FOUNDATIONS OF ALGEBRAIC GEOMETRY CLASS 2

RAVI VAKIL

CONTENTS

Last day: What is algebraic geometry? Crash course in category theory, ending with Yoneda's Lemma.

Today: more examples of things defined using universal properties: inverse limits, direct limits, adjoint functors, groupification. Sheaves: the motivating example of differentiable functions. Definition of presheaves and sheaves.

At the start of the class, everyone filled out a sign-up sheet, giving their name, e-mail address, mathematical interests, and odds of attending. You're certainly not committing yourself to anything by signing it; it will just give me a good sense of who is in the class. Also, I'll use this for an e-mail list. If you didn't fill out the sheet, but want to, e-mail me.

I'll give out some homework problems on Monday, that will be based on this week's lectures. Most or all of the problems will already have been asked in class. Most likely I will give something like ten problems, and ask you to do five of them.

If you have any questions, please ask me, both in and out of class. Other people you can ask include the other algebraic geometers in the class, including Rob Easton, Andy Schultz, Jarod Alper, Joe Rabinoff, Nikola Penev, and others.

1. Some constructions using universal properties, old and new

Last day, I defined categories and functors. I hope I convinced you that you already have a feeling for categories, because you know so many examples.

A key point was **Yoneda's Lemma**, which says informally that you can essentially recover an object in a category by knowing the maps into it. For example, the data of maps to $X \times Y$ are precisely the data of maps to X and to Y. The data of maps to $X \times_Z Y$ are

Date: Wednesday, September 28, 2005. Last edited January 31, 2007.

precisely the data of maps to X and Y that commute with maps to Z. (Explain how the universal property says this.)

1.1. Example: Inverse limits. Here is another example of something defined by universal properties. Suppose you have a sequence

$$\cdots \longrightarrow A_3 \longrightarrow A_2 \longrightarrow A_1$$

of morphisms in your category. Then the **inverse limit** is an object $\lim_{\leftarrow} A_i$ along with commuting morphisms to all the A_i

$$\begin{tikzcd}
 & \lim_{\leftarrow} A_i \arrow[d] \arrow[dr] \arrow[drr] & & \\
\cdots \arrow[r] & A_3 \arrow[r] & A_2 \arrow[r] & A_1
\end{tikzcd}$$

so that any other object along with maps to the A_i factors through $\lim_{\leftarrow} A_i$

$$\begin{tikzcd}
W \arrow[r, dotted, "\exists!"] \arrow[dr] \arrow[drr] \arrow[drrr] & \lim_{\leftarrow} A_i \arrow[d] \arrow[dr] \arrow[drr] & & \\
\cdots \arrow[r] & A_3 \arrow[r] & A_2 \arrow[r] & A_1
\end{tikzcd}$$

You've likely seen such a thing before. How many of you have seen the p-adics ($\mathbb{Z}_p = ? + ?p + ?p^2 + ?p^3 + \cdots$)? Here's an example in the categories of rings.

$$\begin{tikzcd}
 & \mathbb{Z}_p \arrow[d] \arrow[dr] \arrow[drr] & & \\
\cdots \arrow[r] & \mathbb{Z}/p^3 \arrow[r] & \mathbb{Z}/p^2 \arrow[r] & \mathbb{Z}/p
\end{tikzcd}$$

You can check using your universal property experience that if it exists, then it is unique, up to unique isomorphism. It will boil down to the following fact: we know precisely what the maps to $\lim_{\leftarrow} A_i$ are: they are the same as maps to all the A_i's.

A few quick comments.

(1) We don't know in general that they have to exist.

(2) Often you can see that it exists. If these objects in your categories are all sets, as they are in this case of $\mathbb{Z}_p$, you can interpret the elements of the inverse limit as an element of $a_i \in A_i$ for each i, satisfying $f(a_i) = a_{i-1}$. From this point of view, $2 + 3p + 2p^2 + \cdots$ should be understood as the sequence $(2, 2 + 3p, 2 + 3p + 2p^2, \ldots)$.

(3) We could generalize the system in any different ways. We could basically replace it with any category, although this is way too general. (Most often this will be a *partially ordered set*, often called *poset* for short.) If you wanted to say it in a fancy way, you could say the system could be indexed by an arbitrary category. Example: the product is an example of an inverse limit. Another example: the fibered product is an example of an inverse limit. (What are the partially ordered sets in each case?) Infinite products, or indeed products in general, are examples of inverse limits.

1.2. Example: Direct limits. More immediately relevant for us will be the dual of this notion. We just flip all the arrows, and get the notion of a *direct limit*. Again, if it exists, it is unique up to unique isomorphism.

Here is an example. $5^{-\infty}\mathbb{Z} = \lim_{\to} 5^{-i}\mathbb{Z}$ is an example. (These are the rational numbers whose denominators are required to be powers of 5.)

$$\begin{tikzcd} \mathbb{Z} \arrow[r] \arrow[d] & 5^{-1}\mathbb{Z} \arrow[r] \arrow[dl] & 5^{-2}\mathbb{Z} \arrow[r] \arrow[dll] & \cdots \\ 5^{-\infty}\mathbb{Z} & & & \end{tikzcd}$$

Even though we have just flipped the arrows, somehow it behaves quite differently from the inverse limit.

Some observations:

(1) In this example, each element of the direct limit is an element of something upstairs, but you can't say in advance what it is an element of. For example, 17/125 is an element of the $5^{-3}\mathbb{Z}$ (or $5^{-4}\mathbb{Z}$, or later ones), but not $5^{-2}\mathbb{Z}$.

(2) We can index this by any partially ordered set (or *poset*). (Or even any category, although I don't know if we care about this generality.)

1.3. Remark. (3) That first remark applies in some generality for the category of A-modules, where A is a ring. (See Atiyah-Macdonald p. 32, Exercise 14.) We say a partially ordered set I is a *directed set* if for $i, j \in I$, there is some $k \in I$ with $i, j \leq k$. We can show that the direct limit of any system of R-modules indexed by I exists, by constructing it. Say the system is given by M_i ($i \in I$), and $f_{ij} : M_i \to M_j$ ($i \leq j$ in I). Let $M = \oplus_i M_i$, where each M_i is associated with its image in M, and let R be the submodule generated by all elements of the form $m_i - f_{ij}(m_i)$ where $m_i \in M_i$ and $i \leq j$. **Exercise.** Show that M/R (with the inclusion maps from the M_i) is $\lim_{\to} M_i$. (This example will come up soon.) You will notice that the same argument works in other interesting categories, such as: sets; groups; and abelian groups. (Less important question for the experts: what hypotheses do we need for this to work more generally?)

(4) (Infinite) sums are examples of direct limits.

2. Adjoint functors

Let me re-define adjoint functors (Weibel Definition 2.3.9). Two *covariant* functors $L : \mathcal{A} \to \mathcal{B}$ and $R : \mathcal{B} \to \mathcal{A}$ are *adjoint* if there is a natural bijection for all $A \in \mathcal{A}$ and $B \in \mathcal{B}$

$$\tau_{AB} : \mathrm{Hom}_{\mathcal{B}}(L(A), B) \to \mathrm{Hom}_{\mathcal{A}}(A, R(B)).$$

In this instance, let me make precise what "natural" means, which will also let us see why the functors here are covariant. For all $f : A \to A'$ in $\mathcal{A}$, we require

$$\begin{tikzcd} \mathrm{Hom}_{\mathcal{B}}(L(A'), B) \arrow[r, "Lf^*"] \arrow[d, "\tau"] & \mathrm{Hom}_{\mathcal{B}}(L(A), B) \arrow[d, "\tau"] \\ \mathrm{Hom}_{\mathcal{A}}(A', R(B)) \arrow[r, "f^*"] & \mathrm{Hom}_{\mathcal{A}}(A, R(B)) \end{tikzcd}$$

to commute, and for all $g : B \to B'$ in $\mathcal{B}$ we want a similar commutative diagram to commute. (Here f^* is the map induced by $f : A \to A'$, and Lf^* is the map induced by $Lf : L(A) \to L(A')$.)

Exercise. Write down what this diagram should be.

We could figure out what this should mean if the functors were both contravariant. I haven't tried to see if this could make sense.

You've actually seen this before, in linear algebra, when you have seen adjoint matrices. But I've long forgotten how they work, so let me show you another example. (Question for the audience: is there a very nice example out there?)

2.1. Example: groupification. Motivating example: getting a group from a semigroup. A semigroup is just like a group, except you don't require an inverse. Examples: the non-negative integers $0, 1, 2, \ldots$ under addition, or the positive integers under multiplication $1, 2, \ldots$. From a semigroup, you can create a group, and this could be called groupification. Here is a formalization of that notion. If S is a semigroup, then its groupification is a map of semigroups $\pi : S \to G$ such that G is a group, and any other map of semigroups from S to a *group* G' factors *uniquely* through G.

$$\begin{tikzcd} S \arrow[r] \arrow[dr, "\pi"] & G \arrow[d, dashed, "\exists!"] \\ & G' \end{tikzcd}$$

(Thanks Jack for explaining how to make dashed arrows in `\xymatrix`.)

(General idea for experts: We have a full subcategory of a category. We want to "project" from the category to the subcategory. We have $\mathrm{Hom}_{\text{category}}(S, H) = \mathrm{Hom}_{\text{subcategory}}(G, H)$ automatically; thus we are describing the left adjoint to the forgetful functor. How the argument worked: we constructed something which was in the small category, which automatically satisfies the universal property.)

Example of a universal property argument: If a semigroup is *already* a group then groupification is the identity morphism, by the universal property.

Exercise to get practice with this. Suppose R is a ring, and S is a multiplicative subset. Then $S^{-1}R$-modules are a full subcategory of the category of R-modules. Show that $M \to$

$S^{-1}M$ satisfies a universal property. Translation: Figure out what the universal property is.

2.2. Additive and abelian categories. There is one last concept that we will use later. It is convenient to give a name to categories with some additional structure. Here are some definitions.

Initial object of a category. It is an object with a unique map to any other object. (By a universal property argument, if it exists, it is unique up to unique isomorphism.) Example: the empty set, in the category of sets.

Final object of a category. It is an object with a unique map from any other object. (By a universal property argument, if it exists, it is unique up to unique isomorphism.) Question: does the category of sets have a final object?

Exercise. If Z is the final object in a category $\mathcal{C}$, and $X, Y \in \mathcal{C}$, then "$X \times_Z Y = X \times Y$" ("the" fibered product over Z is canonically isomorphic to "the" product). (This is an exercise about unwinding the definition.)

Additive categories (Weibel, p. 5). (I think I forgot to say part of this definition in class.) A category $\mathcal{C}$ is said to be *additive* if it has the following properties. For each $A, B \in \mathcal{C}$, $\mathrm{Hom}_{\mathcal{C}}(A, B)$ is an abelian group, such that composition of morphisms distributes over addition (think about what this could mean). It has a 0-object (= simultaneously initial object and final object), and products (a product $A \times B$ for any pair of objects). (Why is the 0-object called the 0-object?)

Yiannis points out that Banach spaces form an additive category. Another example are R-modules for a ring R, but they have even more structure.

Abelian categories (Weibel, p. 6). I deliberately didn't give a precise definition in class, as you should first get used to this concept before reading the technical definition.

But here it is. Let $\mathcal{C}$ be an additive category. First, a *kernel* of a morphism $f : B \to C$ is a map $i : A \to B$ such that $f \circ i = 0$, and that is universal with respect to this property. (Hence it is unique up to unique isomorphism by universal property nonsense. Note that we said "a" kernel, not "the" kernel.) A *cokernel* is defined dually by reversing the arrows — do this yourself. We say a morphism i in $\mathcal{C}$ is *monic* if $i \circ g = 0$, where the source of g is the target of i, implies $g = 0$. Dually, there is the notion of *epi* — reverse the arrows to find out what that is.

An *abelian category* is an additive category satisfying three properties. 1. Every map has a kernel and cokernel. 2. Every monic is the kernel of its cokernel. 3. Every epi is the cokernel of its kernel.

It is a non-obvious and imprecise fact that every property you want to be true about kernels, cokernels, etc. follows from these three.

An abelian category has kernels and cokernels and images, and they behave the way you expect them to. So you can have exact sequences: we say

$$A \xrightarrow{f} B \xrightarrow{g} C$$

is exact if $\ker g = \operatorname{im} f$.

The key example of an abelian category is the category of R-modules (where R is a ring).

3. Sheaves

I now want to discuss an important new concept, the notion of a *sheaf*. A sheaf is the kind of object you will automatically consider if you are interested in something like the continuous functions on a space X, or the differentiable functions on a space X, or things like that. Basically, you want to consider all continuous functions on all open sets all at once, and see what properties this sort of collection of information has. I'm going to motivate it for you, and tell you the definition. I find this part quite intuitive. Then I will do things with this concept (for example talking about cokernels of maps of sheaves), and things become less intuitive.

3.1. Motivating example: sheaf of differentiable functions. We'll consider differentiable functions on $X = \mathbb{R}^n$, or a more general manifold X. To each open set $U \subset X$, we have a ring of differentiable functions. I will denote this ring $\mathcal{O}(U)$.

If you take a differentiable function on an open set, you can restrict it to a smaller open set, and you'll get a differentiable function there. In other words, if $U \subset V$ is an inclusion of open sets, we have a map $\mathrm{res}_{V,U}; \mathcal{O}(V) \to \mathcal{O}(U)$.

If you take a differentiable function on a big open set, and restrict it to a medium open set, and then restrict that to a small open set, then you get the same thing as if you restrict the differentiable function on the big open set to the small open set all at once. In other words, if $U \hookrightarrow V \hookrightarrow W$, then the following diagram commutes:

$$\begin{tikzcd}
\mathcal{O}(W) \arrow[rr, "\mathrm{res}_{W,V}"] \arrow[dr, "\mathrm{res}_{W,U}"'] & & \mathcal{O}(V) \arrow[dl, "\mathrm{res}_{V,U}"] \\
& \mathcal{O}(U) &
\end{tikzcd}$$

Now say you have two differentiable functions f_1 and f_2 on a big open set U, and you have an open cover of U by some U_i. Suppose that f_1 and f_2 agree on each of these U_i. Then they must have been the same function to begin with. Right? In other words, if $\{U_i\}_{i \in I}$ is a cover of U, and $f_1, f_2 \in \mathcal{O}(U)$, and $\mathrm{res}_{U,U_i} f_1 = \mathrm{res}_{U,U_i} f_2$, then $f_1 = f_2$. In other words, I can *identify* functions locally.

Finally, suppose I still have my U, and my cover U_i of U. Suppose I've got a differentiable function on each of the U_i — a function on U_1, a function on U_2, and so on — and they agree on the overlaps. Then I can glue all of them together to make one function on all of U. Right? In other words: given $f_i \in \mathcal{O}(U_i)$ for all i, such that $\mathrm{res}_{U_i, U_i \cap U_j} f_i = \mathrm{res}_{U_j, U_i \cap U_j} f_j$ for all i, j, then there is some $f \in \mathcal{O}(U)$ such that $\mathrm{res}_{U, U_i} f = f_i$ for all i.

Great. Now I could have done all this with continuous functions. [Go back over it all, with differentiable replaced by continuous.] Or smooth functions. Or just functions. That's the idea that we'll formalize soon into a sheaf.

3.2. Motivating example continued: the germ of a differentiable function. Before we do, I want to point out another definition, that of the germ of a differentiable function at a point $x \in X$. Intuitively, it is a shred of a differentiable function at x. Germs are objects of the form $\{(f, \text{open } U) : x \in U, f \in \mathcal{O}(U)\}$ modulo the relation that $(f, U) \sim (g, V)$ if there is some open set $W \subset U, V$ where $f|_W = g|_W$ (or in our earlier language, $\mathrm{res}_{U,W} f = \mathrm{res}_{V,W} g$). In other words, two functions that are the same here near x but differ way over there have the same germ. Let me call this set of germs $\mathcal{O}_x$. Notice that this forms a ring: you can add two germs, and get another germ: if you have a function f defined on U, and a function g defined on V, then $f + g$ is defined on $U \cap V$. Notice also that if $x \in U$, you get a map

$$\mathcal{O}(U) \to \{\text{germs at } x\}.$$

Aside for the experts: this is another example of a direct limit, and I'll tell you why in a bit.

Fact: $\mathcal{O}_x$ is a local ring. Reason: Consider those germs vanishing at x. That certainly is an ideal: it is closed under addition, and when you multiply something vanishing at x by any other function, you'll get something else vanishing at x. Anything not in this ideal is invertible: given a germ of a function f not vanishing at x, then f is non-zero near x by continuity, so $1/f$ is defined near x. The residue map should map onto a field, and in this case it does: we have an exact sequence:

$$0 \longrightarrow \mathfrak{m} := \text{ideal of germs vanishing at } x \longrightarrow \mathcal{O}_x \xrightarrow{f \mapsto f(x)} \mathbb{R} \longrightarrow 0$$

If you have never seen exact sequences before, this is a good chance to figure out how they work. This is what is called a short exact sequence. **Exercise.** Check that this is an exact sequence, i.e. that the image of each map is the kernel of the next. Show that this implies that the map on the left is an injection, and the one on the right is a surjection.

(Interesting fact, for people with a little experience with a little geometry: $\mathfrak{m}/\mathfrak{m}^2$ is a module over $\mathcal{O}_x/\mathfrak{m} \cong \mathbb{R}$, i.e. it is a real vector space. It turns out to be "naturally" — whatever that means — the cotangent space to the manifold at x. This will turn out to be handy later on, when we define tangent and cotangent spaces of schemes.)

Conclusion: We can interpret the value of a function at a point, or the value of a germ at a point, as an element of the local ring modulo the maximal ideal. (However, this can be a bit more problematic for more general sheaves.)

3.3. Definition of sheaf and presheaf.

We are now ready to formalize these notions.

Definition: Sheaf on a topological space X**.** (A note on language: this is called a sheaf because of an earlier, different perspective on the definition, see Serre's *Faisceaux Algébriques Cohérents*. I'm not going to discuss this earlier definition, so you'll have to take this word without any motivation.)

I will define a sheaf of sets , just to be concrete. But you can have sheaves of groups, rings, modules, etc. without changing the definitions at all. Indeed, if you want to be fancy, you can say that you can have a sheaf with values in any category.

A *presheaf* $\mathcal{F}$ is the following data. To each open set $U \subset X$, we have a set $\mathcal{F}(U)$ (e.g. the set of differentiable functions). (Notational warning: Several notations are in use, for various good reasons: $\mathcal{F}(U) = \Gamma(U, \mathcal{F}) = H^0(U, \mathcal{F})$. I will use them all. I forgot to say this in class, but will say it next day.) The elements of $\mathcal{F}(U)$ are called *sections of* $\mathcal{F}$ *over* U.

For each inclusion $U \hookrightarrow V$, we have a restriction map $\mathrm{res}_{V,U} : \mathcal{F}(V) \to \mathcal{F}(U)$ (just as we did for differentiable functions). The map $\mathrm{res}_{U,U}$ should be the identity. If $U \hookrightarrow V \hookrightarrow W$, then the restriction maps commute, i.e. the following diagram commutes.

$$\begin{tikzcd} \mathcal{F}(W) \arrow[rr, "\mathrm{res}_{W,V}"] \arrow[dr, "\mathrm{res}_{W,U}"'] & & \mathcal{F}(V) \arrow[dl, "\mathrm{res}_{V,U}"] \\ & \mathcal{F}(U) & \end{tikzcd}$$

That ends the definition of a presheaf.

3.4. Useful exercise for experts liking category theory: "A presheaf is the same as a contravariant functor". Given any topological space X, we can get a category, which I will call the "category of open sets". The objects are the open sets. The morphisms are the inclusions $U \hookrightarrow V$. (What is the initial object? What is the final object?) Verify that the data of a presheaf is precisely the data of a contravariant functor from the category of open sets of X to the category of sets.

Note for pedants, which can be ignored by everyone else. An annoying question is: what is $\mathcal{F}(\emptyset)$. We will see that it can be convenient to have $\mathcal{F}(\emptyset) = \{1\}$, or more generally, if we are having sheaves with value in some category $\mathcal{C}$ (such as Groups), we would like $\mathcal{F}(\emptyset)$ to be the final object in the category. This should probably be part of the definition of presheaf. (For example, Weibel, p. 26, takes it as such; Weibel seems to define sheaves with values only in an abelian category.) I hope to be fairly scrupulous in this course, so I hope people who care keep me honest on issues like this.

We add two more axioms to make this into a sheaf.

Identity axiom. If $\{U_i\}_{i \in I}$ is a open cover of U, and $f_1, f_2 \in \mathcal{F}(U)$, and $\mathrm{res}_{U,U_i} f_1 = \mathrm{res}_{U,U_i} f_2$, then $f_1 = f_2$.

(A presheaf + identity axiom is sometimes called a *separated sheaf*, but we will not use that notation here.)

Gluability axiom. If $\{U_i\}_{i \in I}$ is a open cover of U, then given $f_i \in \mathcal{F}(U_i)$ for all i, such that $\mathrm{res}_{U_i, U_i \cap U_j} f_i = \mathrm{res}_{U_j, U_i \cap U_j} f_j$ for all i, j, then there is some $f \in \mathcal{F}(U)$ such that $\mathrm{res}_{U, U_i} f = f_i$ for all i.

(Philosophical note: identity means there is at most one way to glue. Gluability means that there is at least one way to glue.)

Remark for people enjoying category theory for the first time — as opposed to learning it for the first time. The gluability axiom may be interpreted as saying that $\mathcal{F}(\cup_{i \in I} U_i)$ is a certain inverse limit.

Example. If U and V are disjoint, then $\mathcal{F}(U \cup V) = \mathcal{F}(U) \times \mathcal{F}(V)$. (Here we use the fact that $F(\emptyset)$ is the final object, from the "note for pedants" above.)

3.5. Exercise. Suppose Y is a topological space. Show that "continuous maps to Y" form a sheaf of sets on X. More precisely, to each open set U of X, we associate the set of continuous maps to Y. Show that this forms a sheaf.
(Fancier versions that you can try:
(b) Suppose we are given a continuous map $f : Y \to X$. Show that "sections of f" form a sheaf. More precisely, to each open set U of X, associate the set of continuous maps s to Y such that $f \circ s = id|_U$. Show that this forms a sheaf.
(c) (If you know what a topological group is.) Suppose that Y is a topological group. Show that maps to Y form a sheaf of *groups*. (If you don't know what a topological group is, you might be able to guess.)

Example: skyscraper sheaf. Suppose X is a topological space, with $x \in X$, and G is a group. Then $\mathcal{F}$ defined by $\mathcal{F}(U) = G$ if $x \in U$ and $\mathcal{F}(U) = \{e\}$ if $x \notin U$ forms a sheaf. (Check this if you don't see how.) This is called a *skyscraper sheaf*, because the informal picture of it looks like a skyscraper at x.

Important example/exercise: the pushforward . Suppose $f : X \to Y$ is a continuous map, and $\mathcal{F}$ is a sheaf on X. Then define $f_*\mathcal{F}$ by $f_*\mathcal{F}(V) = \mathcal{F}(f^{-1}(V))$, where V is an open subset of Y. Show that $f_*\mathcal{F}$ is a sheaf. This is called a *pushforward sheaf*. More precisely, $f_*\mathcal{F}$ is called the *pushforward of $\mathcal{F}$ by* f.

Example / exercise. (a) Let X be a topological space, and S a set with more than one element, and define $\mathcal{F}(U) = S$ for all open sets U. Show that this forms a presheaf (with the obvious restriction maps), and even satisfies the identity axiom. Show that this needn't form a sheaf. (Actually, for this to work, here we need $\mathcal{F}(\emptyset)$ to be the final object, not S. Without this patch, the constant presheaf *is* a sheaf. You can already see how the empty set is giving me a headache.) This is called the *constant presheaf with values in* S. We will denote this presheaf by $\underline{S}^{\mathrm{pre}}$.
(b) Now let $\mathcal{F}(U)$ be the maps to S that are *locally constant*, i.e. for any point x in U, there is a neighborhood of x where the function is constant. A better description is this: endow S with the discrete topology, and let $\mathcal{F}(U)$ be the continuous maps $U \to S$. Show that this is a *sheaf*. (Here we need $\mathcal{F}(\emptyset)$ to be the final object again, *not* S.) Using the "better description", this follows immediately from Exercise 3.5. We will try to call this the *locally*

constant sheaf. (Unfortunately, in the real world, this is stupidly called the *constant sheaf*.) We will denote this sheaf by $\underline{S}$.

3.6. Stalks. We define **stalk** = set of germs of a (pre)sheaf $\mathcal{F}$ in just the same way as before: Elements are $\{(f, \text{open } U) : x \in U, f \in \mathcal{O}(U)\}$ modulo the relation that $(f, U) \sim (g, V)$ if there is some open set $W \subset U, V$ where $\mathrm{res}_{U,W} f = \mathrm{res}_{V,W} g$. In other words, two section that are the same near x but differ far away have the same germ. This set of germs is denoted $\mathcal{F}_x$.

A useful equivalent definition is as a direct limit, of all $\mathcal{F}(U)$ where $x \in U$:

$$\mathcal{F}_x := \{\text{germs at } x\} = \varinjlim \mathcal{F}(U).$$

(All such U into a partially ordered set using inclusion. People having thought about the category of open sets, §3.4, will have a warm feeling in their stomachs.) This poset is a directed set (§1.3: given any two such sets, there is a third such set contained in both), so these two definitions are the same by Remark/Exercise 1.3. It would be good for you to think this through. Hence by that Remark/Exercise, we can have stalks for sheaves of sets, groups, rings, and other things for which direct limits exist for directed sets.

Let me repeat: it is useful to think of stalks in both ways, as direct limits, and also as something extremely explicit: an element of a stalk at p has as a representative a section over an open set near p.

Caution: Value at a point doesn't yet make sense.

3.7. Exercise. Show that pushforward induces maps of stalks.

E-mail address: `vakil@math.stanford.edu`