Elements of Elementary Topoi

This paper will be an elementary exposition of the basic facts about of Elementary Topoi, Sheaves and Grothendieck Topologies.

An (elementary) topos is a category C, such that:

1. C is cartesian closed.
2. C has all finite limits and colimits.
3. C has a subobject classifier.

The notions of limits and colimits are assumed familiar to the reader, but definitions of a cartesian closed category and of a subobject classifier are probably not as familiar. We shall take some time in this paper to explain these notions, and give some examples of how they connect with sheaves, both over spaces and over categories.

A caretesian closed category (CCC) is a category C which enjoys the following properties:

1. C has all finite products (including the empty product)
2. For any object X, the functor $X \times -$ has a right adjoint, which we denote by $-^X$.

The motivating example of a cartesian closed category is of course the category of Sets. In this category we have that the exponential Y^X is just the functions from X into Y, and if f is a map from Y into Z, then we get a map from Y^X to Z^X by composing any function from X to Y with f.

Examples of Cartesian closed categories abound in nature, especially in computer science and logic. A more mathematical example of a CCC is that of the category of G-Sets. For any given group G, we can look at the collection of all group actions of G as a category. Products here are taken in the obvious way, but exponentials are slightly more counterintuitive. Let X and Y be G-sets, let X^Y be the collection of functions from Y to X. In order to make this into a G-set, we must specify a group action on this set. We do this by insisting that for each such function f, the following diagram commutes:

$$\begin{array}{ccc} Y & \xrightarrow{f} & X \\ \downarrow{\scriptstyle g} & & \downarrow{\scriptstyle g} \\ Y & \xrightarrow{g \cdot f} & X \end{array}$$

If one unravels this definition then it simply asks that:

$$g \cdot f(v) = g \cdot f(g^{-1} \cdot v)$$

One notices here that we introduced the use of inverses when they were note needed in our initial condition. Can one make a definition of exponentials in monoids?

We shall come back to this question later, and turn for now to a discussion of subobject classifiers in a category.

We recall that in any category, a subobject of an object X is an equivalence class of monics with codomain X, where two monics m and m' are considered to be equivalent if there exists an isomorphism f of their domains, such that $fm = m'$. This definition requires that talk of subobjects requires quantifying over all arrows into our given object. The subobject classifier simplifies this somewhat, by giving us a fixed object Ω, and allowing us to understand subobjects of any object X by analyzing the set of maps of X into Ω.

As always, the obvious example is the case of sets, where we can analyze the subsets of any set X by studying the maps of X into the object $2 = 0, 1$. Namely, if Y is any subset of X, then we can identify Y with the map ϕ from X into 2 which sends any element x to 1 if x is in Y and to 0 otherwise. We can express this map by means of the following commutative diagram:

$$\begin{array}{ccc} X & \xrightarrow{!} & 1 \\ \downarrow{\scriptstyle m} & & \downarrow{\scriptstyle T} \\ Y & \xrightarrow{\phi} & 2 \end{array}$$

Where m is the inclusion map, ! is the unique map from any object into the one object set, and T is the map from 1 into 2 which sends 0 to 1. We then see that ϕ certainly makes the above diagram commute. Furthermore, it is the smallest such one, in that it sends the fewest elements to 1 as possible. Categorically we can express this fact by noting that the above diagram is in fact a pullback. To see this, let Z be any map making the diagram commute:

$$\begin{array}{ccc} Z & \xrightarrow{!} & 1 \\ \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle T} \\ Y & \xrightarrow{\phi} & X \end{array}$$

The commutativity of this diagram asserts that $f(Z)$ is a subset of X. Thus f may be thought of a s a map g from Z to X such that gm=f. Obviously it is the unique map with this property. But this precisely states the desired diagram is a pullback.

For an arbitrary category (which we assume has 1 and has pullbacks), we can define a subobject classifier as the an object Ω, and a map T which takes 1 into Ω so that for any monic m from X into Y, there exists a unique map ϕ

making:ect classifier as the an object Ω, and a map T which takes 1 into Ω so that for any monic m from X into Y, there exists a unique map ϕ making:

$$\begin{array}{ccc} X & \xrightarrow{!} & 1 \\ \downarrow{\scriptstyle m} & & \downarrow{\scriptstyle T} \\ Y & \xrightarrow{\phi} & \Omega \end{array}$$

into a pullback. We note here that the map ϕ depends only on the isomorphism class of the monic m, ie let m,m' be two isomorphic monics with an isomorphism f. Let ϕ be the map associated with m, then to show that it is the map associated with m', we must show that in:

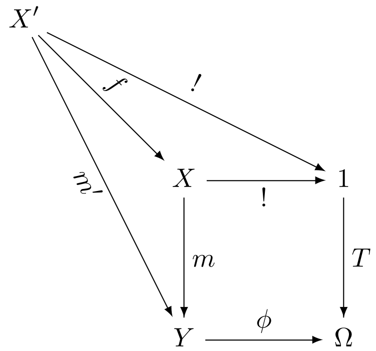

the outer square is a pullback, but this is clear, as the inner square is, and the diagram commutes. Note also that T must always be a monic map (since 1 is the final object). Thus if we have any map ϕ whatsoever into Ω, and our category admits pullbacks, then we may pullback ϕ along T to get a diagram:

$$\begin{array}{ccc} X & \xrightarrow{!} & 1 \\ \downarrow{\scriptstyle m} & & \downarrow{\scriptstyle T} \\ Y & \xrightarrow{\phi} & \Omega \end{array}$$

Because T is monic, so is m, and this allows us to conclude that there is then a bijection between isomorphism classes of monics, and maps into Ω.

Note however that the association of subobjects to a given object is functorial, that is to say that the map sending X to $Sub(X)$ is a contravariant functor into sets. It's action on maps is as follows, if f sends X to Y, we must give a map from subobjects of Y to subobjects of X. So let m be a monic from Y' into Y, then we can form the pullback:

$$\begin{array}{ccc} X & \xrightarrow{\pi} & Y' \\ \downarrow{\scriptstyle m'} & & \downarrow{\scriptstyle m} \\ X & \xrightarrow{f} & Y \end{array}$$

and we are guaranteed that m' will be monic. So the map sending m to (the isomorphism class of) m' will be our desired action on f. It can be easily checked that this association is functorial. With that said, we can now state the bijection between subobjects is natural, that is the subobject functor is isomorphic to the functor $Hom(-, \Omega)$. To see this we must check that if f is any map from Y to X then

$$\begin{array}{ccc} Sub(X) & \xrightarrow{\phi} & Hom(X, \Omega) \\ \downarrow f & & \downarrow f \\ Sub(Y) & \xrightarrow{\phi} & Hom(Y, \Omega) \end{array}$$

commutes. This follows easily from the fact that composing pullback squares gives a pullback square, so that in

$$\begin{array}{ccccc} Y' & \xrightarrow{\pi} & X' & \xrightarrow{!} & 1 \\ \downarrow m' & & \downarrow m & & \downarrow T \\ Y & \xrightarrow{f} & X & \xrightarrow{\phi} & \Omega \end{array}$$

The outer square is a pullback.

Thus we see that the subobject classifier gives us a functor $Hom(-, \Omega)$ which completely encapsulates the notion of subobjects. We say that the subobject functor is *representable* by Ω.

So now that we have some understanding of the basic ideas around the definitions of a topos, let us turn to an important example, which provided much of the original motivation for the study of topos, and is an extreme generalization of our earlier discussion of group actions.

Let C be any category, then we denote by $\hat{C}$ the category of all contravariant functors from C into the category of sets. The morphisms of this category are, of course, the natural transformations between functors. This is called the presheaf category associated with C. It enjoys several important properties (we will discover that it is in fact a topos). Firstly we note that this category will be complete and cocomplete. To define the limit of a diagram D of functors in $\hat{C}$, simply define the functor whose value on an object A is the limit of the diagram evaluated on A. For example, to find the product of F and G, simply take the functor whose value on a given object A is $F(A) \times G(A)$. Any map from A to B will give a map from the diagram over B to the diagram over A and thus maps from the corresponding limits, which defines our functor on maps.

Thus we see that in general, functors from a category C into a category D will have limits when they exist in D, so that since the category of sets has all limits and colimits, so will the category $\hat{C}$. Thus we can think of the category

$\hat{C}$ as being the category-theoretic "completion" of the category C. Part of the value of this fact is that the category C embeds naturally inside the category $\hat{C}$. The map which embeds C inside $\hat{C}$ is the yoneda embedding, which sends an object X to the functor $Hom(-,X)$. We will often denote this map by h. We take a moment here to recall the Yoneda lemma, which states that if F is an object in the category $\hat{C}$, we have:

$$Hom(h(X),F) \cong F(X)$$

This implies that h is in fact an embedding, as:

$$Hom(h(X),h(Y)) \cong h(Y)(X) = Hom(X,Y)$$

So that the functor h is clearly fully faithful.

The Hom functor is rather important in other respects as well. We define now the notion of a subfunctor of the Hom functor, which is a functor F such that $F(X) \subset Hom(X,Y)$ and such that $F(f)$ is $h(Y)(f)$ restricted to $F(X)$. A subfunctor F of the $h(Y)$ functor is also called a sieve. We often identify a sieve on Y with the image of F on the entire category, so that a sieve on Y is just a collection of maps into Y.

An important operation on Sieves is given by pulling back along maps. If f is a map from X to Y, then we can take any sieve on Y to a sieve on X via the pullback operation. So let S be a given sieve on Y, then the pullback along f will be the set of all maps into X which when postcomposed with f lie in S. Sieves will prove to be of great theoretical importance later on.

For now we turn to other aspects of presheaf categories, namely the existence of exponential objects. Let X and Y be presheafs over C. In order for the the exponential to be right adjoint to the product, we need that $Hom(X \times Z, Y) \cong Hom(Z, Y^X)$ for any Z in $\hat{C}$. To find a formula for this, we will first impose the condition that the functor Z is representable;, suppose that it is isomorphic to $h(U)$. Then using the Yoneda lemma we will have that:

$$Hom(Z,Y^X) \cong Hom(h(U),Y^X) \cong Y^X(U)$$

So thus we have that if the exponential is defined, then by the fact that it is right adjoint to the product we must have that it be defined on elements U in C by:

$$Y^X(U) = Hom(X \times h(U), Y)$$

We must show that this formula does indeed give the right adjoint, for obviously in the general case we obviously cannot expect $\hat{C}$ to be isomorphic to C. However, it is always true that any object of $\hat{C}$ is a colimit of representable objects, so that for any Z in $\hat{C}$ we have that:

$$Hom(Z, Y^X) \cong Hom(\lim_{\rightarrow} h(U_i), Y^X) \cong \lim_{\leftarrow} Hom(h(U_i), Y^X)$$

$$\cong \lim_{\leftarrow} Hom(h(U_i) \times X, Y) \cong Hom(\lim_{\rightarrow} h(U_i) \times X, Y)$$

$$\cong Hom(Z \times X, Y)$$

As desired. So therefore we have that the presheaf category does indeed have the desired exponential object.

This of course generalizes the above given example of group representations, and shows that the group structure is almost utterly redundant, as the only structure needed is that the group G is a category.

So since we already know that any presheaf category is complete, this gives us that the category $\hat{C}$ is Cartesian closed.

Thus in order to see that any presheaf category is in fact a topos, we must now only show that the presheaf category admits subobject classifiers. The subobject classifier is, firstly an object, so we must define a contravariant functor on C which we will denote Ω. So let U be an element of C, then we define $\Omega(U)$ to be the set of sieves on U. We can define if f goes from Y to X $\Omega(f)$ must be a map from sieves on X to sieves on Y. This is defined simply by mapping any sieve on Y to it's pullback along f.

It turns out that in the definition of an elementary topos the existence of colimits is redundant, it can be deduced from the existence of exponentials (actually just power sets) and limits. The proof uses Beck's theorem to prove the fact that the functor Ω^X is what is called a "monadic functor" which means that it has an adjoint, and the Eilenberg Moore category associated with the monad pair is equivalent to the image of the functor. Once we have this though, it is clear that the image is the opposite category of our original one, and since the underlying functor from the Eilenberg Moore category creates limits, we have colimits.

So thus we have a very large and interesting class of examples of topoi. However, this is only the simplest and most degenerate case of the historically motivating case of sheaves. So we move on to the other half of the topos theory, the theory of Grothendieck topoi. We start by the historically simplest example, that of sheaves over a topological space.

Let $\mathcal{T}$ be a topological space. The collection of open sets of $\mathcal{T}$ forms a poset, and thus a category which we will, by abuse of notation, also denote by $\mathcal{T}$. The contravariant functors over $\mathcal{T}$ are called presheaves on $\mathcal{T}$. If f is an arrow in $\mathcal{T}$, and $\mathcal{F}$ is a presheaf, then $\mathcal{F}(f)$ is often called a restriction map. These are very obvious things to study. A few examples are in order:

The assignment which takes U to continuous functions on U

The assignment which takes U to smooth functions on U (if $\mathcal{T}$ has a smooth structure)

The assignment which takes U to holomorphic functions on U (if $\mathcal{T}$ has a complex structure)

The assignment which takes U to locally constant functions on U

In all these cases the restriction maps are simply the restriction of functions to a smaller domain, (hence the name restriction maps). The important common feature of all these examples is that in all of these cases, the presheaves are determined by local data. The way to make this more precise is to say that, if we have any open cover U_i of U, and any collection $x_i \in \mathcal{F}(U_i)$ satisfying $x_i = x_j$ on $U_i \cap U_j$ (that is they agree after applying the appropriate restriction maps), then there is a unique x in $\mathcal{F}(U)$ which restricts to the set of x_i. It is clear that all of the above examples are in fact sheaves. The category of sheaves is a full subcategory of the category of presheaves.

The concept of presheaves over a topological space extends, as we have already seen to an arbitrary category. However, Grothendieck was interested in seeing if the notion of a sheaf could be expressed via structure of the category $\mathcal{T}$, in order that he might be able to formulate a sheaf theory (and a sheaf cohomology theory) for a class of structures more general than topological spaces. As we have seen, the main structure which allows us to understand the notion of a sheaf is the concept of an open covering. Therefore, Grothendieck defined what is called a Grothendieck topology, which is a category equipped with a notion of open covers of any open set U, which we think of as collections of maps into U. More explicitly, a Grothendieck topology on a category C is an assignment of, to any element U in the category, a collection of sieves $J(U)$ (the covering sieves) such that:

The maximal sieve is in $J(U)$

The collection is closed under all pullbacks

If $R \in J(U)$ and S is a sieve on U such that the pullback of S along any map in R is a covering sieve, then S is a covering sieve.

The category C together with it's Grothendieck topology is called a site, and once we have a site, we can easily enough define sheaves over that site. Namely, a presheaf $\mathcal{F}$ over a site C is called a sheaf if for any U, and any sieve S in $J(U)$, any natural transformation from S into F extends to a unique transformation from $Hom(-, U)$ into F. Again, the category of sheaves is a full subcategory of the category of presheaves. A category equivalent to the category of sheaves over some site is called a Grothendieck Topos.

We now come to the main point, that any Grothendieck topos is an elementary topos. We have already shown that any presheaf category is a topos, so we have given descriptions of the necessary objects, and so need only show that they are indeed sheaves. So first we must show that the (co)limits of sheaves are again sheaves. So let D be a diagram in $\hat{C}$, and P be the limit. Let S be a sieve in $J(U)$, and let us be given a map f from S to P. Then S forms a cone over the diagram when we compose with f. Thus we have that S extends to a unique map to $h(U)$ into each element of the diagram so that the triangles commute, and thus S extends into a map into the limit P. Conversely, if we have any map into the limit, it is determined by it's associated maps into the diagrams, so since those are unique, the map into P is as well.

The subobject classifier and exponentials are very similar to the case of presheaves, but I do not know of any elementary proof that they constitute sheaves.

REFERENCES

PT Johnstone, "Topos Theory", Academic Press 1977

Saunders Mac Lane, I Moerdijk, "Sheaves in Geometry and Logic", Springer Verlag, 1992.

William Fulton, Joe Harris, "Representation Theory, a first course", Springer Verlag, 1991.