

Topological Spaces and Sheaf Axioms

Linus Vepstas

29 August 2020

Abstract

This chapter presents the conventional (pre-)sheaf axioms and then demonstrates how these describe the graphical/tensor structures that are the primary topic of this series of chapters.

This is a part of a sequence of articles exploring inter-related ideas; it is meant to provide details for a broader context. The current working title for the broader text is “Connectors and Variables”.

Pre-sheaves

This text reviews the concept of a pre-sheaves and sheaves, starting from the conventional definition, and then demonstrating how this applies to the system of knowledge representation being developed in these texts. It is primarily an exercise in examining the axioms using different kinds of notation, as notational differences appear to present one of the primary stumbling blocks to the development of AGI.


The conventional definition of a (pre-)sheaf begins as follows: Let X be a topological space and let C be a category. A presheaf F on X is a functor with values in C having a certain set of properties ... (to be reviewed below). The conception of “seeds” being advanced here play the role of a topological space; this is the primary driver of much of the mathematical machinery. Thus, the re review here starts with topological spaces.

Topological space

Let’s review the conventional definition of a topological space.¹ During the review, keep in mind the conception of “seeds”, as presented earlier. As a reminder, here are

¹This entire section painfully belabors something that is relatively obvious to those who know algebraic topology: it is describing something very much like a “simplicial complex”, and points out that it is a topological space. Perhaps we could go farther: describe the analog of “CW-complexes”, by describing gluing and a weak topology. Yet there are differences: the connectors on seeds are typed, whereas the vertexes of simplicial complexes are simply-typed. The seeds also resemble dual-graphs, in that the connectors connect up neighbors, and so denote where two things come together at a boundary. So we cannot just make this jump, as simplicial complexes are almost but not quite the correct concept. It would also put too great a burden on the reader to tell them that they must first study algebraic topology. The goals here are entirely different: the goals are to gain insight into knowledge representation and AGI.

diagrams of a single seed, a pair of connected seeds, and a hedge-hog: four interconnected seeds.



These begin to look like the open sets of a topology, if one considers the open circles to be the “open” boundary of the set, and the black circles to be the interior. The primary conceptual issue is that topologies are conventionally framed in terms of “sets”, which then presumes “set theory” underneath it all. And yet nothing about a single seed looks particularly set-like. Yes, the middle diagram above is obviously a “union” of two seeds, and the diagram on the right is the “union” of four seeds. But they are not sets.

No matter. A topology is defined in terms of a collection of axioms, and one should argue that anything obeying those axioms is a topology, irregardless of whether the underlying objects are actually sets, or not.² What are the axioms of a topology? Let’s quote directly from Wikipedia, again: “A *topological space* is an ordered pair (X, τ) , where X is a set and τ is a collection of subsets of X , satisfying the following axioms:...”. In the current context, the goal is to cross out the word “set” and replace it by “hedgehog” or “mesh” (the diagram above shows four interconnected seeds; in general it could be a mesh of trillions, or even an infinite mesh. An infinite mesh might be countable, if it was constructed recursively, or even uncountable, if one is willing to entertain a mesh that interconnects elements in a Cantor set.) One point to be careful about, in this analogy: here, X is meant to be the fully-assembled mesh; it is NOT a collection of seeds!³

Let’s continue:

- “The empty set and X itself belong to τ .” Of course: the empty seed is just nothing at all: the seed with no connectors, and no center. And sure, let us assign the the hedgehog to τ .
- “Any arbitrary (finite or infinite) union of members of τ still belongs to τ .” Of course, although one must be careful in reading the word “union”, here. There are two warring senses. In one sense, the word “union” can be taken to mean “the connection or connecting-together of two or more seeds, if they are connectable”. This is meant to be the re-gluing together of pieces of a mesh that were cut out: if one takes some scissors, and cuts out pieces from the mesh, and then joins them back together again, it is still a sub-portion of the mesh. In the other sense, the union is a disjoint union: supposed the two scissored pieces of the mesh were far

²In a certain sense, this is precisely what the Yoneda lemma states. But that sense is “sideways” from the direction we wish to travel in. The Yoneda lemma is also sufficiently abstract that it cannot be made a pre-requisite to understanding this text. It does, however, provide the sign-post to a “classical” formalization of the discussion here.

³Of course, X could be just a collection of seeds, if that is the topological space that one wished to work with; but that would be a very special case, and not the general case.

apart from each other. One is then left with a *set* of two pieces, that cannot be joined together. This is still a union. If one carefully examines the concept of a “union” in conventional set theory, it *always* means “glue them together if you can”. Take, for example, a Venn diagram shown to elementary-school students: the union is always the gluing-together, the overlapping bits being the joint. The overlapping bits are mashed together: this is the “union” of set theory: it is a gluing-together of the common, shared parts.

- “*The intersection of any finite number of members of τ still belongs to τ .*” The reader was already invited to think of taking scissors to the mesh to disconnect it, and also of gluing together pieces. How, exactly, can the word “intersection” be interpreted? The intent is that it be read in the intuitively-obvious sense. If we have two fragments of a mesh, and they have points and links in common, then the intersection are those points and links that are in common between those two fragments. Cut edges then become open connectors: as always, a connector is a “half-edge”.
- “*The elements of τ are called **open sets** and the collection τ is called a **topology** on X .*” That is the conclusion, with emphasis in the original.

It would appear that a mesh: that is, a graph whose edges can be scissored into half-edges aka connectors forms a topology, in that it obeys the axioms of a topology (ignoring the fact that hedge-hogs are not actually sets). This is meant to be self-evident, so that it becomes straight-forward to develop notational devices. For example, given a hedgehog U , the boundary ∂U of U is the collection of connectors on U that have not been connected up, yet. The interior then are the vertexes that are fully connected (and so also, the edges connecting them.)

Non-topological concerns

Of course, the mesh is a fair bit richer than the set: the edges are meant to be typed, whereas in set theory, all sets have the same type: set theory is “simply typed”, there is only one type, and that is the type of the set.

Sets, as commonly presented in textbooks on general topology, are all about uncountable collections of “points”. The imposition of uncountability leads to a wide variety of careful distinctions, e.g. between Hausdorff spaces, which behave in a generally intuitive fashion, and other spaces, which often do not. In the current texts, the number of connectors is meant to be finite, or at least always countable. The size of the graph is meant to be finite, or at most recursively infinite. Expanding beyond this would, of course, lead to a wild bestiary of strange things.

Another, more arcane distinction is that between Boolean algebras and Heyting algebras, which is foundational for the category of Frames and Locales, which in turn form the basic examples for Topoi.

If one cuts up a graph into disjoint, unconnected parts, by cutting a collection of edges, then by naive reasoning, the cut edges belong to neither the one nor the other subgraph. Thus, applying naive notions about cutting result in an algebra of graphs that is not a Boolean algebra: the cut edges are neither in the set, nor in its complement.

The sophisticated way of tackling this is with the concept of a “Subobject Classifier”, which, for Boolean algebras, just boils down to the ordinary membership predicates: functions that take a value in $2 = \{0, 1\} = \{F, T\}$. Either something belongs or doesn’t. It is, or it ain’t. The subobject classifier is conceptually heavyweight, though. The whole point of having half-edges is that it is conceptually simple. A jigsaw-puzzle piece with unconnected tabs is a “thing”, and there is no particular problem of asking whether the unconnected tabs belong to it, or to its set-complement.

TODO: Refining the discussion of Frames and Locales, as it applies to the jigsaw-puzzle metaphor, is arguably an important pre-requisite for the discussion of pre-sheaves and sheaves. However, this is held off until later, as the level of abstraction is more demanding.

Topological conclusions

At any rate, the goal here would seem to have been met: meshes can be made to look like a topology. This means that many of the mechanisms and concepts intended for topologies can be carried forward, transported and applied to meshes.

Presheaf axioms

The conventional definition begins as follows. Let X be a topological space and let C be a category. A presheaf F on X is a functor with values in C having the following properties (quoted from Wikipedia):

- “For each open set U of X , there corresponds an object $F(U)$ in C .”
- “For each inclusion of open sets $V \subseteq U$, there is a corresponding morphism $\text{res}_{V,U} : F(U) \rightarrow F(V)$ in the category C .”
- “For every open set U of X , the restriction morphism $\text{res}_{U,U} : F(U) \rightarrow F(U)$ is the identity morphism on $F(U)$.”
- “If we have three open sets $W \subseteq V \subseteq U$, then the composite $\text{res}_{W,V} \circ \text{res}_{V,U}$ equals $\text{res}_{W,U}$.”

These are rather abstract statements; before they can be given a concrete meaning, an exploration of what the category C might be is required. Since, again, the task at hand is that of providing meaning, the category C can be taken to be “the category of meanings”. But what is a meaning? To prevent the discussion from devolving into abstract nonsense, it is best to make C concrete.

To that end, a suitable example might be the vectorial/tensorial meanings provided by word-sense disambiguation (*e.g.* the “Mihalcea algorithm”), or the compositional concept of meaning that is commonly invoked in artificial neural nets (*e.g.* Skipgram models of language.)

To recap these two different models of meaning ... (is this wise here or does this belong in a different chapter?)

TODO this section is obviously, wildly incomplete. This needs to be started with a realistic non-abstract example.

References