.22. Sheafification of the constant presheaf with value A (Example 2.2) is the sheaf assigning to each open set U the continuous maps from U to A, where A is viewed as a topological space with the discrete topology. To see this, notice that the germ of $a \in A$ at $x \in X$ is just the singleton $\{a\}$. Thus the space Λ_P is A with every singleton being open, that is, with the discrete topology.

2.3 Sheaves 'are' étale bundles

In the previous section we noted that $\Gamma\Lambda$ sheafifies each presheaf. Here we make a similar observation for the reverse composition $\Lambda\Gamma$. This will require the notion of an étale bundle, and will eventually lead to characterisation of sheaves as these special kind of bundles.

Since $\mathbf{Sh}(X)$ is a subcategory of $\mathbf{PSh}(X)$, the functor $\Gamma : \mathbf{Bund}(X) \to \mathbf{Sh}(X)$ as constructed in the previous section can be regarded as a functor from $\mathbf{Bund}(X)$ to $\mathbf{PSh}(X)$. We thus arrive at an important result.

Theorem 2.23. *Let X be a topological space. The functors as constructed in Section 2.2*

$$\Lambda : \mathbf{PSh}(X) \rightleftarrows \mathbf{Bund}(X) : \Gamma$$

are adjoint (left adjoint on the left).

Proof. For a presheaf P and a bundle p, we need to exhibit natural transformations

$$\eta_P : P \to \Gamma\Lambda P$$

$$\epsilon_p : \Lambda\Gamma_p \to p,$$

satisfying the usual triangle identities. This will establish the claimed adjunction; η and ϵ being its unit and counit.

We begin with η. For a presheaf P, the sheafification $\Gamma\Lambda P$ is the sheaf of sections of the bundle $p : \Lambda_P \to X$. Hence given an open set $U \subseteq X$, we need to define a map

$$\eta_{P(U)} : P(U) \to \Gamma\Lambda P(U)$$

taking each section of (the presheaf) P over U to a section of (the bundle) p over U. We only know one way to do this, hence define $\eta_{P(U)}(s) = \dot{s}$.

This assignment of $\dot{s}$ to each $s \in P(U)$ is natural in U; this amounts to the statement that $(\dot{s|_V}) = \dot{s}|_V$ for each open $V \subseteq U$, which is in turn equivalent to $\mathrm{germ}_x s|_V = \mathrm{germ}_x s$ for all elements $x \in V$ and for each open subset V. The last statement is manifestly true.

Assigning such map for each open $U \subseteq X$ yields the desired natural transformation $\eta_P : P \to \Gamma\Lambda P$. Naturality of η is the statement that $(\dot{\alpha_U(s)}) = \Gamma\Lambda\alpha_U(\dot{s})$

for any morphism of presheaves $\alpha : P \to Q$, any open set $U \in \mathcal{O}(X)$ and any section $s \in P(U)$ and it follows by unpacking the definitions.

For a bundle $p : Y \to X$, the corresponding bundle $\Lambda\Gamma_p : \Lambda_{\Gamma_p} \to X$ is the map $\mathrm{germ}_x s \mapsto x$ for a section s of p over some open U and $x \in U$. We thus need to define a continuous map ϵ_p from Λ_{Γ_p} to Y such that the diagram

$$\begin{tikzcd}
\Lambda_{\Gamma_p} \arrow[rr, "\epsilon_p"] \arrow[dr, "\Lambda\Gamma_p"'] & & Y \arrow[dl, "p"] \\
& X &
\end{tikzcd}$$

commutes. Again, we only know one way to get an element of Y from $\mathrm{germ}_x s \in \Lambda_{\Gamma_p}$, hence define $\epsilon_p(\mathrm{germ}_x s) = s(x)$. This is well-defined, as we already noted when defining the notion of a germ that sections of p (which are continuous maps) s and t having the same germ at x implies that $s(x) = t(x)$.

Since s is a section of p, the diagram above evidently commutes. It thus remains to check that ϵ_p is continuous. This follows by noting that for any open set $V \subseteq Y$, the preimage $s^{-1}(V)$ is open for any section s of p, hence the set $\dot{s}(s^{-1}(V)) \in \Lambda_{\Gamma_p}$ is open. The preimage $\epsilon_p^{-1}(V)$ is now the union of these sets over all sections of p, consequently, the preimage is open.

Finally, we easily verify that the triangle identities

$$\begin{tikzcd}
\Lambda P \arrow[r, "\Lambda\eta_P"] \arrow[dr, "\mathrm{id}_{\Lambda P}"'] & \Lambda\Gamma\Lambda P \arrow[d, "\epsilon_{\Lambda P}"] \\
& \Lambda P
\end{tikzcd}
\qquad
\begin{tikzcd}
\Gamma_p \arrow[r, "\eta_{\Gamma_p}"] \arrow[dr, "\mathrm{id}_{\Gamma_p}"'] & \Gamma\Lambda\Gamma_p \arrow[d, "\Gamma\epsilon_p"] \\
& \Gamma_p
\end{tikzcd}
\tag{2.24}$$

hold. The left diagram commutes, as is shown by taking $(\mathrm{germ}_x s) \in \Lambda_P$ and observing that the effect of the maps in the diagram is $\mathrm{germ}_x s \mapsto \mathrm{germ}_x \dot{s} \mapsto \dot{s}(x) = \mathrm{germ}_x s$ for each $x \in X$, and section s over some open neighbourhood of x. Similarly, the right diagram commutes as $s \mapsto \dot{s} \mapsto s$ for any section s of p. Hence η and ϵ are indeed the unit and the counit of the adjunction. □

We have already noticed that the functor $\Gamma\Lambda$ sheafifies each presheaf, thus sending $\mathbf{PSh}(X)$ to its full subcategory $\mathbf{Sh}(X)$. We thus ask whether $\Lambda\Gamma$ sends bundles to some special kind of bundles. This turns indeed out to be the case.

Definition 2.25. (Étale bundle) A bundle $p : Y \to X$ is *étale* if p is a *local homeomorphism*, that is, every $y \in Y$ has an open neighbourhood V such that pV is open in X and the restriction

$$p|_V : V \to pV$$

is a homeomorphism.

Next proposition gives us some intuition about what it means for a bundle to be étale. It will also be useful later, as the last part of the proof of Theorem 2.27

relies crucially on the proposition. This is part of Proposition 1 in [6, II.6 p. 88]. There, however, the proof is omitted; we thus give it here.

Proposition 2.26. *Let $p : Y \to X$ be an étale bundle. Then p and all sections (both local and global) of p are open maps. For each $y \in Y$, there is an open subset U of X and at least one section $s : U \to Y$ of p such that $y \in s(U)$. If $s : U \to Y$ and $t : V \to Y$ are sections of p over open subsets U and V, then the set*

$$W = \{x \in U \cap V : s(x) = t(x)\} \subseteq X$$

is open.

Proof. Let $V \subseteq Y$, and suppose V is open. For each $y \in V$, let V_y be an open neighbourhood such that $p(V_y)$ is open and $p|_{V_y} : V_y \to p(V_y)$ is a homeomorphism. Since V is open, $V_y \cap V$ is open and, since $p|_{V_y}$ is a homeomorphism, so is its image

$$p|_{V_y}(V_y \cap V) \subseteq p(V).$$

Hence each point $p(y) \in p(V)$ has an open neighbourhood contained in $p(V)$, that is, $p(V)$ is open. Thus p is indeed an open map.

To see that sections are open maps, let U be an open subset of an open $E \subseteq X$ and $s : E \to Y$ a section. For each $y \in sU$, let V_y and $p|_{V_y}$ be as before. By continuity of s, the set $s^{-1}(V_y) \cap U$ is open. Moreover, since ps is the inclusion, we have

$$s^{-1}(V_y) \cap U = ps(s^{-1}(V_y) \cap U) \subseteq p(ss^{-1}(V_y) \cap sU) \subseteq p(V_y \cap sU),$$

where the first containment follows by properties of intersections, and the second one by properties of preimage. Since $p(V_y \cap sU) = p|_{V_y}(V_y \cap sU)$ and $p|_{V_y}$ is a homeomorphism, we get

$$p|_{V_y}^{-1}(s^{-1}(V_y) \cap U) \subseteq V_y \cap sU \subseteq sU.$$

It is straightforward to verify that $y \in p|_{V_y}^{-1}(s^{-1}(V_y) \cap U)$ for each $y \in sU$. Hence we have found an open neighbourhood of each $y \in sU$, whence sU is open.

For any $y \in Y$, let V_y be an open neighbourhood such that $p(V_y)$ is open and $p|_{V_y} : V_y \to p(V_y)$ is a homeomorphism. The section that is always guaranteed to exist is the inverse of the homeomorphism, $s = p|_{V_y}^{-1} : p(V_y) \to V_y$.

Let W be the subset of X on which the sections s and t agree. Since $W = ps(W) = pt(W)$ and p is an open map, it is sufficient to show that $s(W) = t(W)$ is open in Y. Again, for each $y \in s(W)$, let V_y be an open neighbourhood such that $p(V_y)$ is open. Since both s and t are open maps, $sp(V_y) \cap tp(V_y)$ is open in Y, moreover, it is contained in $s(W)$. To see this, note that we can rewrite this set as follows.

$$\begin{aligned} z \in sp(V_y) \cap tp(V_y) &\iff \exists x, x' \in p(V_y) \quad \text{such that} \quad s(x) = t(x') = z \\ &\iff \exists x \in p(V_y) \quad \text{such that} \quad s(x) = t(x) = z \\ &\qquad (\text{as } s(x) = t(x') \quad \text{implies} \quad ps(x) = x = x' = pt(x')) \\ &\iff \exists x \in p(V_y) \cap W \quad \text{such that} \quad s(x) = z \\ &\iff z \in s\,(p(V_y) \cap W) \subseteq s(W). \end{aligned}$$

Since $y \in s\,(p(V_y) \cap W)$, we have found an open neighbourhood contained in $s(W)$ for each $y \in s(W)$, showing that $s(W)$ is open. □

We denote the full subcategory of $\mathbf{Bund}(X)$ consisting of étale bundles by $\mathbf{Et}(X)$. In fact, for any presheaf P over X, the bundle

$$\Lambda P = p : \Lambda_P \to X$$

is étale. To see this, let $\mathrm{germ}_x s \in \Lambda_P$. Then there is an open $U \in X$ such that $s \in P(U)$ and $x \in U$. Hence $\dot{s}(U)$ is an open neighbourhood of $\mathrm{germ}_x s$ such that $p\dot{s}(U) = U$, and $\dot{s}$ is the left inverse for p restricted to $\dot{s}U$. Since it is also the right inverse by construction, the restriction of p is a homeomorphism, and so p is indeed étale. Hence the functor

$$\Lambda\Gamma : \mathbf{Bund}(X) \to \mathbf{Et}(X)$$

can be seen as 'étalification' sending each bundle to an étale bundle. The similarity of this to the sheafification functor points to the main result of the chapter.

Theorem 2.27. *Let X be a topological space. The adjunction in Theorem 2.23 restricts to an equivalence of categories*

$$\mathbf{Sh}(X) \rightleftarrows \mathbf{Et}(X).$$

Hence sheaves over X are (in the equivalence of categories sense) local homeomorphisms over X. This further reinforces our intuition that sheaves are capturing the notion of locality. In order to prove Theorem 2.27, we will need the following general result for an adjunction.

Lemma 2.28. *Let $\mathcal{P}$ and $\mathcal{B}$ be categories and*

$$\Lambda : \mathcal{P} \rightleftarrows \mathcal{B} : \Gamma$$

an adjunction (left adjoint on the left) with unit η and counit ϵ. Let $\mathcal{P}_0$ be the full subcategory of $\mathcal{P}$ containing those objects $P \in \mathcal{P}$ for which η_P is an isomorphism, and dually, let $\mathcal{B}_0$ contain those $B \in \mathcal{B}$ for which ϵ_B is an isomorphism. Then the adjunction $(\Lambda, \Gamma, \eta, \epsilon)$ restricts to an equivalence of $\mathcal{P}_0$ and $\mathcal{B}_0$.

Proof. The only nontrivial part of the proof is to show that for any object P in $\mathcal{P}_0$ we have $\Lambda P \in \mathcal{B}_0$, and similarly that $B \in \mathcal{B}_0$ implies $\Gamma B \in \mathcal{P}_0$. Hence let $P \in \mathcal{P}_0$ so that $\eta_P : P \to \Gamma\Lambda P$ is an isomorphism. From functoriality of Λ, it follows that $\Lambda\eta_P$ is an isomorphism. Then, by the left triangle identity in (2.24), $\epsilon_{\Lambda_P} \circ \Lambda\eta_P = \mathrm{id}_{\Lambda_P}$ implying $\epsilon_{\Lambda_P} = \Lambda\eta_P^{-1}$, which being inverse of an isomorphism is itself an isomorphism. Hence ΛP is indeed an object of $\mathcal{B}_0$. The result for $B \in \mathcal{B}_0$ follows by duality. Since natural transformations restrict componentwise, this completes the proof. □

We now have all the ingredients to prove the equivalence of sheaves and étale bundles asserted in Theorem 2.27.

Proof (of Theorem 2.27). We wish to use Lemma 2.28. It thus suffices to show that a presheaf P is a sheaf if and only if $\eta_P : P \to \Gamma\Lambda P$ is an isomorphism, and similarly, a bundle $p : Y \to X$ is étale if and only if $\epsilon_p : \Lambda\Gamma_p \to p$ is an isomorphism.

First, let P be a presheaf such that η_P is an isomorphism. Let $(U_i)_{i\in I}$ be a family of open sets and write $U = \bigcup_i U_i$. Now $\Gamma\Lambda P$ is a sheaf, hence, by Proposition 2.5, the top row in the diagram

$$\begin{tikzcd}
\Gamma\Lambda P(U) \arrow[r, dashed, "e"] & \prod_i \Gamma\Lambda P(U_i) \arrow[r, shift left, "r"] \arrow[r, shift right, "q"'] & \prod_{i,j} \Gamma\Lambda P(U_i \cap U_j) \\
P(U) \arrow[u, "\eta_{P(U)}"] \arrow[r, dashed, "e'"] & \prod_i P(U_i) \arrow[u, "\left(\eta_{P(U_i)}\right)_{i\in I}"] \arrow[r, shift left, "r'"] \arrow[r, shift right, "q'"'] & \prod_{i,j} P(U_i \cap U_j) \arrow[u, "\left(\eta_{P(U_i\cap U_j)}\right)_{i,j\in I}"]
\end{tikzcd}$$

is an equaliser. Here the horizontal maps are the relevant products of restrictions as defined in Proposition 2.5, and the vertical maps are products of components of η_P. Since η_P is an isomorphism, each component in each of the vertical maps is invertible, and so the vertical maps are invertible. For clarity, we will write η_i for $\left(\eta_{P(U_i)}\right)_{i\in I}$ and η_{ij} for $\left(\eta_{P(U_i\cap U_j)}\right)_{i,j\in I}$. By naturality of η_P, the square on the left commutes, and likewise, the squares on the right commute via upper and lower paths, that is, $r\eta_i = \eta_{ij}r'$ and $q\eta_i = \eta_{ij}q'$. In order to show that P is a sheaf, we need to show that the bottom row is an equaliser. Hence suppose that $f : Z \to \prod_i P(U_i)$ is a map such that $r'f = q'f$. It then follows that $\eta_{ij}r'f = \eta_{ij}q'f$, whence $r\eta_i f = q\eta_i f$. Then, since the top row is an equaliser, there exists a unique map $g : Z \to \Gamma\Lambda P(U)$ such that $eg = \eta_i f$. Hence define $\hat{f} \coloneqq \eta_{P(U)}^{-1} \circ g$. We then have

$$e'\hat{f} = e' \circ \eta_{P(U)}^{-1} \circ g = \eta_i^{-1} \circ e \circ g = \eta_i^{-1} \circ \eta_i \circ f = f.$$

Moreover, $\hat{f}$ is unique with this property; if we suppose that $e'f' = f$ for some f', then $e\eta_{P(U)}f' = \eta_i e'f' = \eta_i f$, implying $\eta_{P(U)}f' = g$ by uniqueness of g, and so $f' = \eta_{P(U)}^{-1} \circ g = \hat{f}$. Thus the bottom row is indeed an equaliser.

Conversely, suppose that P is a sheaf. We wish to show that $\eta_{P(U)}$ is an isomorphism for each open $U \subseteq X$, which is to say that it is a bijection. For injectivity, suppose that

$$\eta_{P(U)}(s) = \dot{s} = \dot{t} = \eta_{P(U)}(t)$$

for some $s, t \in P(U)$. That is, $\mathrm{germ}_x s = \mathrm{germ}_x t$ for all $x \in U$. By the definition of a germ, it follows that each $x \in U$ has an open neighbourhood U_x such that $s|_{U_x} = t|_{U_x}$. We thus have $\bigcup_x U_x = U$ and $s|_{U_x}$ agree on the intersections in the sense of condition (2.4). Hence, since P is a sheaf, it follows by Definition 2.3 that there is a unique $\hat{s} \in P(U)$ such that $\hat{s}|_{U_x} = s|_{U_x} = t|_{U_x}$ for all $x \in U$. Since both s and t satisfy this, it follows by uniqueness that $s = \hat{s} = t$.

Similarly, using the definition of a sheaf, one can show that η_P is surjective, that is, for any section $h : U \to \Lambda_P$ of ΛP (where $U \subseteq X$ open) there is an $s \in P(U)$ such that $\dot{s} = h$. For a proof, see [6, II.5 p. 86]. This completes the proof that P is a sheaf if and only if η_P is an isomorphism.

Next, let $p : Y \to X$ be a bundle such that ϵ_p is an isomorphism. That is, $\epsilon_p : \Lambda_{\Gamma_P} \to Y$ is a homeomorphism, and so is its inverse. In particular, ϵ_p^{-1} is a local homeomorphism. The fact that p is étale follows immediately, as we have $\Lambda\Gamma_p = p\epsilon_p$, implying $p = \Lambda\Gamma_p \circ \epsilon_p^{-1}$. Since $\Lambda\Gamma_p$ and ϵ_p^{-1} are both local homeomorphisms, so is p.

Conversely, suppose that $p : Y \to X$ is étale. By Proposition 2.26, for each point $y \in Y$ there is a section $s : U \to Y$ such that $y \in sU$. We can thus define $\theta_p : Y \to \Lambda_{\Gamma_p}$ by $\theta_p(y) = \mathrm{germ}_{p(y)} s$. This map is well-defined, that is, independent of the choice of section s. This follows again from Proposition 2.26; if $t : V \to Y$ is another section such that $y \in tV$, then $s(x) = t(x) = y$ for $x = p(y) \in U \cap V$, and the set on which s and t agree is open. Hence, s and t have the same germ at $x = p(y)$, that is, $\mathrm{germ}_{p(y)} s = \mathrm{germ}_{p(y)} t$.

The map θ_p is evidently a two-sided inverse for ϵ_p. It thus remains to show that θ_p is continuous, or equivalently, that ϵ_p is an open map. This follows quickly by noting that if $\Sigma = \bigcup_{s \in S} \dot{s}U$ is an open set in Λ_{Γ_p} for some collection S of sections of p, then

$$\epsilon_p(\Sigma) = \bigcup_s \epsilon_p \dot{s}U = \bigcup_s sU,$$

which is open, since each s is an open map, once more, by Proposition 2.26. Thus ϵ_p is indeed an isomorphism. □

Chapter 3

Connections with Logic

We mentioned in the beginning of Section 1.2 that we view a topos as a generalisation of the category **Set**. One important property of a set is that its subsets have some sort of 'logical operations'. We say that an element is in the union of two sets if it is in one *or* in the other, an element is in the intersection if it is in one set *and* in the other one, an element is in the complement of a set if it is *not* in the set. The notion of a Boolean algebra defined in Section 3.1 makes this precise.

It turns out that the fact that we can use classical logic in defining union, intersection and complement is very special to the category of sets (namely, it is a Boolean topos as defined in Section 3.3). In a general topos, the corresponding 'algebra of subobjects' is given by a Heyting algebra, whereof Boolean algebra is a special case. The first two sections in this chapter will define the Heyting algebra of subobjects in an arbitrary topos. The final section will connect this to our discussion of sheaves; we will see that the category of sheaves is a topos, and that the algebra of subobjects is given by the poset of open sets of the underlying topological space.

3.1 Boolean and Heyting algebras

Boolean algebra provides a model for a classical logic, in the sense that the law of double negation, or equivalently the law of excluded middle is valid in any such algebra. The simplest example is given by the classical two-valued logic, each statement is either true or false, and the binary operations 'and', 'or' and 'implication', as well as unary operation 'negation' are defined by their usual truth tables. Heyting algebra generalises this to model an intuitionistic logic, in which the law of double negation is not valid.

Both Boolean and Heyting algebras can be defined purely equationally; we have, however, chosen an approach which uses a poset structure to define logical operations, as it requires fewer definitions and exposes connections between concepts more explicitly.

Definition 3.1. Let $(P, \leq)$ be a poset, and let $x, y \in P$. An *infimum* (or *meet*) of x and y is an element of P, denoted by $x \wedge y$, such that

i. $x \wedge y \leq x$ and $x \wedge y \leq y$,

ii. $z \leq x$ and $z \leq y$ implies $z \leq x \wedge y$ for all $z \in L$.

Dually, a *supremum* (or *join*) of x and y is an element of P, denoted by $x \vee y$, such that

i. $x \leq x \vee y$ and $y \leq x \vee y$,

ii. $x \leq z$ and $y \leq z$ implies $x \vee y \leq z$ for all $z \in L$.

Infima and suprema are not guaranteed to exist for an arbitrary poset. However, if they do exist, they are unique. The easiest way to see this is to notice that infimum and supremum are the product and the coproduct in the poset viewed as a category (defined after Definition 1.7). Indeed, the definition of a lattice infimum states that whenever there are maps from z to both x and y, they factor (uniquely, as there is at most one map) through $x \wedge y$. Hence infimum is the categorical product of x and y. Dually, supremum is the coproduct. Since products and coproducts are unique up to an isomorphism, and the only isomorphisms is a poset are the identities, it follows that supremum and infimum are unique elements of P if they exist.

A poset which has all suprema and all infima is so important it has a special name.

Definition 3.2. (Lattice) A *lattice* is a poset $(L, \leq)$ such that any two elements $x, y \in L$ have the infimum denoted $x \wedge y$ and the supremum $x \vee y$.

We will simply write L for the lattice with underlying poset $(L, \leq)$. A concise way to define a lattice is to require the poset category to have all binary products and coproducts.

Two immediate consequences of the definition of infimum (or the fact that it is the product) are

$$\begin{aligned} x \wedge y &= y \wedge x, \quad \text{and} \\ (x \wedge y) \wedge z &= x \wedge (y \wedge z) \end{aligned}$$

for all $x, y, z \in L$, that is, $\wedge$ is symmetric and associative. We can thus interchange the order of the elements and drop the brackets without ambiguity. The same is, of course, also true for supremum.

A lattice L is said to be *bounded* if there are elements $0, 1 \in L$ such that $0 \leq x \leq 1$ for all $x \in L$. Hence 0 and 1 are initial and terminal objects of L, respectively. The initial object 0 is called the *bottom object* of L and 1 correspondingly the *top object*. It follows from the definitions in a straightforward manner that the following identities hold in any bounded lattice.

$$\begin{aligned} x \wedge x &= x = x \vee x, \\ 1 \wedge x &= x = 0 \vee x, \\ x \wedge (y \vee x) &= x = (x \wedge y) \vee x. \end{aligned} \tag{3.3}$$

For example, the first identity on the last line follows by noting that $x \leq x$ and $x \leq y \vee x$ and so $x \leq x \wedge (y \vee x) \leq x$. Moreover, given the equations (3.3), we can recover the partial order on L by noticing that $x \leq y$ if and only if $x \wedge y = x$ (or equivalently, $x \vee y = y$). Hence we could equivalently define a bounded lattice as a set L having two special elements $0, 1 \in L$ together with associative and symmetric binary operations $\wedge, \vee : L \times L \to L$ such that the equations (3.3) hold; this induces a partial order on L by the observed equivalence. The interpretation of $\wedge$ and $\vee$ as infimum and supremum then follows by proving the properties of Definition 3.2. For example, $x \wedge y \leq x$ follows from

$$x \wedge y \wedge x = x \wedge x \wedge y = x \wedge y,$$

and the second property of supremum follows since $x \vee z = z$ and $y \vee z = z$ implies

$$(x \vee y) \vee z = x \vee (y \vee z) = x \vee z = z.$$

A lattice is said to be *distributive* if either of the following equivalent identities holds for all $x, y, z \in L$,

$$x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z),$$

$$x \vee (y \wedge z) = (x \vee y) \wedge (x \vee z).$$

In order to see the equivalence of these, one simply expands the right-hand sides using the other equation, the expression then simplifies using the identities 3.3. For instance, let us assume the second equation, the right-hand side of the first one then expands to

$$\begin{aligned}
(x \wedge y) \vee (x \wedge z) &= ((x \wedge y) \vee x) \wedge ((x \wedge y) \vee z) \\
&= x \wedge ((x \vee z) \wedge (y \vee z)) \\
&= (x \wedge (x \vee z)) \wedge (y \vee z) \\
&= x \wedge (y \vee z),
\end{aligned}$$

as required.

Definition 3.4. (Heyting algebra) A *Heyting algebra* is a bounded lattice H such that for any pair of elements $x, y \in H$ there is an exponential y^x, written $x \Rightarrow y$.

Using the terminology introduced in Chapter 1, a Heyting algebra H is then a poset category which is Cartesian closed and has all finite coproducts. The initial object 0 can then be seen as the empty coproduct, analogously to the terminal object being the empty product. Unpacking the definition of exponential (Definition 1.2), we get that

$$z \wedge x \leq y \quad \text{if and only if} \quad z \leq x \Rightarrow y \tag{3.5}$$

for all $z \in H$. The 'if' direction follows since $z \leq x \Rightarrow y$ implies

$$z \wedge x \leq (x \Rightarrow y) \wedge x \leq y,$$

where the last 'inequality' is just the evaluation morphism of the exponential as in Definition 1.2. The exponential $\Rightarrow$ is called *implication* and in an algebra describing classical logic it will turn out to have its usual meaning.

Proposition 3.6. *Any Heyting algebra is a distributive lattice.*

Proof. We will show that

$$x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z)$$

for all $x, y, z \in H$. Let $u \in H$, we then have the following sequence of equivalent statements, following the proof in Bell [1, p. 4].

$$\begin{aligned} x \wedge (y \vee z) \leq u \quad & \text{iff} \quad y \vee z \leq x \Rightarrow u \\ & \text{iff} \quad y \leq x \Rightarrow u \text{ and } z \leq x \Rightarrow u \\ & \text{iff} \quad y \wedge x \leq u \text{ and } z \wedge x \leq u \\ & \text{iff} \quad (y \wedge x) \vee (z \wedge x) \leq u. \end{aligned}$$

Then first taking $u = x \wedge (y \vee z)$ and then $u = (y \wedge x) \vee (z \wedge x)$ gives the result. □

Since a Heyting algebra is supposed to capture some kind of propositional logic, we would like it to have a negation operation. For an element x in a Heyting algebra H, we thus define its *negation* or *pseudocomplement* as

$$\neg x := (x \Rightarrow 0).$$

Hence, the 'negation of x' means 'x implies falsity' or 'x implies absurdity' [6, p. 53]. To get a better idea of the definition, we can reformulate this using (3.5). For $z \in H$ we have

$$z \leq \neg x \quad \text{if and only if} \quad z \wedge x = 0. \tag{3.7}$$

If we think of the elements in a Heyting algebra entailing the ones above them, this tells us that z entails the pseudocomplement of x if and only if z and x are disjoint in the sense that no element apart from zero entails both of them. From this characterisation of pseudocomplements it immediately follows that $\neg 0 = 1$, $\neg 1 = 0$ and $\neg x \wedge x = 0$. Note, however, that we do not necessarily have $\neg x \vee x = 1$. The last identity is known as the law of excluded middle, and if it holds, we say that $\neg x$ is the *complement* of x.

Importantly, in general it is not the case that $\neg\neg x = x$ (see Example 3.8 below), which is to say that the law of double negation does not hold. We do have, however, that $x \leq \neg\neg x$ since $x \wedge \neg x = 0$. Interestingly, it turns out that the law of double negation is equivalent to the law of excluded middle, as proved in Proposition 3.11.

The canonical example of a Heyting algebra is the so called 'algebra of opens'.

Example 3.8. (Algebra of opens) Let X be a topological space. The open sets $\mathcal{O}(X)$ of X ordered by inclusion form a Heyting algebra, with infimum being the intersection and supremum the union of two sets. The bottom object is $0 = \varnothing$ and the top one $1 = X$. By (3.5) it follows that the implication $U \Rightarrow V$ is the interior of $U^c \cup V$. Then, taking $V = \varnothing$, it follows that the pseudocomplement is given by the interior of the complement. In general, the law of double negation does not hold in $\mathcal{O}(X)$. Take, for example $X = (-1, 1)$ as a subset of $\mathbb{R}$, and take $U = X \setminus \{0\}$. Then $\neg\neg U = X \neq U$.

The same example demonstrates that one of the De Morgan's laws (see Proposition 3.9) does not hold in a general Heyting algebra. Take $V = (-1, 0)$ and $W = (0, 1)$. Then $\neg(V \cap W) = X$, but $\neg V \cup \neg W = U \neq X$.

Proposition 3.9. *(De Morgan's laws) In any Heyting algebra H, the following holds for all $x, y \in H$,*

$$\neg(x \vee y) = \neg x \wedge \neg y,$$

$$\neg x \vee \neg y \leq \neg(x \wedge y).$$

Proof. For $z \in H$, we have the following sequence of equivalent statements,

$$\begin{aligned} z \leq \neg(x \vee y) &\iff z \wedge (x \vee y) = 0 \\ &\iff (z \wedge x) \vee (z \wedge y) = 0 \\ &\iff z \wedge x = 0 \text{ and } z \wedge y = 0 \\ &\iff z \leq \neg x \text{ and } z \leq \neg y \\ &\iff z \leq \neg x \wedge \neg y, \end{aligned}$$

from which the first identity follows. For the second one, it is equivalent to show that $(\neg x \vee \neg y) \wedge (x \wedge y) = 0$. Thus compute

$$(\neg x \vee \neg y) \wedge (x \wedge y) = (\neg x \wedge x \wedge y) \vee (\neg y \wedge y \wedge x) = 0 \vee 0 = 0,$$

as required. □

Remark 3.10. The usual De Morgan's law $\neg(x \wedge y) = \neg x \vee \neg y$ does not hold in a general Heyting algebra. Only the containment as in Proposition 3.9 is true in general. The failure of the opposite containment is demonstrated by Example 3.8.

For an algebra to model classical logic, we would certainly want both De Morgan's laws to hold and to have complements (not just pseudocomplements) for each element. This motivates the next proposition and the definition following it.

Proposition 3.11. *In any Heyting algebra H, the following are equivalent.*

i. $\neg\neg x = x$ *for all* $x \in H$ *(law of double negation),*

ii. $\neg x \vee x = 1$ *for all* $x \in H$ *(law of excluded middle).*

Proof. First suppose that (i) holds. Then, by Proposition 3.9,

$$\begin{aligned} \neg(x \vee \neg x) &= \neg x \wedge \neg\neg x \\ &= \neg x \wedge x \\ &= 0. \end{aligned}$$

Hence $\neg\neg(x \vee \neg x) = \neg 0$, implying $x \vee \neg x = 1$.

Now suppose that (ii) holds. We already observed that $x \leq \neg\neg x$ in any Heyting algebra; to show that the law of double negation holds it therefore suffices to show that $\neg\neg x \leq x$. For a $z \in H$, we have the following sequence of implications,

$$\begin{aligned} z \leq \neg\neg x &\implies z \wedge \neg x = 0 \\ &\implies (z \wedge \neg x) \vee x = x \\ &\implies (z \vee x) \wedge (\neg x \vee x) = x \\ &\implies z \vee x = x \\ &\implies z \leq x, \end{aligned}$$

from which it follows that $\neg\neg x \leq x$. □

Definition 3.12. (Boolean algebra) A *Boolean algebra* is a Heyting algebra such that one of the equivalent conditions in Proposition 3.11 holds.

Hence in a Boolean algebra all pseudocomplements are complements and negation is an involution. Furthermore, in any Boolean algebra B, the second De Morgan's law holds, which is to say that $\neg(x \wedge y) = \neg x \vee \neg y$ for all $x, y \in B$. To see this, observe that for a $z \in B$ we have the following sequence of equivalent statements,

$$\begin{aligned} z \leq \neg(x \wedge y) &\iff z \wedge x \wedge y = 0 \\ &\iff z \wedge (\neg\neg x \wedge \neg\neg y) = 0 \\ &\iff z \wedge \neg(\neg x \vee \neg y) = 0 \\ &\iff z \leq \neg\neg(\neg x \vee \neg y) \\ &\iff z \leq \neg x \vee \neg y, \end{aligned}$$

from which the equality follows. Hence both De Morgan's laws, the law of double negation and the law of excluded middle all hold in a Boolean algebra. Boolean algebras therefore provide a suitable language for classical logic. We conclude our discussion of Boolean and Heyting algebras with a canonical example of a Boolean algebra.

Example 3.13. Let X be a set. Then the power set $\mathcal{P}(X)$ ordered by inclusion is a Boolean algebra under intersection and union, with $0 = \varnothing$ and $1 = X$. As in Example 3.8, by (3.5) it follows that the implication $U \Rightarrow V$ is $U^c \cup V$. It then follows by taking $V = \varnothing$ that negation is the set-theoretic complement. Since the complement of a complement is the identity, $\mathcal{P}(X)$ is indeed a Boolean algebra.

3.2 Logic in a topos

We would like to apply the discussion of the previous section to define a 'logic' an arbitrary category. We ask, how much of the poset, lattice and Heyting algebra structure still makes sense for subobjects rather than subsets? It turns out that 'posets' can always be formed, but in order to get a lattice we need more structure that provided by an arbitrary category. We will thus focus on toposes, which are meant to generalise the category of sets and functions. Surprisingly, the algebra that emerges from a topos is not in general Boolean, but rather a Heyting algebra. The topos **Set** is thus special in the sense that it gives rise to a Boolean algebra, as described below.

Let $\mathcal{C}$ be a category and C an object of $\mathcal{C}$. Let $\mathrm{Sub}(C)$ be the category whose objects are all subobjects of C and morphisms all monics between the subobjects (as maps in $\mathcal{C}/C$). We can then define a partial order on $\mathrm{Sub}(C)$ by letting $m \leq n$ if and only if there exists a monic from m to n. This is one of the motivations to define a subobject as an isomorphism class of monics, had we defined a subobject simply as a monic, this definition would yield a preorder rather than an order. Now, however, $m \leq n \leq m$ implies that m and n are isomorphic and thus belong to the same subobject. The category $\mathrm{Sub}(C)$ is often called the *poset of subobjects* of C, although it may well be a partially ordered proper class.

The poset of subobjects gives us a way to define the image of a morphism in a general category.

Definition 3.14. (Image) Let $f : D \to C$ be a morphism in a category $\mathcal{C}$. An *image* of f is a subobject $m \in \mathrm{Sub}(C)$ such that $f = me$ for some e, and whenever $f = m'e'$ for some $m' \in \mathrm{Sub}(C)$, then $m \leq m'$.

Hence an image m is the smallest subobject of the codomain C through which f factors. Note that although the definition of an image makes sense in a general category, the image is not guaranteed to exist. As one anticipates for definitions via universal properties, the image is unique up to an isomorphism if it exists. Moreover, in any topos the image always exists, and the map e for which $f = me$ is an epimorphism (see Proposition 1 in section IV.6 in Mac Lane and Moerdijk [6, p. 185]). In particular, in **Set** this definition of an image reduces to that of an image of a function.

Although the construction of $\mathrm{Sub}(C)$ makes sense in any category, it is by no means guaranteed to have a lattice structure. If, however, the category we start with is a topos, $\mathrm{Sub}(C)$ is always a lattice, whose infimum and supremum operations we will now construct. For the rest of this section, let $\mathcal{E}$ be a topos with subobject classifier $(\Omega, \mathtt{t})$.

Definition 3.15. (Conjunction) *Conjunction*

$$\frown : \Omega \times \Omega \to \Omega$$

in $\mathcal{E}$ is the characteristic morphism of the product of two `true` maps, $\mathtt{t} \times \mathtt{t} : 1 \to \Omega \times \Omega$.

It is worth writing this definition out in detail. The product morphism $\mathtt{t} \times \mathtt{t} : 1 \to \Omega \times \Omega$ is certainly monic, hence by the definition of subobject classifier (Definition 1.8), there is a unique morphism $\chi_{\mathtt{t}\times\mathtt{t}} : \Omega \times \Omega \to \Omega$ such that the diagram

$$\begin{tikzcd} 1 \arrow[r, "\mathtt{t}\times\mathtt{t}"] \arrow[d] & \Omega\times\Omega \arrow[d, "\chi_{\mathtt{t}\times\mathtt{t}}"] \\ 1 \arrow[r, "\mathtt{t}"] & \Omega \end{tikzcd}$$

is a pullback. We thus define $\frown := \chi_{\mathtt{t}\times\mathtt{t}}$. This definition is motivated by the case where $\mathcal{E} = \mathbf{Set}$ and $\Omega = 2 = \{0,1\}$. Then the product map $\mathtt{t}\times\mathtt{t}$ 'picks' the element $(1,1)$ from $\Omega\times\Omega$, and so the characteristic function sends $(1,1)$ to 1 (true) and everything else to 0 (false), which is precisely the classical truth table for conjunction.

Definition 3.16. (Disjunction) Let $f : \Omega \amalg \Omega \to \Omega\times\Omega$ be the unique map induced by $\mathtt{t}\times\mathrm{id}_\Omega : \Omega \to \Omega\times\Omega$ and $\mathrm{id}_\Omega\times\mathtt{t} : \Omega\to\Omega\times\Omega$. We then define *disjunction*

$$\smile : \Omega\times\Omega\to\Omega$$

in $\mathcal{E}$ as the the characteristic morphism of the image of f.

Again, let us write this definition in more detail. Firstly, we have the following coproduct diagram.

$$\begin{tikzcd} & \Omega\times\Omega & \\ \Omega \arrow[ur, "\mathrm{id}_\Omega\times\mathtt{t}"] \arrow[r, "\amalg_1"'] & \Omega\amalg\Omega \arrow[u, dashed, "f"'] & \Omega \arrow[ul, "\mathtt{t}\times\mathrm{id}_\Omega"'] \arrow[l, "\amalg_2"] \end{tikzcd}$$

Here $\amalg_1$ and $\amalg_2$ are the relevant coprojection maps. Next, f factors as

$$\begin{tikzcd} \Omega\amalg\Omega \arrow[rr, "f"] \arrow[dr, "e"', two heads] & & \Omega\times\Omega \\ & \operatorname{im} f \arrow[ur, "m"', hook] & \end{tikzcd}$$

where m is a monic. Hence, by the definition of the subobject classifier, there is a unique morphism $\chi_m : \Omega\times\Omega\to\Omega$ such that the diagram

$$\begin{tikzcd} \operatorname{im} f \arrow[r, "m"] \arrow[d] & \Omega\times\Omega \arrow[d, "\chi_m"] \\ 1 \arrow[r, "\mathtt{t}"] & \Omega \end{tikzcd}$$

is a pullback. The disjunction is thus defined as $\smile := \chi_m$. The motivation for this rather convoluted definition is, again, that it reduces to the usual disjunction in the topos of sets and functions. Indeed, if $\Omega = 2$, then $\mathtt{t}\times\mathrm{id}_\Omega$ is the map $1\mapsto(1,1)$ and $0\mapsto(1,0)$. Similarly, $\mathrm{id}_\Omega\times\mathtt{t}$ maps $1\mapsto(1,1)$ and $0\mapsto(0,1)$. Hence $\operatorname{im} f = \{(1,1);(1,0);(0,1)\}$, and so $\smile = \chi_m$ maps $(0,0)\in\Omega\times\Omega\setminus\operatorname{im} f$ to 0 (false) and everything else to 1 (true), as one would expect disjunction to behave.

Definition 3.17. (Conditional) Let $e : \circledleq \to \Omega \times \Omega$ be the equaliser of $\cap$ and π_1 as in the diagram

$$\begin{tikzcd} \circledleq \arrow[r, "e"] & \Omega \times \Omega \arrow[r, shift left, "\cap"] \arrow[r, shift right, "\pi_1"'] & \Omega \end{tikzcd},$$

where π_1 is the first projection map. The *conditional*

$$\Rrightarrow: \Omega \times \Omega \to \Omega$$

in $\mathcal{E}$ is the characteristic morphism of e.

Once more, it is easy to see that in **Set** this coincides with the usual notion of the conditional. The equaliser of $\cap$ and π_1 is the set

$$\circledleq = \{(0,0);(0,1);(1,1)\} \subseteq 2 \times 2$$

with e being the inclusion, thus the characteristic function of $\circledleq$ reproduces the truth table for the conditional. Note that in this case the set $\circledleq$ can be described as all those pairs whose first component is below (or equal to) the second one, explaining the notation used (which is adapted from Goldblatt [2, ch. 6.6]).

Thus far, we have generalised the familiar logical operators to the subobject of an arbitrary topos. In order to define a lattice structure on the poset of subobjects, we will next pull back the conjunction, disjunction and conditional morphisms from the subobject classifier to the poset. We follow the presentation of Goldblatt [2, ch. 7]. Fix an object C of $\mathcal{E}$, and let $m : A \to C$ and $n : B \to C$ be subobjects of C.

Intersection The intersection of n and m is the subobject

$$n \wedge m : A \wedge B \to C$$

which is the pullback of $\mathtt{t}$ along $\chi_m \cap \chi_n := \cap \circ (\chi_m \times \chi_n)$. In other words, the characteristic morphism of $m \wedge n$ is $\chi_{m\wedge n} = \chi_m \cap \chi_n$, as in the left diagram below.

Union Analogously to the intersection, the union is the subobject

$$n \vee m : A \vee B \to C$$

which is the pullback of $\mathtt{t}$ along $\chi_m \cup \chi_n := \cup \circ (\chi_m \times \chi_n)$. Hence $\chi_{m\vee n} = \chi_m \cup \chi_n$, as in the middle diagram.

Implication The implication

$$n \Rightarrow m : A \Rightarrow B \to C$$

is the pullback of $\mathtt{t}$ along $\chi_m \Rrightarrow \chi_n := \Rrightarrow \circ (\chi_m \times \chi_n)$. Hence $\chi_{m\Rightarrow n} = \chi_m \Rrightarrow \chi_n$, as in the right diagram.

$$
\begin{tikzcd}
A \wedge B \arrow[r, "m\wedge n"] \arrow[d] & C \arrow[d, "\chi_m \frown \chi_n"] \\
1 \arrow[r, "\mathtt{t}"] & \Omega
\end{tikzcd}
\qquad
\begin{tikzcd}
A \vee B \arrow[r, "m\vee n"] \arrow[d] & C \arrow[d, "\chi_m \smile \chi_n"] \\
1 \arrow[r, "\mathtt{t}"] & \Omega
\end{tikzcd}
\qquad
\begin{tikzcd}
A \Rightarrow B \arrow[r, "m\Rightarrow n"] \arrow[d] & C \arrow[d, "\chi_m \Mapsto \chi_n"] \\
1 \arrow[r, "\mathtt{t}"] & \Omega
\end{tikzcd}
$$

The following result will be very useful in all subsequent discussion.

Lemma 3.18. *Let $m : A \to C$ and $n : B \to C$ be subobjects of C in some topos, and let*

$$
\begin{tikzcd}
D \arrow[r, "m^*"] \arrow[d, "n^*"'] \arrow[dr, "\alpha"] & B \arrow[d, "n"] \\
A \arrow[r, "m"] & C
\end{tikzcd}
\tag{3.19}
$$

be their pullback, where $\alpha = nm^* = mn^*$. Then $\alpha = m \wedge n$, so that $D \cong A \wedge B$.

Proof. We claim that then the inner square in the diagram

$$
\begin{tikzcd}
X \arrow[dr, dashed, "k"] \arrow[drr, bend left, "f"] \arrow[ddr, bend right, "g"'] & & & \\
& D \arrow[r, "\alpha"] \arrow[d] & C \arrow[d, "\chi_m \times \chi_n"] & \\
& 1 \arrow[r, "\mathtt{t}\times\mathtt{t}"] & \Omega \times \Omega \arrow[r, "\pi_i"] & \Omega
\end{tikzcd}
$$

is also a pullback. Here π_i is the projection map for component $i = 1, 2$.

Let f and g be some maps as in the diagram such that the outer square commutes (although for g there is, of course, no choice). Since $\pi_1 \circ (\chi_m \times \chi_n) = \chi_m$, this implies that $\chi_m f = tg$. Now, since the square defining χ_m is a pullback (Definition 1.8), there is a unique map $u : X \to A$ such that $f = mu$. Similarly, using the second projection map, we get a unique map $u' : X \to B$ such that $f = nu'$; hence in particular we have $mu = nu'$. Since (3.19) is a pullback, there is a unique map $k : X \to D$ such that $u' = m^* k$ and $u = n^* k$, whence $\alpha k = nm^* k = nu' = f$. Uniqueness of such map follows since α is a monic. Hence the inner square in the above diagram is indeed a pullback.

It follows that in the diagram below, both smaller squares are pullbacks; the lower one is just the definition of $\frown$.

$$
\begin{tikzcd}
D \arrow[r, "\alpha"] \arrow[d] & C \arrow[d, "\chi_m \times \chi_n"] \\
1 \arrow[r, "\mathtt{t}\times\mathtt{t}"] \arrow[d] & \Omega \times \Omega \arrow[d, "\frown"] \\
1 \arrow[r, "\mathtt{t}"] & \Omega
\end{tikzcd}
$$

By the pullback lemma 1.6, the outer rectangle is also a pullback. Hence, by the definition of characteristic morphisms, $\chi_\alpha = \cap \circ (\chi_m \times \chi_n) = \chi_m \cap \chi_n$. By uniqueness of characteristic morphisms, we conclude that α belongs to the same subobject with $m \wedge n$. $\square$

Theorem 3.20. *Let $\mathcal{E}$ be a topos. Then $(Sub(C), \leq, \wedge, \vee)$ is a bounded lattice for any $C \in \mathcal{E}$. The bottom object is $0_C : 0 \to C$, the unique map from the initial object of $\mathcal{E}$ into C. The top object is id_C.*

Proof. Since by Proposition 1.5 a pullback of a monic is monic, it follows by Lemma 3.18 that $m \wedge n \leq m$ and $m \wedge n \leq n$. The second property of the infimum, that is, $k \leq m$ and $k \leq n$ imply $k \leq m \wedge n$ for all $m, n, k \in \mathrm{Sub}(C)$, follows from the fact that (3.19) is a pullback if we take $D = A \wedge B$. Universality of the square then guarantees existence of a monic from k to $m \wedge n$.

The proof that $\vee$ is the lattice supremum is rather lengthy and requires a technical result, we therefore refer the reader to Theorem 1 in section 7.2 in Goldblatt [2, p. 151] for the proof.

Boundedness of the lattice follows immediately from observing that the monics 0_C and id_C are initial and terminal in $\mathrm{Sub}(C)$, respectively. $\square$

In order to prove that $\mathrm{Sub}(C)$ is a Heyting algebra, we will need the following results, which are Lemma 1 and its corollary in section 7.5 in Goldblatt [2, p. 163].

Lemma 3.21. *Let $m, n, k \in Sub(C)$ in some topos. Then*

$$m \wedge k = n \wedge k \quad \textit{if and only if} \quad \chi_m k = \chi_n k.$$

Proof. We have the following commutative diagrams.

$$\begin{tikzcd}
A \wedge D \arrow[r, "m^*"] \arrow[d, "k^*"'] \arrow[dr, "m \wedge k"] & D \arrow[d, "k"] \\
A \arrow[r, "m"] \arrow[d] & C \arrow[d, "\chi_m"] \\
1 \arrow[r, "\mathtt{t}"] & \Omega
\end{tikzcd}
\qquad
\begin{tikzcd}
B \wedge D \arrow[r, "n'"] \arrow[d, "k'"'] \arrow[dr, "n \wedge k"] & D \arrow[d, "k"] \\
B \arrow[r, "n"] \arrow[d] & C \arrow[d, "\chi_n"] \\
1 \arrow[r, "\mathtt{t}"] & \Omega
\end{tikzcd}$$

The upper squares are pullbacks by Lemma 3.18 and the lower ones by definition of the subobject classifier. Hence by the pullback lemma 1.6, the outer rectangles are also pullbacks. By uniqueness of characteristic functions, we thus have $\chi_{m^*} = \chi_m k$ and $\chi_{n'} = \chi_n k$. Hence $\chi_m k = \chi_n k$ if and only if $m^* = n'$ as subobjects of C. Since k is a monic, this is equivalent to $km^* = kn'$, that is, $m \wedge k = n \wedge k$. $\square$

Corollary 3.22. *For $m, n, k \in Sub(C)$,*

$$k \wedge m \leq n \quad \textit{if and only if} \quad \chi_m k = \chi_{m \wedge n} k.$$

Proof. We have $k \wedge m \leq n$ if and only if $k \wedge m \wedge n = k \wedge m$, which by the above lemma is equivalent to $\chi_m k = \chi_{m \wedge n} k$. $\square$

Theorem 3.23. *(Theorem 1.(1) in [2, 7.5 p. 164]) The lattice as defined in Theorem 3.20 is a Heyting algebra, with implication given by $\Rightarrow$ as defined on page 34.*

Proof. Let $m : A \to C$ and $n : B \to C$ be subobjects of C. We will show that condition (3.5) holds, that is, for any subobject $k : D \to C$ we have

$$k \wedge m \leq n \quad \text{if and only if} \quad k \leq m \Rightarrow n.$$

Consider the diagram,

$$\begin{tikzcd}
A \Rightarrow B \arrow[r, "m \Rightarrow n"] \arrow[d, dashed] & C \arrow[d, "\chi_m \times \chi_n"] & \\
\textcircled{$\leq$} \arrow[r, "e"] \arrow[d] & \Omega \times \Omega \arrow[r, shift left, "\cap"] \arrow[r, shift right, "\pi_1"'] \arrow[d, "\Mapsto"] & \Omega \\
1 \arrow[r, "\mathtt{t}"] & \Omega &
\end{tikzcd}$$

where the middle row is the equaliser in the definition of conditional (Definition 3.17), and the bottom square is a pullback by the same definition. Now the outer rectangle is a pullback by the definition of implication $\Rightarrow$. Since the bottom square is a pullback, the dotted arrow making the top square commute exists (and is unique). It follows by the pullback lemma 1.6 that the top square is a pullback.

Now $k \leq m \Rightarrow n$ if and only if there is a monic $u : D \to A \Rightarrow B$ such that $(m \Rightarrow n) \circ u = k$. Since the top square is a pullback, this happens if and only if there is a map $f : D \to \textcircled{$\leq$}$ such that $ef = (\chi_m \times \chi_n) \circ k$. By the universal property of the equaliser (Definition 1.1), such map exists if and only if

$$\pi_1 \circ (\chi_m \times \chi_n) \circ k = \cap \circ (\chi_m \times \chi_n) \circ k.$$

Since $\pi_1 \circ (\chi_m \times \chi_n) = \chi_m$ and $\cap \circ (\chi_m \times \chi_n) = \chi_{m \wedge n}$, this is equivalent to

$$\chi_m k = \chi_{m \wedge n} k,$$

which by Corollary 3.22 is equivalent to $k \wedge m \leq n$. $\square$

Theorem 3.23 allows us to apply the discussion on Heyting algebras in the previous section to $\mathrm{Sub}(C)$ in any topos. In particular, it shows that we have implicitly defined the negation $\neg m = m \Rightarrow 0_C$, for a subobject m of C. However, $m \Rightarrow 0_C$ is defined as a pullback of $\mathtt{t}$ along $\chi_m \Mapsto \chi_{0_C}$, which is itself defined in terms of $\Mapsto$, the characteristic morphism of the equaliser of $\cap$ and a projection map. This convoluted definition of negation is somewhat unsatisfactory, as negation is perhaps the simplest logical operation, and our intuition suggests that this should not result in such a complicated notion in a topos. More importantly, this definition tells us absolutely nothing about what negation does. The next proposition addresses these two issues in a particular case.

Proposition 3.24. *Let* $(\Omega, \mathtt{t})$ *be a subobject classifier in some topos. Then* $\neg\mathtt{t} = \mathtt{f}$*, where* $\mathtt{f}$ *is the map* $\mathtt{false}$ *and* $\mathtt{t}$ *is viewed as a subobject of* Ω*.*

Proof. Since in any Heyting algebra $\mathtt{t} \wedge \neg\mathtt{t} = 0_\Omega$, by Lemma 3.18 we have the pullback diagram.

$$\begin{tikzcd}
0 \arrow[r, "\mathtt{t}^*"] \arrow[d] \arrow[dr, "\mathtt{t}\wedge\neg\mathtt{t}"] & 1 \Rightarrow 0 \arrow[d, "\neg\mathtt{t}"] \\
1 \arrow[r, "\mathtt{t}"] & \Omega
\end{tikzcd}$$

Observe that if we show that $1 \Rightarrow 0 \cong 1$, then uniqueness of the characteristic map and definition of $\mathtt{false}$ will yield the desired result, $\neg\mathtt{t} = \mathtt{f}$. Suppose therefore, that there is a map $k : 1 \to 1 \Rightarrow 0$, which is necessarily monic as its domain is a terminal object. It follows that $\neg\mathtt{t}k$ is a subobject of Ω with $\neg\mathtt{t}k \leq \neg\mathtt{t}$, which by (3.7) implies $\mathtt{t} \wedge (\neg\mathtt{t}k) = 0_\Omega$. Then, again by Lemma 3.18, the diagram

$$\begin{tikzcd}
0 \arrow[r] \arrow[d] \arrow[dr, "\mathtt{t}\wedge(\neg\mathtt{t}k)"] & 1 \arrow[d, "\neg\mathtt{t}k"] \\
1 \arrow[r, "\mathtt{t}"] & \Omega
\end{tikzcd}$$

is a pullback. But then $\neg\mathtt{t}k$ is nothing but the characteristic function of $0_1 : 0 \to 1$, which by definition is the map $\mathtt{false}$.

Hence, since $\neg\mathtt{t}$ is a monic, we have shown that if there is a map from 1 to $1 \Rightarrow 0$, then it is unique. To see that there is such a map, note that $\mathtt{t} \wedge \mathtt{f} = 0_\Omega$, which by (3.7) implies that $\mathtt{f} \leq \neg\mathtt{t}$, whence existence follows. Since 1 is terminal, there is also a unique map from $1 \Rightarrow 0$ to 1. Uniqueness of these two maps implies they must be mutually inverse, and so indeed $1 \Rightarrow 0 \cong 1$. □

It is important to note that $\mathrm{Sub}(C)$ is not itself an object in the topos $\mathcal{E}$. We are thus taking an 'outside' perspective when considering this lattice structure. For this reason, $\mathrm{Sub}(C)$ is called an *external* Heyting algebra (of subobjects of C). However, Golblatt invites us to think of a topos (and of a category in general) as "a universe for a particular kind of mathematical discourse" [2, p. 1]. Hence, although all standard set-theoretic constructions exist in any topos, different toposes give rise to different results about these structures. It is therefore of interest to know what kind of statements are true *internally*, within that particular discourse.

In the category of sets, an internal equivalent of the lattice of subobjects is the collection of maps from a set S to the subobject classifier $\Omega = 2$, which corresponds to the power set of S. This generalises to an arbitrary topos by replacing the collection of maps with the exponential. As opposed to $\mathrm{Sub}(C)$, the exponential Ω^C is an object in any topos and any object C. The internal lattice structure is then given by symmetric, associative maps

$$\wedge, \vee : \Omega^C \times \Omega^C \to \Omega^C$$

'satisfying' the equations defining a lattice (3.3) in the sense that they make the corresponding diagrams commute. Similarly, one can define implication in purely equational and hence diagrammatic terms.

For details on distinction between internal and external algebras and logic, see section 7.4 in Goldblatt [2], and for construction of an internal Heyting algebra see section IV.8 in Mac Lane and Moerdijk [6]. We have, in fact, defined an internal Heyting algebra on the subobject classifier $\Omega \cong \Omega^1$ given by $(\frown, \smile, \Rrightarrow)$.

We say that a topos is *Boolean* if the Heyting algebra of subobjects of any object in the topos is in fact a Boolean algebra. We illustrate that there is a strong connection between internal and external algebras with the following result.

Proposition 3.25. *Let $\mathcal{E}$ be a topos with subobject classifier Ω. Then $\mathcal{E}$ is Boolean if and only if $\Omega \cong \Omega^1$ is an internal Boolean algebra.*

The proof of this fact can be found in Proposition 1 in section VI.1 in Mac Lane and Moerdijk [6, p. 270].

3.3 The algebra of sheaves

The main aim of this section is to construct a subobject classifier for sheaves. We begin by quoting two important results.

Theorem 3.26. *The category of presheaves $\mathbf{PSh}(X)$ is a topos for any topological space X.*

This follows from a more general fact that for any small category $\mathcal{C}$, the functor category $\mathbf{Set}^{\mathcal{C}^{op}}$ is a topos. Existence of finite limits in such category follows from the fact that limits are taken 'pointwise', meaning that the limit of each functor F can be defined in terms of limits of its components $F(U)$ in $\mathbf{Set}$ for each $U \in \mathcal{C}^{op}$. Then, since $\mathbf{Set}$ has finite limits, so does the functor category. To see how exponentials can be constructed, see section I.6 in Mac Lane and Moerdijk [6], in particular the proof of Proposition 1 on page 46. For construction of the subobject classifier, we once again refer the reader to Mac Lane and Moerdijk, section I.4 [6, p. 37].

Proposition 3.27. *The category of sheaves $\mathbf{Sh}(X)$ has all finite limits, and exponentials for any pair of objects.*

Existence of finite limits is Proposition 2 in section II.2 [6, p. 71], while Proposition 1 in section II.8 [6, p. 97] guarantees that the exponential of presheaves F^P is in fact a sheaf whenever F is a sheaf, both in Mac Lane and Moerdijk.

Since the category of presheaves is a topos, we expect the sheaves to form a topos as well. In fact, the by Proposition 3.27 sheaves are almost a sheaf, we are only missing the subobject classifier. We will next reverse engineer what a subobject classifier should look like in $\mathbf{Sh}(X)$, and then prove that it indeed is one.

We begin by noting that the presheaf sending each open set of X to the one element set 1 is terminal in $\mathbf{PSh}(X)$, let us denote it by $\mathbb{1}$. Then $\mathbb{1}$ is rather trivially a sheaf. Next, suppose $(\Omega, \mathtt{t})$ is a subobject classifier for $\mathbf{Sh}(X)$. That is, Ω is a sheaf and $\mathtt{t}$ is a natural transformation $\mathtt{t} : \mathbb{1} \to \Omega$. Suppose further that $\eta : F \to G$ is a natural transformation of sheaves F and G which is monic, that is, η is a subobject of G. Let $U, V \in \mathcal{O}(X)$ such that $U \subseteq V$. We then have a commutative diagram.

$$\begin{tikzcd}
F(V) \arrow[r, "\eta_V"] \arrow[d] & G(V) \arrow[d] \arrow[ddd, bend left, "\chi_{\eta_V}"] \\
F(U) \arrow[r, "\eta_U"] \arrow[d] & G(U) \arrow[d, "\chi_{\eta_U}"] \\
\mathbb{1}(U) \arrow[r, "\mathtt{t}_U"] \arrow[d, equal] & \Omega(U) \\
\mathbb{1}(V) \arrow[r, "\mathtt{t}_V"] & \Omega(V) \arrow[u]
\end{tikzcd}$$

The unlabeled maps in the top and bottom squares are the restrictions, in the middle square the unlabeled map is the unique function from $F(U)$ to the one element set. Hence for each open subset $U \subseteq V$, the map $\mathtt{t}_U$ picks an element in $\Omega(U)$ such that $\chi_{\eta_U}(\eta_U(q))$ is equal to that element for each $q \in F(U)$. Moreover, this is done in a consistent way, so that $(\chi_{\eta_V}(s))\,|_U = \chi_{\eta_U}(s|_U)$ for each $s \in G(V)$.

We thus have a family of sections $(\chi_{\eta_V}(s))\,|_U$ for $U \in \mathcal{O}(V)$ which agree on the overlaps of the open subsets of V. Since Ω is a sheaf, there is a unique section $r \in \Omega(V)$ such that $r|_U = \chi_{\eta_U}(s|_U)$, and we must have $\mathtt{t}_V = \chi_{\eta_V}(s) = r$ whenever $s \in \mathrm{im}\ \eta_V$, for $s \in G(V)$. Notice that, by a slight abuse of notation, we identify the map from the one-element set with its value at the unique element of the set. Similarly, $r|_U = \mathtt{t}_U$ whenever $s|_U \in \mathrm{im}\ \eta_U$. Since we can repeat this process for any open set of V, we get a poset structure on $\Omega(V)$ corresponding one-to-one with $\mathcal{O}(V)$.

These observations suggest that we define $\Omega(V)$ as the set of all open sets of V. The restriction $\Omega(V) \to \Omega(U)$ is then given by the intersection $(S)|_U = S \cap U$. Since the truth map $\mathtt{t}_V$ picks the maximal element of the lattice for each open V, we have $\mathtt{t}_V = V$. It follows that for a subobject $\eta : F \to G$, the characteristic morphism $\chi_{\eta_V} : G(V) \to \Omega(V)$ is

$$\chi_{\eta_V}(s) = \left\{ \bigcup_i U_i : U_i \in \mathcal{O}(V) \quad \text{such that} \quad s|_{U_i} \in \mathrm{im}\ \eta_{U_i} \right\}. \tag{3.28}$$

Note that if $s \in \mathrm{im}\ \eta_V$, then the union in the above definition necessarily becomes all of V, and so we indeed have $\chi_{\eta_V}(s) = \mathtt{t}_V = V$.

Lemma 3.29. *The presheaf Ω sending each open set V to its set of open sets is a sheaf.*

Proof. Let $V_i \subseteq V$ be a family of open sets such that $\bigcup_i V_i = V$. Let $U_i \subseteq V_i$ be a family of open sets such that U_i's agree on the intersections of V_i's, that is, $U_i \cap V_j = U_j \cap V_i$ for all i and j. Then

$$V_i \cap \bigcup_j U_j = \bigcup_j V_i \cap U_j = \bigcup_j V_j \cap U_i = V \cap U_i = U_i$$

for all i. Since $\bigcup_j U_j$ is the unique such set, Ω is indeed a sheaf. □

Theorem 3.30. *The sheaf Ω as in Lemma 3.29 is a subobject classifier for* $\mathbf{Sh}(X)$*, with the truth map $\mathtt{t} : \mathbb{1} \to \Omega$ given by $\mathtt{t}_V = V$.*

Proof. Let $\eta : F \to G$ be a monic natural transformation of sheaves. We claim that the characteristic morphism of η is given by $\chi_\eta : G \to \Omega$ as defined in 3.28. To see this, consider the pullback of $\mathtt{t} : \mathbb{1} \to \Omega$ along χ_η. Hence for each open set V we have the pullback diagram,

$$\begin{tikzcd} S(V) \arrow[r, "\mathtt{t}_V^*"] \arrow[d] & G(V) \arrow[d, "\chi_{\eta_V}"] \\ \mathbb{1}(V) \arrow[r, "\mathtt{t}_V"] & \Omega(V) \end{tikzcd}$$

where $S(V)$ is the subset of $G(V)$ so that for each $s \in S(V)$ we have $\chi_{\eta_V}(s) = \mathtt{t}_V = V$ and $\mathtt{t}_V^*$ is the inclusion. But $\chi_{\eta_V}(s) = V$ means the open sets U_i in 3.28 cover V, and so $s|_{U_i}$ glue together to form a unique global section in $G(V)$ which restricts to $s|_{U_i}$ for each i. By uniqueness, this section must be s, and since each $s|_{U_i} \in \text{im } \eta_{U_i}$, we must have $s \in \eta_V$. When constructing χ_η we already observed that $s \in \eta_V$ implies $\chi_{\eta_V}(s) = V$. We have thus shown that $S(V) = \text{im } \eta_V$, and so $\mathtt{t}^*$ and η belong to the same subobject.

It remains to show uniqueness of the characteristic map. But if we take $\mathtt{t}_V^* = \eta_V$ in the above diagram and require the square to be a pullback, uniqueness then follows by the discussion preceding Lemma 3.29, as after fixing $\Omega(V)$ to be the set of open sets we had no choice in defining χ_{η_V}. Hence Ω is indeed a subobject classifier for $\mathbf{Sh}(X)$. □

Corollary 3.31. *The category* $\mathbf{Sh}(X)$ *is a topos for any topological space X.*

By the results of the previous section, the lattice of subobjects $\text{Sub}(G)$ is a Heyting algebra for any sheaf G on a topological space X. This gives us another perspective on sheaves, whereas in Chapter 2 sheaves were viewed as a special kind of functors, the results of this chapter suggest that sheaves have some interesting set-like properties with their own logic and internal structure. This prompts a question; to what extent does this analogy hold? It turns out that, in some sense, the topos of sheaves is quite different from **Set**. One manifestation of this is that $\text{Sub}(G)$ is in general not Boolean. To see this, take $X = (-1, 1)$ as in Example 3.8 and take $G = \Omega$, the subobject classifier for $\mathbf{Sh}(X)$. Then every subobject of

Ω can be identified with an open subset of X, and so $\mathrm{Sub}(\Omega)$ is the poset of open sets of X. By example 3.8, this algebra is not Boolean.

Since the algebra of sheaves is in general not Boolean, there is a connection between sheaf theory and intuitionistic logic, and with constructivist mathematics more generally. For more details on this, see Chapter 8 in Goldblatt [2].

We conclude our discussion by providing a necessary and sufficient condition for the topos $\mathbf{Sh}(X)$ to be Boolean.

Proposition 3.32. *The topos of sheaves* $\mathbf{Sh}(X)$ *is Boolean if and only if open and closed sets of* X *coincide.*

Proof. By Proposition 3.25, it is equivalent to show that Ω is an internal Boolean algebra if and only if open and closed sets of X coincide. But Ω is a Boolean algebra if and only if every element of $\Omega(X)$ has a complement. Since elements of $\Omega(X)$ are open subsets of X, this happens if and only if the (set-theoretic) complement U^c of every open subset U is also open, that is, every open set is closed. The last condition says that a set is open if and only if it is closed. □

Epilogue

We have scratched the surface of sheaf and topos theories, paving the road for future studies. We first introduced sheaves as particular kind of functors useful for distinguishing local and global properties. After this we proved that sheaves can be described in purely topological terms as étale bundles. We then took a step back to introduce Boolean and Heyting algebras. We discovered that the former corresponds to our idea of a classical logic, while the latter arises for any topos. We concluded with showing that the algebra of subobjects of a sheaf on a topological space is in general not Boolean.

Since we took two related but distinct perspectives on sheaves, there are at least two possible directions for further development. The very last observation, that is the fact that the topos of sheaves in not Boolean, suggests there is a connection between sheaf theory and intuitionistic logic, leading to the question what role do sheaves play in constructivist mathematics and mathematical logic in general. It indeed turns out many results in mathematical logic can be reformulated using sheaf theory. This line of development is taken, for instance, in Chapter VI in Mac Lane and Moerdijk [6], where inter alia it is shown that the continuum hypothesis and the axiom of choice are independent of the Zermelo-Frænkel set theory axioms using methods from sheaf theory.

As remarked when introducing presheaves, there is a more general notion of a presheaf which does not need to be a presheaf on a topological space. The definition of a sheaf depended, however, on the set-like structure of the underlying space, allowing us to talk about subsets, unions and intersections, as well as on the topology of the space, allowing us to talk about open sets. Thus in order to generalise the notion of a sheaf, something more is needed that just an arbitrary category. This 'something more' turns out to be a generalisation of a topological space; indeed there is a way to define a 'topology' on a category rather than on a set. This is called a Grothendieck topology, and a category equipped with a Grothendieck topology is called a site. Presheaves and sheaves can then be defined similarly to those on an ordinary topological space. This is developed in detail in Chapter III in Mac Lane and Moerdijk [6].

Finally, a natural question to ask is whether there is a notion of a map between sheaves on different sites or topological spaces. In the case of sheaves on topological spaces, we have no choice, as a continuous map between topological spaces induces a pair of adjoint functors between the corresponding categories of sheaves. This is then taken as the definition of a map between sheaves on sites. The map is called a geometric morphism, and is introduced in Chapter VII in Mac Lane and Moerdijk [6].

One of the objectives set at the beginning of this work was to understand the precise way in which sheaves distinguish local and global properties. We have certainly been successful in achieving this aim. More than that, by studying Heyting algebras of a topos, and by learning that sheaves are in fact a topos, we have revealed that sheaves have a rich structure of their own. We have attempted to shed some light on reasons why sheaves have connections to logic and set theory; in doing so, we have hopefully convinced the reader that sheaves are more general than either of these fields, forming a universe worth exploring in its own right.

Bibliography

[1] J. L. Bell. *Set theory. Boolean-valued models and independence proofs*. Oxford University Press, 2011. ISBN: 0199609161.

[2] Robert Goldblatt. *Topoi, the categorial analysis of logic*. North-Holland Publishing Company, 1984. ISBN: 0444867112.

[3] Tom Leinster. *An informal introduction to topos theory*. arXiv:1012.5647v3. 2011. URL: https://arxiv.org/pdf/1012.5647.pdf.

[4] Tom Leinster. *Basic category theory*. Cambridge studies in advanced mathematics ; 143. 2014. ISBN: 9781107360068. URL: https://arxiv.org/pdf/1612.09375.pdf.

[5] Saunders Mac Lane. *Categories for the working mathematician*. New York: Springer-Verlag, 1971. ISBN: 0387900357.

[6] Saunders Mac Lane and Ieke Moerdijk. *Sheaves in geometry and logic. A first introduction to topos theory*. Springer, 1992. ISBN: 0387977104.

[7] Colin McLarty. *Elementary categories, elementary toposes*. Oxford: Clarendon Press, 1992. ISBN: 0198514735.

[8] Emily Riehl. *Category theory in context*. Mineola, New York: Dover Publications, 2016. ISBN: 9780486809038. URL: http://www.math.jhu.edu/~eriehl/context.pdf.

[9] B. R. Tennison. *Sheaf theory*. Cambridge: Cambridge University Press, 1975. ISBN: 0521207843.