Toposes *pour les nuls*

Steve Vickers
Imperial College

Introduction

The very first sentence of Mac Lane and Moerdijk [5] says:

> A topos can be considered both as a “generalized space” and as a “generalized universe of sets”.

The “generalized universe of sets” aspect of toposes is relatively easy to understand and is well documented in the literature: start with Goldblatt [1] and proceed via Mac Lane and Moerdijk [5], or MacLarty [4], to Johnstone [2]. The basic trick is to use categorical properties to characterize set-theoretic constructions in the category of sets, and thence to transfer them to other categories that are sufficiently similar.

The generalized spaces, on the other hand, though present in ideas of toposes right from their introduction by Grothendieck, are somewhat mysterious. Much of this is because the generalized universes of sets are not direct expressions of the spatial idea but represent it by a mathematical duality. My aim here is at least to present a clear picture of how intuitions of generalized spaces fit into a mathematical framework of generalized categories of sets.

To try to be clear, I shall use the word topos only for the view as generalized space. When I consider it as a generalized universe of sets, I shall call it a G-frame (standing for Giraud/Grothendieck-frame). By the duality which I shall explain, the G-frame is used to represent the topos.

Continuous functions

If $f(x)$ is a real-valued function of a real number x, we have a simple pictorial intuition of what it means for f to be “continuous” — namely that its graph has no gaps or instantaneous jumps in it. Consider, for instance, the function

$$f(x) = \begin{cases} 0 & \text{if } x < 0 \\ 1 & \text{if } x \geq 0 \end{cases}$$

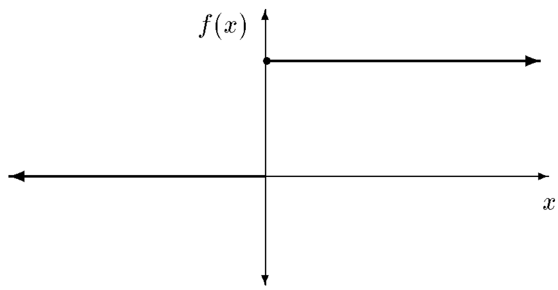

This has a gap in it — or, if you feel you could fill that in with a vertical line, it’s an instantaneous jump. This function is discontinuous at $x = 0$.

The definition in terms of drawing graphs and looking for gaps or jumps is not a rigorous one, and it was made more precise as follows. If a function is continuous at a point x_0, then there is a surrounding neighbourhood, one that goes a little way beyond x_0 on each side, within which $f(x)$ doesn’t stray too far from $f(x_0)$. How big this neighbourhood can be depends on what you think “too far” means, but as long as you are prepared to allow $f(x)$ some positive amount of latitude then you can also allow x some positive amount of latitude. To express it in the famous "$\epsilon - \delta$" formulation, f is continuous at x_0 iff

$$\forall\epsilon > 0.\exists\delta > 0.\forall x.(|x - x_0| < \delta \Rightarrow \\ |f(x) - f(x_0)| < \epsilon)$$

Topological spaces

That looks very heavily dependent on the fact that we are working with real numbers, but really it depends solely on the notion of "neighbourhood". For the real numbers, we say that a set N is a neighbourhood of x_0 iff it contains not only x_0 itself, but also all the numbers close to x_0 within some unspecified positive distance. Then f is continuous at x_0 iff for every neighbourhood N of $f(x_0)$ (corresponding to ϵ) there is a neighbourhood M of x_0 (corresponding to δ) that is mapped into N by f. It is possible to axiomatize the notion of neighbourhood in an abstract way and thereby define continuity for functions in contexts other than the real line. An alternative axiomatization, and in many ways a more useful one, is of open sets, those that are neighbourhoods of all their elements. Then a function is continuous (everywhere) iff the inverse image of every open set is open.

A set equipped with such a structure of neighbourhoods, or, equivalently, of open sets, is called a *topological space*.

Sheaf = continuous set-valued function

What could it mean for a set-valued function $S(x)$ to be continuous? Let us try to apply the same intuition as we had before. S is continuous at x_0 iff there is some neighbourhood of x_0 within which $S(x)$ doesn't stray too far from $S(x_0)$. What neighbourhood is needed depends on what part — which element — of the set $S(x_0)$ we are looking at. What we want to formalize is that each $y_0 \in S(x_0)$ is still in $S(x)$ as long as x is close to x_0, but that if we stray too far we start gaining or losing elements. Hence there is some neighbourhood N of x_0 and some selection $y_x \in S(x)$ for the values $x \in N$ such that $(y_x)_{x \in N}$ represents y_0 "as you move around a little". Such a "continuous set-valued function" is called a sheaf. What I have written is admittedly still vague, but it is the idea behind the definition of sheaf as local homeomorphism that you will see in the standard texts. Here is an example of a sheaf on the real line:

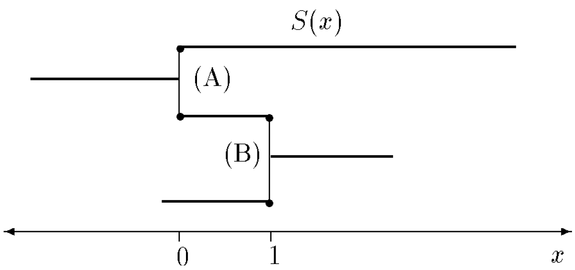

Notice the forking structure at (A) ($x = 0$) and (B) ($x = 1$). The two blobs above $x = 0$, near (A), represent two distinct elements of $S(0)$ (and there's a third one, not blobbed, lower down on the lowest thick horizontal line). To the right, for x just greater than 0, the two blobs maintain their separate identities along horizontal lines. However, to the left, for x just less than 0, they become equal so that $S(x)$ has only two elements instead of three. (B) is similar.

We also get a natural notion of morphism between sheaves S and T: it will have for each x a function from $S(x)$ to $T(x)$, together with conditions to ensure that these functions fit together in a continuous way.

A category of sheaves is a generalized universe of sets

If X is a topological space, then the sheaves over X (the continuous set-valued functions on X) are the objects of a category $\mathcal{S}X$. It is a G-frame.

Since a sheaf is a parameterized set ($S(x)$ parameterized by point x), we can consider doing set-theoretic constructions on sheaves by doing them pointwise on the sets. For instance, if $S(x)$ and $T(x)$ are two sheaves, then we can define a product $(S \times T)(x) = S(x) \times T(x)$. This is indeed still a sheaf, as are the results of a number of constructions such as disjoint unions — even infinitary ones — and (to use categorical language) equalizers and coequalizers. Some constructions, such as function spaces and power sets, do not yield sheaves when applied pointwise. Nonetheless, it turns out that there are sensible interpretations of these constructions in the category of sheaves (making it an "elementary topos").

Geometric constructions

The constructions that do yield sheaves when done pointwise are called geometric, and there is a corresponding geometric fragment of logic (its connectives are $\bigvee, \wedge, \exists$ and $=$). Categorically, the geometric constructions are those that can be described as colimits of finite limits.

Given two categories of sheaves, we are particularly interested in the functors between them that preserve the geometric constructions (i.e. that preserve all colimits, and finite limits). I shall call such functors "G-frame homomorphisms".

Subsheaves of 1 correspond to open sets

If X is a space, then the sheaf 1 — the terminal object in $\mathcal{S}X$ — has $1(x)$ a singleton for all x. This is because 1 is a finite (nullary) product, and hence geometric, so it is constructed pointwise. A subsheaf S of 1 — a subobject in $\mathcal{S}X$ — has $S(x)$ always a subset of a singleton, and so is determined by the set of points x at which $S(x)$ contains its only possible element. By the continuity condition, this set is an open subset of X. In fact, the subsheaves of 1 correspond exactly to the open subsets of X.

Note that we instantly lose classical logic! If U is an open subset of X, its complement might not be open (for instance, in the real line the set of negative reals is open, but its complement, the set of zero-or-positive reals, is not). As a consequence, if we consider U as a subsheaf of 1 we don't necessarily have another subsheaf V such that $U \cup V = 1$ and $U \cap V = \emptyset$ ($\cup$ and $\cap$ are interpreted pointwise). Thinking of $\mathcal{S}X$ as a generalized universe of sets, the subsheaves of 1 are the subsets of a singleton and correspond to logical truth values. The upshot is that we lose the law of excluded middle, $P \vee \neg P$.

Categories of sheaves are dual to spaces (more or less)

An ordinary set S is a disjoint union of copies of the singleton set 1 — one copy for each element of S —, and a sheaf is a colimit of subsheaves of 1. Hence any G-frame homomorphism from one category of sheaves, $\mathcal{S}Y$, to another, $\mathcal{S}X$, is defined by its action on the subsheaves of 1. Moreover, since it preserves finite limits, it preserves monomorphisms and hence maps the subsheaves of 1 in one category to subsheaves of 1 in the other and hence gives a function from the open subsets of one space to the open subsets of the other. One can follow the argument further to show that this function preserves finite intersections and arbitrary unions and is exactly the inverse image function on open sets for a continuous map from X to Y. In other words, we can represent continuous maps between spaces as, exactly, functors between the categories of sheaves that preserve the geometric constructions.

Well, that's not quite true. But it works if X and Y are "sober", which decent spaces are. Continuous maps from X to Y are equivalent to G-frame homomorphisms from $\mathcal{S}Y$ to $\mathcal{S}X$. (Note the reversal of direction! That's why it's a "duality".)

Generalizing the duality

I haven't defined "G-frame" exactly, but there are plenty of categories other than the categories of sheaves that support the geometric constructions sufficiently well to be admitted as G-frames. They generalize the categories of sheaves as generalized categories of sets, and we try to understand the above equivalence as generalizing on the space side.

The morphisms between the generalized spaces, which generalize continuous maps and by definition are dual to G-frame homomorphisms between the generalized universes of sets, are called *geometric morphisms*.

The "generalized space" of sets

The most obvious example comes out of the very notion of sheaves. We have already motivated them as "continuous set-valued functions", so a sheaf S on X should be a continuous map from X to "the space of sets", hence, to generalize the equivalence, a G-frame homomorphism from a G-frame $\mathcal{E}$ to $\mathcal{S}X$. Because S is just an object of $\mathcal{S}X$, and this is to determine the entire G-frame homomorphism from $\mathcal{E}$, $\mathcal{E}$ should be generated geometrically by an object S_0 — every other object is constructed from S_0 by the geometric constructions. I shall not go into the details of the structure of $\mathcal{E}$, but it can be constructed. It is what Johnstone calls the "object classifier". I shall denote it $\mathcal{S}[set]$. It has an object S_0, the *generic* set, which has no properties other than those which follow from the fact that it is a set. If you wanted to describe a way of constructing something or other out of an arbitrary set, and you started your description by saying "Let S_0 be a set ...", then — at least if your construction is geometric — S_0 is really the generic set in $\mathcal{E}$. It has no known elements (as morphisms from 1 to S_0), but on the other hand it has no isomorphism with the initial object $\emptyset$. $\mathcal{S}[set]$ has the property that G-frame homomorphisms from $\mathcal{S}[set]$ to a category of sheaves $\mathcal{S}X$ are equivalent to objects of $\mathcal{S}X$.

An ordinary category of sheaves, $\mathcal{S}X$, corresponds to an ordinary space, X. Our generalized category of sheaves $\mathcal{S}[set]$ does not correspond to an ordinary space $[set]$ of sets, but nonetheless we can see a sense in which we know what the points of $[set]$ are — they are just sets, and a continuous map from X to $[set]$ (defined as a G-frame homomorphism from $\mathcal{S}[set]$ to $\mathcal{S}X$) is a continuous set-valued function. It's just that the extra structure on this space, used to define what "continuous" means, is not the usual topological structure, expressed in terms of neighbourhoods or open sets.

We can understand this in terms of the core, generating structure of the category of sheaves, from which everything else is constructed geometrically. Continuity between topological spaces can be defined in terms of the open sets, which correspond to the subsheaves of 1, so there is no real need to consider the whole of $\mathcal{S}X$: the subobjects of 1 are enough. In $\mathcal{S}[set]$, on the other hand, the essential part is the generic set S_0 which lies beyond the subobjects of 1. For ordinary topological spaces the subobjects of 1 are all we need to consider; for generalized spaces the rest of the G-frame is also important.

Classifying toposes for geometric theories

Suppose a logical theory is presented "using the geometric constructions". There are various ways of imposing appropriate presentational constraints, but a simple one is to say that the theory is presented using sorts, function symbols (including constants), predicate symbols (including propositions) and axioms of the form $\phi \vdash \psi$ where ϕ and ψ are logical formulae built up from the language ingredients by using finitary conjunction $\wedge$, arbitrary disjunction $\bigvee$, existential quantification $\exists$, and equality $=$.

If T is a geometric theory, I shall write $[T]$ for the generalized "space of models" of T, known technically as the classifying topos of T. Its points are the models of T, and, just as for $[set]$, the structure needed to capture the idea of continuity is given by a G-frame $\mathcal{S}[T]$. It is constructed by taking a "generic model" of T and adding everything that can be constructed geometrically from it.

You can find the construction in more detail in Johnstone or in Mac Lane and Moerdijk (described as the classifying topos, but this is topos as generalized universe of sets, i.e. G-frame). What makes it work is that for any G-frame $\mathcal{E}$, the models of T in $\mathcal{E}$ are equivalent to G-frame homomorphisms from $\mathcal{S}[T]$ to $\mathcal{E}$: first map the generic model of T in $\mathcal{S}[T]$ to the given model in $\mathcal{E}$, and extend this to the whole of $\mathcal{S}[T]$ by applying the same geometric constructions on both sides.

Continuity = geometricity + genericity

Taking $\mathcal{E}$ to be $\mathcal{S}[T']$, then a G-frame homomorphism from $\mathcal{S}[T]$ to $\mathcal{S}[T']$ is a model of T in $\mathcal{S}[T']$, in other words a model of T that's constructed geometrically from the generic model of T' (because everything in $\mathcal{S}[T']$ is made geometrically from the generic model of T'). How do you describe this? You say "Let M be a model of T'" — M is now your generic model of T', since you have assumed nothing about it other than what follows from its being a model of T' — and then you proceed to construct a model of T geometrically. Of course, this is just a particularly disciplined way of getting functions from the class of models of T' (points of $[T']$) to the class of models of T (points of $[T]$), so we can think of it as a "continuous" map from $[T']$ to $[T]$ and "continuous" now refers to the genericity of the description and the geometricity of the construction.

As an illustration, an object of $\mathcal{S}[T]$ is a G-frame homomorphism from $\mathcal{S}[set]$ to $\mathcal{S}[T]$, i.e. geometric morphism from $[T]$ to $[set]$. Each is defined by saying "Let M be a model of T" and then, geometrically, constructing a set.

Geometric theories for spaces

If X is a topological space, then we can present a propositional geometric theory T as follows. It has no sorts, functions or predicates except for some nullary predicates (i.e. propositional symbols), namely a proposition P_U for each open set U and axioms

$$\begin{array}{ll} P_U \vdash P_V & \text{(whenever } U \subseteq V) \\ true \vdash P_X & \\ P_U \wedge P_V \vdash P_{U \cap V} & \\ P_{\bigcup S} \vdash \bigvee_{U \in S} P_U & \end{array}$$

(S here is any set of opens — note the possibly infinite disjunction!)

Now any point x gives a model for this propositional theory, assigning the value true to the proposition P_U iff $x \in U$. If X is sober, then in fact the points of X correspond exactly to the models of the theory, i.e. the points of the classifying topos $[T]$. It can also be proved that the category $\mathcal{S}X$ of sheaves over X is equivalent to $\mathcal{S}[T]$: for any G-frame $\mathcal{E}$, the G-frame homomorphisms from $\mathcal{S}X$ to $\mathcal{E}$ are equivalent to models of T in $\mathcal{E}$. (Bearing in mind that propositional symbols must be interpreted in a model as subobjects of 1, the generic model of T in $\mathcal{S}X$ interprets each P_U as the subsheaf of 1 corresponding to the open set U.) All in all, we might as well consider $X = [T]$.

Locales classify propositional geometric theories

The sobriety condition on the spaces can be circumvented entirely if you deal directly with propositional geometric theories instead of spaces. The classifying toposes for propositional geometric theories are called *locales*, and these are the more genuinely spatial of the generalized spaces. They are the toposes that are determined by the subobjects of 1 with no need to consider the rest of the G-frame.

The tendency in topos theory is to study locales in place of topological spaces, and you can read more about them in Vickers [6] and Johnstone [3].

Summary

A topos-as-generalized-space is the space of models for a geometric theory. The whole story is in answer to the question of what "space" means here.

If the theory is presented as T, then its topos — its classifying topos — is denoted $[T]$.

The points of $[T]$ are the models of T.

Associated with each topos $[T]$ is a generalized category of sets, $\mathcal{S}[T]$. It is got by taking a "generic" model of T and including everything that can be constructed "geometrically" from it.

The geometric morphisms (generalizing continuous maps) from $[T]$ to $[T']$ are the transformations of models of T into models of T' that can be described generically and geometrically. They are equivalent to models of T' in $\mathcal{S}[T]$, or to functors from $\mathcal{S}[T']$ to $\mathcal{S}[T]$ that preserve finite limits and arbitrary colimits. In effect, "space" of models means "class" of models together with whatever structure is needed to impose this constraint on the transformations.

Sober topological spaces are the spaces of models for (certain) propositional geometric theories. Geometric morphisms between their classifying toposes correspond to continuous maps between the spaces. A sheaf over X is a continuous map from X to $[set]$.

Locales are the classifying toposes for general propositional geometric theories.

References

[1] Robert Goldblatt. *Topoi: The Categorical Analysis of Logic.* North-Holland Publishing Company, Amsterdam, 1979.

[2] Peter T. Johnstone. *Topos Theory.* Academic Press, London, 1977.

[3] Peter T. Johnstone. *Stone Spaces. Vol. 3, Cambridge Studies in Advanced Mathematics.* Cambridge University Press, 1982.

[4] Colin McLarty. *Elementary Categories, Elementary Toposes.* Oxford University Press, Oxford, 1992.

[5] Saunders Mac Lane and Ieke Moerdijk. *Sheaves in geometry and logic.* LNMS. Springer Verlag, 1992.

[6] Steven J. Vickers. *Topology Via Logic. Vol. 5, Cambridge Tracts in Theoretical Computer Science.* Cambridge University Press, 1989.