The category of sheaves is a topos: part one

Patrick Elliott

In these two talks we will prove (finally) that the category $\mathrm{Sh}(\mathcal{C})$ sheaves of sets on a site $(\mathcal{C}, \tau)$ is a topos. We will begin in this talk by showing that the category $\mathrm{PSh}(\mathcal{C})$ of presheaves of sets on a small category $\mathcal{C}$ is a topos. From here, we will show how we can canonically upgrade a presheaf on a site $(\mathcal{C}, \tau)$ to a sheaf, using the sheafification functor. In our second talk we will use these results to show that categories of sheaves are topoi.

Recall that a topos is a cartesian closed category which has all finite limits and also possesses a subobject classifier.

1 $\mathrm{PSh}(\mathcal{C})$ is a topos

Let $\mathcal{C}$ be a small category, and write $PSh(\mathcal{C})$ for the functor category $\mathbf{Set}^{\mathcal{C}^{op}}$ of presheaves of sets on $\mathcal{C}$. Let us recall a few notions from the talk Sheaves of Sets Part One:

Recall 1.1. The category $\mathcal{C}$ embeds into $\mathrm{PSh}(\mathcal{C})$ via the Yoneda embedding: an object C is sent to the *representable presheaf* h_C.

To describe the subobjects of h_C, we found it useful to introduce the following notion: a *sieve* S on C is a set of morphisms with codomain C such that if $f \in S$ then $f \circ h \in S$ whenever the composition is defined. We showed that we can identify sieves on C with subfunctors of h_C, and hence with isomorphism classes of subobjects of h_C in $\mathrm{PSh}(\mathcal{C})$.

With these notions in hand, we arrive at a natural candidate for a subobject classifier in $\mathrm{PSh}(\mathcal{C})$, namely the presheaf defined on objects by

$$\Omega(C) := \{S \mid S \text{ is a sieve on } C \text{ in } \mathcal{C}\}.$$

and on arrows $g : C \to D$ by

$$\Omega(g) : \Omega(D) \to \Omega(C), \qquad \Omega(g)(S) =: S|_g = \{h \mid g \circ h \in S\}.$$

For any object $C \in \mathrm{ob}\,\mathcal{C}$, the set $t(C)$ of all arrows into C is a sieve, called the *maximal sieve* on C. We can therefore define a natural transformation

$$\mathrm{true} : 1 \to \Omega$$

by $\mathrm{true}_C : 1 \to \Omega(C),\ * \mapsto t(C)$.

Lemma 1.1. *The monomorphism* $\mathrm{true} : 1 \to \Omega$ *defined above is a subobject classifier in* $\mathrm{PSh}(\mathcal{C})$.

Proof. Suppose we are given a presheaf $\mathcal{F}$ on $\mathcal{C}$, and a sub-presheaf $\mathcal{G}$ of $\mathcal{F}$. For each morphism $f : C \to D$ in $\mathcal{C}$, we obtain a function $\mathcal{F}(f) : \mathcal{F}(D) \to \mathcal{F}(C)$, $x \mapsto x|_f$ in $\mathbf{Set}$ which may or may not take an element $x \in \mathcal{F}(D)$ into $\mathcal{G}(C) \subseteq \mathcal{F}(D)$. Given $x \in \mathcal{F}(D)$, we write

$$\phi_D(x) = \{g \mid \mathrm{cod}(g) = D,\ x|_g \in \mathcal{G}(\mathrm{dom(g)})\}$$

Then $\phi_D(x)$ is a sieve on D, and $\phi : \mathcal{F} \to \Omega$ is a natural transformation of presheaves. Moreover, $\phi_D(x)$ is the maximal sieve $t(D)$ iff $x \in \mathcal{G}(D)$, so the subfunctor $\mathcal{G}$ is the pullback along ϕ of the map true : $1 \to \Omega$:

$$\begin{array}{ccc} \mathcal{G} & \longrightarrow & 1 \\ \big\downarrow & & \big\downarrow \text{true} \\ \mathcal{F} & \underset{\phi}{\longrightarrow} & \Omega \end{array}$$

This shows that ϕ is a candidate for the characteristic map for the sub-presheaf $\mathcal{G}$ of $\mathcal{F}$. What remains is to show that it is unique among such natural transformations making the diagram above into a pullback.

Suppose $\theta : \mathcal{F} \to \Omega$ is another natural transformation with this property. Then, with f and x as above, the pullback condition implies that $x|_f \in \mathcal{G}(C)$ iff

$$\theta_C(x|_f) = \mathrm{true}_C(1) = t(C).$$

By the naturality of θ, this is equivalent to

$$\theta_D(x)|_f = t(C),$$

and this in turn means that $f \in \theta_D(x)$. But this implies that $\theta_D(x) = \phi_D(x)$ for all D and all $x \in \mathcal{F}(D)$. Thus ϕ is the unique natural transformation satisfying the pullback condition, and the monomorphism true : $1 \to \Omega$ is a subobject classifier for $\mathrm{PSh}(\mathcal{C})$.

□

Next we will construct exponentials in $\mathrm{PSh}(\mathcal{C})$. Recall first that the product of two presheaves is defined object-wise, so

$$(\mathcal{F} \times \mathcal{G})(C) = \mathcal{F}(C) \times \mathcal{G}(C),$$

where $C \in \mathrm{ob}\,\mathcal{C}$, $\mathcal{F}$ and $\mathcal{G}$ are presheaves of sets on $\mathcal{C}$, and the "$\times$" on the right hand side is the usual cartesian product in **Set**. Unfortunately it is not so straightforward to define exponentials in presheaves, namely because the naive element-wise definition $\mathcal{G}^{\mathcal{F}}(C) = \mathrm{Hom}(\mathcal{F}(C), \mathcal{G}(C))$ is not a functor of C.

To find a the right notion of the exponential in presheaves, we will first assume one exists and then unwind the characterising adjunction to obtain a candidate definition.

Let $\mathcal{F}$ and $\mathcal{G}$ be presheaves of sets on a small category $\mathcal{C}$. If the exponential $\mathcal{G}^{\mathcal{F}}$ exists it must satisfy the adjunction condition

$$\mathrm{Hom}(\mathcal{H} \times \mathcal{F}, \mathcal{G}) \cong \mathrm{Hom}(\mathcal{H}, \mathcal{G}^{\mathcal{F}}),$$

for every presheaf $\mathcal{H}$. In particular, when $\mathcal{H}$ is a representable functor h_C for some $C \in \mathrm{ob}\,\mathcal{C}$, this isomorphism composed with the Yoneda embedding gives

$$\mathcal{G}^{\mathcal{F}}(C) \cong \mathrm{Hom}(h_C \times \mathcal{F}, \mathcal{G}).$$

This is a well-defined presheaf, with restriction map inherited from the restriction map of h_C, so we will take it as a *definition* of the exponential $\mathcal{G}^{\mathcal{F}}$. The next lemma proves that this is not crazy.

Lemma 1.2. *The presheaf* $\mathcal{G}^{\mathcal{F}}$ *is an exponential in* $\mathrm{PSh}(\mathcal{C})$.

Proof. We need to verify the adjoint condition. First, we write an evaluation map $e : \mathcal{G}^{\mathcal{F}} \times \mathcal{F} \to \mathcal{G}$ with components

$$e_c(\theta, y) := \theta_C(1_C, y) \in \mathcal{G}(C)$$

for $C \in \mathrm{ob}\,\mathcal{C}$, $\theta : h_C \times \mathcal{F} \to \mathcal{G}$, and $y \in \mathcal{F}(C)$. It follows that e is a natural transformation. Moreover, to any natural transformation $\phi : \mathcal{H} \times \mathcal{F} \to \mathcal{G}$ we can find a unique natural transformation $\phi' : \mathcal{H} \to \mathcal{G}^{\mathcal{F}}$ such that the following diagram of natural transformations is commutative:

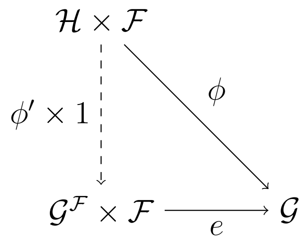

In detail, for $C \in \mathrm{ob}\,\mathcal{C}$ and $u \in \mathcal{H}(C)$, we define a natural transformation $\phi'_C(u) : h_C \times \mathcal{F} \to \mathcal{G}$ with components

$$(\phi'_C(u))_D : \mathrm{Hom}_{\mathcal{C}}(D, C) \times \mathcal{F}(D) \to \mathcal{G}(D), \qquad (f, x) \mapsto \phi_D\left(\mathcal{H}(f)(u), x\right)$$

It is clear that ϕ' is a natural transformation, and moreover by the definition of the evaluation e, we have

$$e_C(\phi'_C(u), y) = (\phi'_C(u))_C\,(1_C, y) = \phi_C(u, y),$$

so the triangle above does indeed commute. Therefore our candidate exponential is a bona fide adjoint, as required.

□

Example 1.1. Consider the case when the category $\mathcal{C} = M$ is a monoid or one-object category, Then an object of $\mathrm{PSh}(M)$ is just a set X together with a right action of M on X, and a morphism from X to another object Y is a function $f : X \to Y$ such that

$$f(xm) = f(x)m, \qquad x \in X,\ m \in M.$$

Unpacking the definitions, we see that a natural transformation $\phi : h_M \times X \to Y$ gives a functor $\phi_M : M \times X \to Y$ and that the action of M is only on M. Hence, X^Y is the set $\mathrm{Hom}(M \times X, Y)$ with action

$$(fm)(m', x) = f(mm', x), \qquad m, m' \in M, x \in X.$$

Putting lemmas 1.1 and 1.2 together we arrive at our first theorem:

Theorem 1.1. *If* $\mathcal{C}$ *is a small category then the category* $\mathrm{PSh}(\mathcal{C})$ *of presheaves of sets on* $\mathcal{C}$ *is a topos.*

2 Sheafification

Recall that a site is a pair $(\mathcal{C}, \tau)$ consisting of a small category $\mathcal{C}$ and a Grothendieck topology τ on $\mathcal{C}$. We saw in Sheaves of Sets Part 2 that the notion of a presheaf on a site can be refined to that of a sheaf, which is essentially a presheaf which respects the local data of the Grothendieck topology. In this section we will show that the canonical inclusion functor

$$\iota : \mathrm{Sh}_\tau(\mathcal{C}) \to \mathrm{PSh}(\mathcal{C})$$

has a right adjoint, namely the *sheafification* functor which canonically upgrades a presheaf to a sheaf.

Before we proceed we will introduce some new terminology.

Definition 2.1. Let $\mathcal{F}$ be a presheaf on $\mathcal{C}$, and let $S \in \tau(C)$ be a covering sieve of some object $C \in \mathcal{C}$.

1. A *matching family* for D of elements of $\mathcal{F}$ is a function which assigns to each $f : D \to C$ in S an element $x_f \in \mathcal{F}(D)$ such that

$$x_f|_g = x_{fg}$$

for all morphisms $g : D' \to D$ in $\mathcal{C}$.

2. An *amalgamation* of a matching family is an element $c \in \mathcal{F}(C)$ such that

$$x|_f = x_f$$

for all $f \in S$.

Remark 2.1. When $(\mathcal{C}, \tau)$ is a site corresponding to a topological space, a matching family of a cover is simply a collection of sections which agree on intersections, and an amalgamation is global section which realises every local section via restriction.

With the above remark in mind, we arrive at yet another definition (or rephrasing of the definition) of a sheaf on a site:

Definition 2.2. A presheaf $\mathcal{F}$ on a site $(\mathcal{C}, \tau)$ is a sheaf when every matching family for every cover of any object of $\mathcal{C}$ has as *unique* amalgamation.

Remark 2.2. As a sanity check, observe that since a sieve S on C is the same thing as a subfunctor of h_C, a matching family $f \mapsto x_f$ for $f \in S$ is the same thing as a natural transformation $S \to \mathcal{F}$. Likewise, using the Yoneda lemma, an amalgamation of a matching family $\{x_f\}_{f \in S}$ is a natural transformation $h_C \to \mathcal{F}$ such that the associated element $x \in \mathcal{F}(C)$ satisfies $x|_f = x_f$ for ever $f \in S$.

It follows that $\mathcal{F}$ is a sheaf iff, for every covering sieve of objects $C \in \mathrm{ob}\,\mathcal{C}$, any natural transformation $S \to \mathcal{F}$ lifts uniquely to a natural transformation $h_C \to \mathcal{F}$. This precisely means that the inclusion $S \to h_C$ induces an isomorphism

$$\mathrm{Hom}(S, \mathcal{F}) \cong \mathrm{Hom}(h_C, \mathcal{F}),$$

meaning that Definition 2.2 agrees with the definition of a sheaf given in previous talks.

Aside from being another interesting and useful rephrasing of the sheaf condition, the language of matching families and amalgamations shine a light on how we can upgrade a presheaf to a sheaf on a site: we must adjoin unique amalgamations to every matching family. This is done using a functor called the *plus construction*.

Definition 2.3. Let $(\mathcal{C}, \tau)$ be a site, and $\mathcal{F}$ a presheaf on $\mathcal{C}$. We write

$$\mathcal{F}^+(C) := \mathrm{colim}_{S \in \tau(C)} \mathrm{Match}(S, \mathcal{F}),$$

where $\mathrm{Match}(S, \mathcal{F})$ is the set of matching families for the cover S of C, and the colimit is taken over all covering sieves of C, ordered by reverse inclusion.

Unwinding this definition, an element of $\mathcal{F}^+(C)$ is an equivalence class of families

$$\mathbf{x} = \{x_f \mid f : D \to C \in S\}, \qquad x_f \in \mathcal{F}(D), \qquad x_f|_k = x_{fk}$$

for all $l : E \to D$, where two such families $\mathbf{x} = \{x_f \mid f \in S\}$ and $\mathbf{y} = \{y_g \mid g \in R\}$ are equivalent when there is a common refinement $T \subseteq R \cap S$ with $T \in \tau(C)$ such that $x_h = y_h$ for every $h \in T$. It is not hard to check that $\mathcal{F}^+$ so defined is a presheaf with restriction along a morphism $h : D \to C$ given by

$$\{x_f \mid f \in R\}|_h := \{x_{hf'} \mid f' \in h^*(R)\},$$

and that for $x \in \mathcal{F}(C)$ the map

$$\eta_C(x) = \{x|_f \mid f \in t(C)\}$$

defines a canonical morphism of presheaves $\mathcal{F} \to \mathcal{F}^+$.

Unfortunately, the plus construction does not give us a sheaf. But it does make a step in the right direction:

Lemma 2.1. *If $\mathcal{F}$ is a presheaf then $\mathcal{F}^+$ is a separated presheaf: all matching families have at most one amalgamation.*

Heuristically this is because the equivalence relation on $\mathcal{F}^+(C)$ identifies all amalgamations of a given matching family.

It may come as a surprise then that applying the plus construction *twice* does in fact give a sheaf with respect to a given Grothendieck topology τ.

Theorem 2.1. *If $\mathcal{F}$ is a separated presheaf on a site $(\mathcal{C}, \tau)$, then $\mathcal{F}^+$ is a sheaf.*

This shows that we can functorially enhance a presheaf to a sheaf via $\mathcal{F} \mapsto (\mathcal{F}^+)^+$. The next lemma shows that this is in fact universal:

Lemma 2.2. *If $\mathcal{G}$ is a sheaf and $\mathcal{F}$ is a presheaf, then any morphism $\phi : \mathcal{F} \to \mathcal{G}$ factors uniquely as*

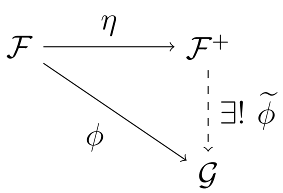

Proof. An element if $\mathcal{F}^+(C)$ is represented by a matching family $\{x_f \mid f \in S\}$ of $\mathcal{F}$ for some covering sieve S of C. For any $h : D \to C$ in S, we have

$$\eta_D(x_h) = \{x_h|_k \mid k \in t(D)\}.$$

By the matching property we have $\{x_f \mid f \in S\}|_h = \{x_{hf'} \mid f' \in h^*S\}$. But $h^*(S)$ is the maximal sieve $t(D)$ so

$$\{x_f \mid f \in S\}|_h = \eta_D(x_h), \qquad D = \mathrm{dom}(h).$$

It follows that if $\widetilde{\phi}$ were to exist, it would have to map a matching family $\mathbf{x} = \{x_f \mid f \in S\}$ to the unique element $y \in \mathcal{G}(C)$ with

$$y|_f = \widetilde{\phi}(\mathbf{x})|_f = \widetilde{\phi}(\mathbf{x}|_f) = \phi(x_f), \qquad \forall f \in S.$$

But $\mathcal{G}$ is a sheaf, so the matching family $\{\phi(x_f) \mid f \in S\}$ amalgamates uniquely to such an element $y \in \mathcal{G}(C)$, and so $\widetilde{\phi}$ exists and is unique.

□

We are now ready to assemble these facts into the main theorem of the section:

Theorem 2.2. *The inclusion functor* $\iota : \mathrm{Sh}(\mathcal{C}, \tau) \to \mathrm{PSh}(\mathcal{C})$ *has a left adjoint*

$$\boldsymbol{a} : \mathrm{PSh}(\mathcal{C}) \to \mathrm{Sh}_\tau(\mathcal{C}),$$

called sheafification, or the associated sheaf functor. Moreover, this functor commutes with finite limits.

Proof. Clearly $\mathbf{a}(\mathcal{F}) = (\mathcal{F}^+)^+$ is well defined as above. Now we have a composite morphism

$$\mathcal{F} \xrightarrow{\eta_{\mathcal{F}}} \mathcal{F}^+ \xrightarrow{\eta_{\mathcal{F}^+}} \left(\mathcal{F}^+\right)^+,$$

so be two applications of Lemma 2.2, this composite is universal among morphisms of $\mathcal{F}$ to a sheaf. Thus $\mathbf{a}$ is indeed the required left adjoint to the inclusion ι, and $\eta_{\mathcal{F}^+} \circ \eta_{\mathcal{F}}$ is the unit of the adjunction.

It remains to show that $\mathbf{a}$ preserves finite limits. To do this, it suffices to show that the plus construction preserves finite limits. To this end, observe first that for any $C \in \mathrm{ob}\,\mathcal{C}$, any covering sieve $S \in \tau(C)$, and any presheaf $\mathcal{F}$ on $\mathcal{C}$, there is a natural isomorphism

$$\mathrm{Match}_C(S, \mathcal{F}) \cong \mathrm{Hom}(S, \mathcal{F}),$$

where on the right we regard S as a presheaf. Clearly Hom preserves all limits, so the functor $\mathcal{F} \mapsto \mathrm{Match}_C(S, \mathcal{F})$ also preserves all limits. Finally, it is a general fact that finite limits commute with filtered colimits, so for any finite index category I we have

$$\begin{aligned}
(lim_{i\in I}\mathcal{F}_i)^+ &= \mathrm{colim}_{S\in\tau(C)}\mathrm{Match}_C\left(S, \lim_{i\in I}\mathcal{F}_i\right) \\
&= \mathrm{colim}_{S\in\tau(C)} \lim_{i\in I} \mathrm{Match}_C(S, \mathcal{F}_i) \\
&= \lim_{i\in I} \mathrm{colim}_{S\in\tau(C)}\mathrm{Match}_C(S, \mathcal{F}_i) \\
&= \lim_{i\in I} \left(\mathcal{F}_i\right)^+ .
\end{aligned}$$

Therefore the sheafification functor preserves finite limits.

□