Math 203B: Algebraic Geometry
UCSD, winter 2016, Kiran S. Kedlaya
Sheaves in the language of categories

Before giving the construction of schemes, let me take a moment to briefly recall the definition of sheaves. In the process, I would like to take the opportunity to bake in a little bit of the language of categories; this will both give me an excuse to take this a bit slowly, and provide an opportunity to get the language set up right for later use (e.g., when we start talking about sheaves of modules instead of rings).

Categories

A *category* consists of:

- a "collection" (more on this below) of sets, called the *objects* of the category;
- for every pair X, Y of (not necessarily) distinct objects, a set $\mathrm{Mor}(X, Y)$, called the *morphisms* from X to Y;
- for every object X, an element $\mathrm{id}_X \in \mathrm{Mor}(X, X)$, called the *identity morphism* on X;
- for every triple X, Y, Z of objects, a binary operation

  $$\circ : \mathrm{Mor}(Y, Z) \times \mathrm{Mor}(X, Y) \to \mathrm{Mor}(X, Z),$$

  called *composition*, which is associative and has the id_* as identities on either side.

The canonical example is supposed to be the category **Set** of all sets, in which $\mathrm{Mor}(X, Y)$ is the set of all functions from X to Y, id_X is the usual identity morphism, and $\circ$ is the usual composition of functions (whence the order of the factors). This example illustrates a foundational nuisance: recall that according to the usual rules of set theory, there is no set of all sets![1] One can finesse this by defining a concept of *classes* of sets, which obey most but not all of the axioms of sets (notably omitting the power axiom). This is sufficient for our purposes.[2]

While categories in principle could be defined totally abstractly, using more or less arbitrary associative binary operations, the examples we will have most use for have as their objects just sets equipped with certain extra structures, and for morphisms the maps between those sets that respect the extra structures. For example:

[1] The usual explanation for this is Russell's paradox: if there were a set X of all sets, then one could form $Y = \{U \in X : U \notin U\}$ for which neither $Y \in Y$ nor $Y \notin Y$ holds. A fancier explanation uses Cantor's diagonalization theorem: for no set X does there exist a bijection between X and its power set $P(X)$, but if X were the set of all sets such a bijection would exist by the Schröder-Bernstein construction.

[2] What this doesn't let you do is define, say, a category of all categories. There are good reasons to want to do this, especially in algebraic topology, but for this one must build more of a hierarchy of set-like objects and develop a theory of ∞-*categories*.

- **Ab**: abelian groups, group homomorphisms;
- **Ring**: rings, ring homomorphisms;
- **Top**: topological spaces, continuous maps.

This observation naturally leads us also to the analogue of a "function between categories", which is called a *functor*. Given two categories $\mathcal{C}_1, \mathcal{C}_2$, a *covariant functor* $F : \mathcal{C}_1 \to \mathcal{C}_2$ consists of the following "functions" (again modulo set-theoretic difficulties):

- one assigning to each object $X \in \mathcal{C}_1$ an object $F(X) \in \mathcal{C}_2$;
- one assigning to each morphism $f \in \mathrm{Mor}_{\mathcal{C}_1}(X, Y)$ a morphism $F(f) \in \mathrm{Mor}_{\mathcal{C}_2}(F(X), F(Y))$ in a fashion preserving identities and composition.

For example, there are "forgetful functors" **Ab** $\to$ **Set**, **Ring** $\to$ **Ab**, **Top** $\to$ **Set** that retain some underlying structure and forget the rest. One can also define a *contravariant functor* that takes $\mathrm{Mor}_{\mathcal{C}_1}(X, Y)$ to $\mathrm{Mor}_{\mathcal{C}_2}(F(Y), F(X))$; this can be viewed as a covariant functor from the *opposite category* of $\mathcal{C}_1$ (i.e., the category with the same objects but with the direction of all morphisms switched) to $\mathcal{C}_2$.

Warning:[3] Category theory means never having to say two objects are equal. In practice, two objects which are *isomorphic* (i.e., there are morphisms both ways which compose both ways to the respective identities) are treated as if they were "the same." The catch is, this only works if you keep track of *which* isomorphisms you are using, since an object can be isomorphic to itself in many ways (e.g., permutations of a set). I'll emphasize this point the first couple of times it comes up, but after a little while this will hopefully become a nonissue.

Presheaves and sheaves

Let X be a topological space. Recall that a *presheaf of rings* $\mathcal{F}$ on X consists of the following data:

- for each open subset $U \subseteq X$, a ring $\mathcal{F}(U)$;
- for each inclusion $U \subseteq V$ of open sets, a homomorphism $\mathrm{Res}_{V,U} : \mathcal{F}(U) \to \mathcal{F}(V)$;

such that:[4]

- $\mathrm{Res}_{U,U}$ is always the identity homomorphism;
- for any inclusions $U \subseteq V \subseteq W$, $\mathrm{Res}_{V,U} \circ \mathrm{Res}_{W,V} = \mathrm{Res}_{W,U}$.

[3] Paraphrased from a quote by Ravi Vakil. If you don't understand the allusion, ask Wikipedia about "Love Story."

[4] Some authors also require $\mathcal{F}(\emptyset) = 0$, but this is not essential. See below.

I hope it is now obvious how to restate this definition in the language of categories. Namely, associate to X a category $\underline{X}$ whose objects are the open subsets of X, and whose morphisms are

$$\mathrm{Mor}_{\underline{X}}(U,V) = \begin{cases} \{*\} & V \subseteq U \\ \emptyset & V \not\subseteq U. \end{cases}$$

(Note that there is absolutely no ambiguity about how to define composition.) Then a presheaf of rings is nothing but a contravariant functor from $\underline{X}$ to **Ring**. With this in mind, it is obvious that we can change the target functor to define *presheaves of sets*, *presheaves of abelian groups*, etc.

The standard examples of presheaves of sets are the ones given by fixing another topological space Y (e.g., the real numbers) and taking $\mathcal{F}(U)$ to be the set of continuous maps $U \to Y$. Such presheaves have the property that the elements of a given $\mathcal{F}(U)$ (which we call *sections of* $\mathcal{F}$ *over* U) can be constructed locally, i.e., by specifying sections on an open cover of U which agree on overlaps. We accordingly define a *sheaf of sets* to be a presheaf $\mathcal{F}$ with the property that for every open subset $U \subseteq X$ and every covering of U by open subsets $\{U_i\}_{i \in I}$, the map

$$\mathcal{F}(U) \to \prod_{i \in I} \mathcal{F}(U_i), \qquad s \mapsto (\mathrm{Res}_{U_i,U}(s))_{i \in I}$$

defines a bijection of $\mathcal{F}(U)$ with the set

$$\{(s_i)_{i \in I} \in \prod_{i \in I} \mathcal{F}(U_i) : \mathrm{Res}_{U_i \cap U_j, U_i}(s_i) = \mathrm{Res}_{U_i \cap U_j, U_j}(s_j) \text{ for all } i,j \in I\}.$$

We may similarly define a *sheaf of abelian groups* or a *sheaf of rings*.

By the way, what is $\mathcal{F}(\emptyset)$? Well, the empty set is a subset of X which is covered by the empty covering, i.e., the covering in which the index set is itself empty. So we must ask ourselves: what is an empty product in the category of sets (or abelian groups or rings)? Let's answer this by thinking categorically, just as we did for a product of two objects. The product of a bunch of objects X_i indexed by $i \in I$ should be an object X with maps $X \to X_i$, such that a family of maps from some other object Y to each X_i should arise uniquely from a map $Y \to X$. If I is empty, then X doesn't have any maps out; the universal condition is that any object Y provided with *no extra structure* should map uniquely to X. That is, X should be a *final object* in the category, and for sets any one-element set has that property. (Likewise for abelian groups, the one-element group; and for rings, the zero ring. Note: to make this work, we have to agree that the zero ring is a ring even though $0 = 1$.)

While we will mostly be interested in sheaves, it sometimes will happen that a natural example of a presheaf does not in fact form a sheaf. There is a natural way to fix this using an operation called *sheafification*; we will come back to this later.

Stalks

In topology, geometry, and analysis, one often has reason to talk about the *germ* of a function at a point, i.e., one is interested in the function on some unspecified neighborhood of the point, and two functions that look the same in a neighborhood of the point should be considered equal even if they differ someplace far away. This is captured in the theory of sheaves via the definition of stalks.

Let $\mathcal{F}$ be a sheaf of sets on a topological space X. For each $x \in X$, we define the stalk $\mathcal{F}_x$ to be the *direct limit* of the sets $\mathcal{F}(U)$ as U varies among all of the open subsets of X containing x. In other words, take the disjoint union of all of the $\mathcal{F}(U)$, then form the equivalence relation in which $s_1 \in \mathcal{F}(U_1)$ equals $s_2 \in \mathcal{F}(U_2)$ if there exists some open $V \subseteq U_1 \cap U_2$ such that $\mathrm{Res}_{V,U_1}(s_1) = \mathrm{Res}_{V,U_2}(s_2)$.

Note that the stalk carries much more information than the value of a function at a point. For instance, if $\mathcal{F}$ consists of continuous functions to some Y, then two elements of the stalk $\mathcal{F}_x$ are equal only if they are defined by continuous functions which agree not only at x itself, but on a whole open neighborhood around x.

One may similarly define the stalk of $\mathcal{F}$ at any set $W \subseteq X$, by taking the direct limit of $\mathcal{F}(U)$ as U varies among all of the open subsets of X containing W. Two extreme cases: if $W = \{x\}$ this is the stalk at a point we just defined; if W is open, then this just gives $\mathcal{F}(W)$ itself.