FOUNDATIONS OF ALGEBRAIC GEOMETRY CLASS 3

RAVI VAKIL

CONTENTS

1. Kernels, cokernels, and exact sequences: A brief introduction to abelian categories 1
2. Sheaves 7
3. Motivating example: The sheaf of differentiable functions. 7
4. Definition of sheaf and presheaf 9

Last day: category theory in earnest. Universal properties. Limits and colimits. Adjoints.

Today: abelian categories: kernels, cokernels, and all that jazz.

Here are some additional comments on last day's material. The details of Yoneda's lemma don't matter so much; what matters most is that you understand how universal properties determine objects up to unique isomorphism.

It doesn't matter much, but limits and colimits needn't be indexed only by categories where there is at most one morphism between any two objects. I gave an example involving a G-action on a set X (where G is a finite group). The G-invariants can be interpreted as limit.

Tony Licata gave a nice argument that $\otimes$ is right-exact using a universal property argument.

1. KERNELS, COKERNELS, AND EXACT SEQUENCES: A BRIEF INTRODUCTION TO ABELIAN CATEGORIES

Since learning linear algebra, you have been familiar with the notions and behaviors of kernels, cokernels, etc. Later in your life you saw them in the category of abelian groups, and later still in the category of A-modules. Each of these notions generalizes the previous one. The notion of abelian category formalizes kernels etc.

Date: Monday, October 1, 2007. Updated November 4, 2007 to add espace étalé construction.

We now briefly introduce a few notions about abelian categories. We will soon define some new categories (certain sheaves) that will have familiar-looking behavior, reminiscent of that of modules over a ring. The notions of kernels, cokernels, images, and more will make sense, and they will behave "the way we expect" from our experience with modules. This can be made precise through the notion of an abelian category. We will see enough to motivate the definitions that we will see in general: monomorphism (and subobject), epimorphism, kernel, cokernel, and image. But we will avoid having to show that they behave "the way we expect" in a general abelian category because the examples we will see will be directly interpretable in terms of modules over rings.

Abelian categories are the right general setting in which one can do "homological algebra", in which notions of kernel, cokernel, and so on are used, and one can work with complexes and exact sequences.

Two key examples of an abelian category are the category **Ab** of abelian groups, and the category $\mathbf{Mod}_A$ of A-modules. As stated earlier, the first is a special case of the second (just take $A = \mathbb{Z}$). As we give the definitions, you should verify that $\mathbf{Mod}_A$ is an abelian category, and you should keep these examples in mind always.

We first define the notion of *additive category*. We will use it only as a stepping stone to the notion of an abelian category.

1.1. Definition. A category $\mathcal{C}$ is said to be *additive* if it satisfies the following properties.

- Ad1. For each $A, B \in \mathcal{C}$, $\mathrm{Mor}(A, B)$ is an abelian group, such that composition of morphisms distributes over addition. (You should think about what this means — it translates to two distinct statements).
- Ad2. $\mathcal{C}$ has a zero-object, denoted 0. (Recall: this is an object that is simultaneously an initial object and a final object.)
- Ad3. It has products of two objects (a product $A \times B$ for any pair of objects), and hence by induction, products of any finite number of objects.

In an additive category, the morphisms are often called homomorphisms, and Mor is denoted by Hom. In fact, this notation Hom is a good indication that you're working in an additive category. A functor between additive categories preserving the additive structure of Hom, and sending the 0-object to the 0-object, is called an *additive functor*. (It is a consequence of the definition that additive functors send 0-objects to 0-objects, and preserve products.)

1.2. *Remarks.* It is a consequence of the definition of additive category that finite direct products are also finite direct sums=coproducts (the details don't matter to us). The symbol $\oplus$ is used for this notion.

One motivation for the name 0-object is that the 0-morphism in the abelian group $\mathrm{Hom}(A, B)$ is the composition $A \to 0 \to B$.

Real (or complex) Banach spaces are an example of an additive category. The category $\mathbf{Mod}_A$ of A-modules is another example, but it has even more structure, which we now formalize as an example of an abelian category.

1.3. Definition. Let $\mathcal{C}$ be an additive category. 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. Diagramatically:

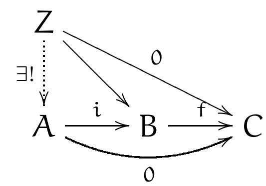

(Note that the kernel is not just an object; it is a morphism of an object to B.) Hence it is unique up to unique isomorphism by universal property nonsense. A *cokernel* is defined dually by reversing the arrows — do this yourself. Notice that the kernel of $f : B \to C$ is the limit

$$\begin{array}{ccc} & & 0 \\ & & \downarrow \\ B & \xrightarrow{f} & C \end{array}$$

and similarly the cokernel is a colimit.

A morphism $i : A \to B$ in $\mathcal{C}$ is *monic* if for all $i \circ g = 0$, where the tail of g is A, implies $g = 0$. Diagramatically,

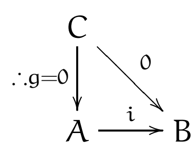

(Once we know what an abelian category is — in a few sentences — you may check that a monic morphism in an abelian category is a monomorphism.) If $i : A \to B$ is monic, then we say that A is a *subobject* of B, where the map i is implicit. Dually, there is the notion of *epi* — reverse the arrows to find out what that is. The notion of *quotient object* is defined dually to subobject.

An *abelian category* is an additive category satisfying three additional properties.

(1) Every map has a kernel and cokernel.
(2) Every monic morphism is the kernel of its cokernel.
(3) Every epi morphism is the cokernel of its kernel.

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

The *image* of a morphism $f : A \to B$ is defined as $\mathrm{im}(f) = \ker(\mathrm{coker}\, f)$. It is the unique factorization

$$A \xrightarrow{\text{epi}} \mathrm{im}(f) \xrightarrow{\text{monic}} B$$

It is the cokernel of the kernel, and the kernel of the cokernel. The reader may want to verify this as an exercise. It is unique up to unique isomorphism.

We will leave the foundations of abelian categories untouched. The key thing to remember is that if you understand kernels, cokernels, images and so on in the category of modules over a ring $\mathbf{Mod}_A$, you can manipulate objects in any abelian category. This is made precise by Freyd-Mitchell Embedding Theorem. However, the abelian categories we'll come across will obviously be related to modules, and our intuition will clearly carry over. For example, we'll show that sheaves of abelian groups on a topological space X form an abelian category. The interpretation in terms of "compatible germs" will connect notions of kernels, cokernels etc. of sheaves of abelian groups to the corresponding notions of abelian groups.

1.4. Complexes, exactness, and homology.

If you aren't familiar with these notions, you should definitely read this section closely!

We say

$$A \xrightarrow{f} B \xrightarrow{g} C \tag{1}$$

is a complex if $g \circ f = 0$, and is exact if $\ker g = \operatorname{im} f$. If (1) is a complex, then its homology is $\ker g / \operatorname{im} f$. We say that $\ker g$ are the *cycles*, and $\operatorname{im} f$ are the *boundaries*. Homology (resp. cohomology) is denoted by H, often with a subscript (resp. superscript), and it should be clear from the context what the subscript means (see for example the discussion below).

An exact sequence

$$A^\bullet : \qquad \cdots \longrightarrow A^{i-1} \xrightarrow{f^{i-1}} A^i \xrightarrow{f^i} A^{i+1} \xrightarrow{f^{i+1}} \cdots \tag{2}$$

can be "factored" into short exact sequences

$$0 \longrightarrow \ker f^i \longrightarrow A^i \longrightarrow \ker f^{i+1} \longrightarrow 0$$

which is helpful in proving facts about long exact sequences by reducing them to facts about short exact sequences.

More generally, if (2) is assumed only to be a complex, then it can be "factored" into short exact sequences

$$0 \longrightarrow \ker f^i \longrightarrow A^i \longrightarrow \operatorname{im} f^i \longrightarrow 0$$

$$0 \longrightarrow \operatorname{im} f^{i-1} \longrightarrow \ker f^i \longrightarrow H^i(A^\bullet) \longrightarrow 0$$

1.A. EXERCISE. Suppose

$$0 \xrightarrow{d^0} A^1 \xrightarrow{d^1} \cdots \xrightarrow{d^{n-1}} A^n \xrightarrow{d^n} 0$$

is a complex of k-vector spaces (often called $A^\bullet$ for short). Show that $\sum(-1)^i \dim A^i = \sum(-1)^i h^i(A^\bullet)$. (Recall that $h^i(A^\bullet) = \dim \ker(d^i)/\operatorname{im}(d^{i-1})$.) In particular, if $A^\bullet$ is exact, then $\sum(-1)^i \dim A^i = 0$. (If you haven't dealt much with cohomology, this will give you some practice.)

1.B. IMPORTANT EXERCISE. Suppose $\mathcal{C}$ is an abelian category. Define the category $\mathbf{Com}_{\mathcal{C}}$ as follows. The objects are infinite complexes

$$A^\bullet : \qquad \cdots \longrightarrow A^{i-1} \xrightarrow{f^{i-1}} A^i \xrightarrow{f^i} A^{i+1} \xrightarrow{f^{i+1}} \cdots$$

in $\mathcal{C}$, and the morphisms $A^\bullet \to B^\bullet$ are commuting diagrams

$$\begin{array}{ccccccccc}
A^\bullet : & \cdots \longrightarrow & A^{i-1} & \xrightarrow{f^{i-1}} & A^i & \xrightarrow{f^i} & A^{i+1} & \xrightarrow{f^{i+1}} & \cdots \\
\downarrow & & \downarrow & & \downarrow & & \downarrow & & \\
B^\bullet : & \cdots \longrightarrow & B^{i-1} & \xrightarrow{f^{i-1}} & B^i & \xrightarrow{f^i} & B^{i+1} & \xrightarrow{f^{i+1}} & \cdots
\end{array}$$

Show that $\mathbf{Com}_{\mathcal{C}}$ is an abelian category. Show that a short exact sequence of complexes

$$\begin{array}{ccccccccc}
0 : & \cdots \longrightarrow & 0 & \longrightarrow & 0 & \longrightarrow & 0 & \longrightarrow & \cdots \\
\downarrow & & \downarrow & & \downarrow & & \downarrow & & \\
A^\bullet : & \cdots \longrightarrow & A^{i-1} & \xrightarrow{f^{i-1}} & A^i & \xrightarrow{f^i} & A^{i+1} & \xrightarrow{f^{i+1}} & \cdots \\
\downarrow & & \downarrow & & \downarrow & & \downarrow & & \\
B^\bullet : & \cdots \longrightarrow & B^{i-1} & \xrightarrow{g^{i-1}} & B^i & \xrightarrow{g^i} & B^{i+1} & \xrightarrow{g^{i+1}} & \cdots \\
\downarrow & & \downarrow & & \downarrow & & \downarrow & & \\
C^\bullet : & \cdots \longrightarrow & C^{i-1} & \xrightarrow{h^{i-1}} & C^i & \xrightarrow{h^i} & C^{i+1} & \xrightarrow{h^{i+1}} & \cdots \\
\downarrow & & \downarrow & & \downarrow & & \downarrow & & \\
0 : & \cdots \longrightarrow & 0 & \longrightarrow & 0 & \longrightarrow & 0 & \longrightarrow & \cdots
\end{array}$$

induces a *long exact sequence in cohomology*

$$\cdots \longrightarrow H^{i-1}(C^\bullet) \longrightarrow$$

$$H^i(A^\bullet) \longrightarrow H^i(B^\bullet) \longrightarrow H^i(C^\bullet) \longrightarrow$$

$$H^{i+1}(A^\bullet) \longrightarrow \cdots$$

1.5. *Exactness of functors.* If $F : \mathcal{A} \to \mathcal{B}$ is a covariant additive functor from one abelian category to another, we say that F is *right-exact* if the exactness of

$$A' \longrightarrow A \longrightarrow A'' \longrightarrow 0,$$

in $\mathcal{A}$ implies that

$$F(A') \longrightarrow F(A) \longrightarrow F(A'') \longrightarrow 0$$

is also exact. Dually, we say that F is *left-exact* if the exactness of

$$0 \longrightarrow A' \longrightarrow A \longrightarrow A'' \qquad \text{implies}$$

$$0 \longrightarrow F(A') \longrightarrow F(A) \longrightarrow F(A'') \qquad \text{is exact.}$$

A contravariant functor is *left-exact* if the exactness of

$$A' \longrightarrow A \longrightarrow A'' \longrightarrow 0 \qquad \text{implies}$$

$$0 \longrightarrow F(A'') \longrightarrow F(A) \longrightarrow F(A') \qquad \text{is exact.}$$

The reader should be able to deduce what it means for a contravariant functor to be *right-exact*.

A covariant or contravariant functor is *exact* if it is both left-exact and right-exact.

1.6. $\star$ **Interactions of adjoints, (co)limits, and (left and right) exactness.** There are some useful properties of adjoints that make certain arguments quite short. This is intended only for experts, and can be ignored by most people in the class, so this won't be said during class. We present them as three facts. Suppose $(F : \mathcal{C} \to \mathcal{D}, G : \mathcal{D} \to \mathcal{C})$ is a pair of adjoint functors.

Fact 1. F commutes with colimits, and G commutes with limits.

We prove the second statement here. The first is the same, "with the arrows reversed". We begin with a useful fact.

1.C. EXERCISE: $\operatorname{Mor}(X, \cdot)$ COMMUTES WITH LIMITS. Suppose A_i ($i \in \mathcal{I}$) is a diagram in $\mathcal{D}$ indexed by $\mathcal{I}$, and $\varprojlim A_i \to A_i$ is its limit. Then for any $X \in \mathcal{D}$, $\operatorname{Mor}(X, \varprojlim A_i) \to \operatorname{Mor}(X, A_i)$ is the limit $\varprojlim \operatorname{Mor}(X, A_i)$.

We are now ready to prove (one direction of) Fact 1.

1.7. *Proposition (right-adjoints commute with limits). — Suppose* $(F : \mathcal{C} \to \mathcal{D}, G : \mathcal{D} \to \mathcal{C})$ *is a pair of adjoint functors. If* $A = \varprojlim A_i$ *is a limit in* $\mathcal{D}$ *of a diagram indexed by* I*, then* $GA = \varprojlim GA_i$ *(with the corresponding maps* $GA \to GA_i$*) is a limit in* $\mathcal{C}$*.*

Proof. We must show that $GA \to GA_i$ satisfies the universal property of limits. Suppose we have maps $W \to GA_i$ commuting with the maps of $\mathcal{I}$. We wish to show that there exists a unique $W \to GA$ extending the $W \to GA_i$. By adjointness of F and G, we can restate this as: Suppose we have maps $FW \to A_i$ commuting with the maps of $\mathcal{I}$. We wish to show that there exists a unique $FW \to A$ extending the $FW \to A_i$. But this is precisely the universal property of the limit. □

Suppose now further that $\mathcal{C}$ and $\mathcal{D}$ are abelian categories, and F and G are additive functors. Kernels are limits and cokernels are colimits (§1.3), so we have **Fact 2.** F commutes with cokernels and G commutes with kernels.

Now suppose

$$M' \xrightarrow{f} M \longrightarrow M'' \longrightarrow 0$$

is an exact sequence in $\mathcal{C}$, so $M'' = \operatorname{coker} f$. Then by Fact 2, $FM'' = \operatorname{coker} Ff$. Thus

$$FM' \to FM \to FM'' \to 0$$

so: **Fact 3.** Left-adjoint additive functors are right-exact, and right-adjoint additive functors are left-exact. For example, the fact that $(\cdot \otimes_A N, \operatorname{Hom}_A(N, \cdot))$ are an adjoint pair (from the A-**Mod** to itself) imply that $\cdot \otimes_A N$ is right-exact (an exercise from last week) and $\operatorname{Hom}(N, \cdot)$ is left-exact.

2. Sheaves

It is perhaps suprising that geometric spaces are often best understood in terms of (nice) functions on them. For example, a differentiable manifold that is a subset of $\mathbb{R}^n$ can be studied in terms of its differentiable functions. Because geometric spaces can have few functions, a more precise version of this insight is that the structure of the space can be well understood by undestanding all functions on all open subsets of the space. This information is encoded in something called a *sheaf*. We will define *sheaves* and describe many useful facts about them. Sheaves were introduced by Leray in the 1940s. The reason for the name is from an earlier, different perspective on the definition, which we shall not discuss.

We will begin with a motivating example to convince you that the notion is not so foreign.

One reason sheaves are often considered slippery to work with is that they keep track of a huge amount of information, and there are some subtle local-to-global issues. There are also three different ways of getting a hold of them.

- in terms of open sets (the definition §4) — intuitive but in some way the least helpful
- in terms of stalks
- in terms of a base of a topology.

Knowing which idea to use requires experience, so it is essential to do a number of exercises on different aspects of sheaves in order to truly understand the concept.

3. Motivating example: The sheaf of differentiable functions.

We will consider differentiable functions on the topological $X = \mathbb{R}^n$, although you may consider a more general manifold X. The sheaf of differentiable functions on X is the data of all differentiable functions on all open subsets on X; we will see how to manage this data, and observe some of its properties. To each open set $U \subset X$, we have a ring of differentiable functions. We denote this ring $\mathcal{O}(U)$.

Given a differentiable function on an open set, you can restrict it to a smaller open set, obtaining a differentiable function there. In other words, if $U \subset V$ is an inclusion of open sets, we have a map $\operatorname{res}_{V,U} : \mathcal{O}(V) \to \mathcal{O}(U)$.

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. The result is the same as if you restrict the differentiable function on the big open set directly to the small open set. In other words, if $U \hookrightarrow V \hookrightarrow W$, then the following diagram commutes:

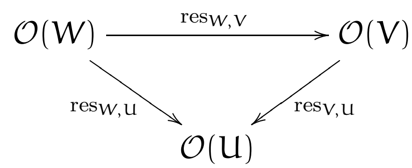

Next take two differentiable functions f_1 and f_2 on a big open set U, and an open cover of U by some U_i. Suppose that f_1 and f_2 agree on each of these U_i. Then they must have been the same function to begin with. In other words, if $\{U_i\}_{i \in I}$ is a cover of U, and $f_1, f_2 \in \mathcal{O}(U)$, and $\operatorname{res}_{U,U_i} f_1 = \operatorname{res}_{U,U_i} f_2$, then $f_1 = f_2$. Thus I can *identify* functions on an open set by looking at them on a covering by small open sets.

Finally, given the same U and cover U_i, take a differentiable function on each of the U_i — a function f_1 on U_1, a function f_2 on U_2, and so on — and they agree on the pairwise overlaps. Then they can be "glued together" to make one differentiable function on all of U. In other words, given $f_i \in \mathcal{O}(U_i)$ for all i, such that $\operatorname{res}_{U_i, U_i \cap U_j} f_i = \operatorname{res}_{U_j, U_i \cap U_j} f_j$ for all i, j, then there is some $f \in \mathcal{O}(U)$ such that $\operatorname{res}_{U,U_i} f = f_i$ for all i.

The entire example above would have worked just as well with continuous function, or smooth functions, or just functions. Thus all of these classes of "nice" functions share some common properties; we will soon formalize these properties in the notion of a sheaf.

3.1. Motivating example continued: the germ of a differentiable function. Before we do, we first 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 $W \subset U, V$ containing x where $f|_W = g|_W$ (or in our earlier language, $\operatorname{res}_{U,W} f = \operatorname{res}_{V,W} g$). In other words, two functions that are the same in a neighborhood of x but (but may differ elsewhere) have the same germ. We 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 $U \cap V$. Moreover, $f + g$ is well-defined: if f' has the same germ as f, meaning that there is some open set W containing x on which they agree, and g' has the same germ as g, meaning they agree on some open W' containing x, then $f' + g'$ is the same function as $f + g$ on $U \cap V \cap W \cap W'$.

Notice also that if $x \in U$, you get a map $\mathcal{O}(U) \to \mathcal{O}_x$. Experts may already see that this is secretly a colimit.

We can see that $\mathcal{O}_x$ is a local ring as follows. Consider those germs vanishing at x, which we denote $\mathfrak{m}_x \subset \mathcal{O}_x$. They certainly form an ideal: $\mathfrak{m}_x$ is closed under addition, and when you multiply something vanishing at x by any other function, the result also vanishes 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. We check that this ideal is maximal by showing that the quotient map is a field:

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

3.A. EXERCISE (FOR THOSE FAMILIAR WITH DIFFERENTIABLE FUNCTIONS). Show that this is the only maximal ideal of $\mathcal{O}_x$.

Note that 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. (We will see that this doesn't work for more general sheaves, but *does* work for things behaving like sheaves of functions. This will be formalized in the notion of a *locally ringed space*, which we will see only briefly later.)

Side fact for those with more geometric experience. Notice that $\mathfrak{m}/\mathfrak{m}^2$ 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 insight will prove handy later, when we define tangent and cotangent spaces of schemes.

4. DEFINITION OF SHEAF AND PRESHEAF

We now formalize these notions, by defining presheaves and sheaves. Presheaves are simpler to define, and notions such as kernel and cokernel are straightforward — they are defined "open set by open set". Sheaves are more complicated to define, and some notions such as cokernel require more thought (and the notion of sheafification). But we like sheaves are useful because they are in some sense geometric; you can get information about a sheaf locally.

4.1. Definition of sheaf and presheaf on a topological space X.

To be concrete, we will define sheaves of sets. However, **Sets** can be replaced by any category, and other important examples are abelian groups **Ab**, k-vector spaces, rings, modules over a ring, and more. Sheaves (and presheaves) are often written in calligraphic font, or with an underline. The fact that $\mathcal{F}$ is a sheaf on a topological space X is often written as

$$\begin{array}{c} \mathcal{F} \\ \Big| \\ X \end{array}$$

4.2. Definition: Presheaf. A *presheaf* $\mathcal{F}$ on a topological space X 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: $\mathcal{F}(U) = \Gamma(U, \mathcal{F}) = H^0(U, \mathcal{F})$. We will use them all.) The elements of $\mathcal{F}(U)$ are called *sections of* $\mathcal{F}$ *over* U.

• For each inclusion $U \hookrightarrow V$ of open sets, 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 $\mathrm{res}_{U,U}$ is the identity: $\mathrm{res}_{U,U} = \mathrm{id}_{\mathcal{F}(U)}$.

• If $U \hookrightarrow V \hookrightarrow W$ are inclusions of open sets, then the restriction maps commute, i.e.

$$\begin{array}{ccc} \mathcal{F}(W) & \xrightarrow{\mathrm{res}_{W,V}} & \mathcal{F}(V) \\ & {\scriptstyle \mathrm{res}_{W,U}} \searrow \quad \swarrow {\scriptstyle \mathrm{res}_{V,U}} & \\ & \mathcal{F}(U) & \end{array}$$

commutes.

4.A. Interesting exercise for category-lovers: "A presheaf is the same as a contravariant functor". Given any topological space X, we can get a category, called the "category of open sets" (discussed last week), where the objects are the open sets and the morphisms are inclusions. 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. (This interpretation is suprisingly useful.)

4.3. Definition: Stalks and germs. We define the stalk of a sheaf at a point in two different ways. In essense, one will be hands-on, and the other will be categorical using universal properties (as a colimit).

4.4. We will define the *stalk of* $\mathcal{F}$ *at* x to be the set of *germs* of a presheaf $\mathcal{F}$ at a point x, $\mathcal{F}_x$, as in the example of §3.1. 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 $\mathrm{res}_{U,W} f = \mathrm{res}_{V,W} g$. Elements of the stalk correspond to sections over some open set containing x. Two of these sections are considered the same if they agree on some smaller open set.

4.5. A useful (and better) equivalent definition of a stalk is as a colimit of all $\mathcal{F}(U)$ over all open sets U containing x:

$$\mathcal{F}_x = \varinjlim \mathcal{F}(U).$$

(Those having thought about the category of open sets will have a warm feeling in their stomachs.) The index category is a directed set (given any two such open sets, there is a third such set contained in both), so these two definitions are the same. 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.

Elements of the stalk $\mathcal{F}_x$ are called *germs*. If $x \in U$, and $f \in \mathcal{F}(U)$, then the image of f in $\mathcal{F}_x$ is called the *germ of* f.

I repeat that it is useful to think of stalks in both ways, as colimits, and also explicitly: a germ at p has as a representative a section over an open set near p.

If $\mathcal{F}$ is a sheaf of rings, then $\mathcal{F}_x$ is a ring, and ditto for rings replaced by abelian groups (or indeed any category in which colimits exist).

(Warning: the value at a point of a section doesn't make sense.)

4.6. Definition: Sheaf. A presheaf is a *sheaf* if it satisfies two more axioms, which will use the notion of when some open sets cover another.

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

(A presheaf satisfying the identity axiom is sometimes called a *separated presheaf*, but we will not use that notation in any essential way.)

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 $\operatorname{res}_{U_i, U_i \cap U_j} f_i = \operatorname{res}_{U_j, U_i \cap U_j} f_j$ for all i, j, then there is some $f \in \mathcal{F}(U)$ such that $\operatorname{res}_{U,U_i} f = f_i$ for all i.

(For experts, and scholars of the empty set only: an additional axiom sometimes included is that $F(\varnothing)$ is a one-element set, and in general, for a sheaf with values in a category, $F(\varnothing)$ is required to be the final object in the category. As pointed out by Kirsten, this actually follows from the above definitions, assuming that the empty product is appropriately defined as the final object.)

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 $F(\varnothing)$ is the final object.)

The *stalk of a sheaf* at a point is just its stalk as a presheaf; the same definition applies.

Philosophical note. In mathematics, definitions often come paired: "at most one" and "at least one". In this case, identity means there is at most one way to glue, and gluability means that there is at least one way to glue.

4.B. UNIMPORTANT EXERCISE FOR CATEGORY-LOVERS. The gluability axiom may be interpreted as saying that $\mathcal{F}(\cup_{i \in I} U_i)$ is a certain limit. What is that limit?

We now give a number of examples of sheaves.

4.7. *Example.* (a) Verify that the examples of §3 are indeed sheaves (of differentiable functions, or continuous functions, or smooth functions, or functions on a manifold or $\mathbb{R}^n$).
(b) Show that real-valued continuous functions on (open sets of) a topological space X form a sheaf.

4.8. *Important Example: Restriction of a sheaf.* Suppose $\mathcal{F}$ is a sheaf on X, and $U \subset$ is an open set. Define the *restriction of $\mathcal{F}$ to* U, denoted $\mathcal{F}|_U$, to be the collection $\mathcal{F}|_U(V) = \mathcal{F}(V)$ for all $V \subset U$. Clearly this is a sheaf on U.

4.9. *Important Example: skyscraper sheaf.* Suppose X is a topological space, with $x \in X$, and S is a set. Then S_x defined by $\mathcal{F}(U) = S$ if $x \in U$ and $\mathcal{F}(U) = \{e\}$ if $x \notin U$ forms a sheaf. Here $\{e\}$ is any one-element set. (Check this if it isn't clear to you.) This is called a *skyscraper sheaf*, because the informal picture of it looks like a skyscraper at x. There is an analogous definition for sheaves of abelian groups, except $\mathcal{F}(U) = \{0\}$ if $x \in U$; and for sheaves with values in a category more generally, $\mathcal{F}(U)$ should be a final object. (Warning: the notation S_x is not ideal, as the subscript of a point will also used to denote a stalk.)

4.C. IMPORTANT EXERCISE: CONSTANT PRESHEAF AND LOCALLY CONSTANT SHEAF. (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. We denote this presheaf $\underline{S}^{pre}$. Show that this needn't form a sheaf. This is called the *constant presheaf with values in* S.
(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. Show that this is a *sheaf*. (A better description is this: endow S with the discrete topology, and let $\mathcal{F}(U)$ be the continuous maps $U \to S$. Using this description, this follows immediately from Exercise 4.E below.) We will call this the *locally constant sheaf*. This is usually called the *constant sheaf*. We denote this sheaf $\underline{S}$.

4.D. UNIMPORTANT EXERCISE: MORE EXAMPLES OF PRESHEAVES THAT ARE NOT SHEAVES. Show that the following are presheaves on $\mathbb{C}$ (with the usual topology), but not sheaves: (a) bounded functions, (b) holomorphic functions admitting a holomorphic square root.

4.E. 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. (Example 4.7(b), with $Y = \mathbb{R}$, and Exercise 4.C(b), with $Y = S$ with the discrete topology, are both special cases.)

4.F. EXERCISE. This is a fancier example of the previous exercise.
(a) 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. (For those who have heard of vector bundles, these are a good example.)
(b) (This exercise is for those who know what a topological group is. If you don't know what a topological group is, you might be able to guess.) Suppose that Y is a topological group. Show that maps to Y form a sheaf of *groups*. (Example 4.7(b), with $Y = \mathbb{R}$, is a special case.)

4.10. ⋆ *The espace étalé of a (pre)sheaf.* Depending on your background, you may prefer the following perspective on sheaves, which we will not discuss further. Suppose $\mathcal{F}$ is a presheaf (e.g. a sheaf) on a topological space X. Construct a topological space Y along with a continuous map to X as follows: as a set, Y is the disjoint union of all the stalks of X. This also describes a natural set map $Y \to X$. We topologize Y as follows. Each section s of $\mathcal{F}$ over an open set U determines a section of $Y \to X$ over U, sending s to each of its germs for each $x \in U$. The topology on Y is the weakest topology such that these sections are continuous. This is called the **espace étalé** of the Then the reader may wish to show that (a) if $\mathcal{F}$ is a sheaf, then the sheaf of sections of $Y \to X$ (see the previous exercise 4.F(a) can be naturally identified with the sheaf $\mathcal{F}$ itself. (b) Moreover, if $\mathcal{F}$ is a presheaf, the sheaf of sections of $Y \to X$ is the *sheafification* of $\mathcal{F}$ (to be defined later).

4.G. IMPORTANT EXERCISE: THE DIRECT IMAGE SHEAF OR PUSHFORWARD SHEAF. 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 *direct image sheaf* of *pushforward sheaf*. More precisely, $f_*\mathcal{F}$ is called the *pushforward of $\mathcal{F}$ by* f.

The skyscraper sheaf (Exercise 4.9) can be interpreted as follows as the pushforward of the constant sheaf $\underline{S}$ on a one-point space x, under the morphism $f : \{x\} \to X$.

Once we realize that sheaves form a category, we will see that the pushforward is a functor from sheaves on X to sheaves on Y.

4.H. EXERCISE (PUSHFORWARD INDUCES MAPS OF STALKS). Suppose $\mathcal{F}$ is a sheaf of sets (or rings or A-modules). If $f(x) = y$, describe the natural morphism of stalks $(f_*\mathcal{F})_y \to \mathcal{F}_x$. (You can use the explicit definition of stalk using representatives, §4.4, or the universal property, §4.5. If you prefer one way, you should try the other.)

4.11. Important Example: Ringed spaces, and $\mathcal{O}_X$-modules.. Suppose $\mathcal{O}_X$ is a sheaf of *rings* on a topological space X (i.e. a sheaf on X with values in the category of **Rings**). Then $(X, \mathcal{O}_X)$ is called a *ringed space*. The sheaf of rings is often denoted by $\mathcal{O}_X$; this is pronounced "oh-of-X". This sheaf is called the *structure sheaf* of the ringed space. We now define the notion of an $\mathcal{O}_X$*-module*. The notion is analagous to one we've seen before:

just as we have modules over a ring, we have $\mathcal{O}_X$-modules over the structure sheaf (of rings) $\mathcal{O}_X$.

There is only one possible definition that could go with this name. An $\mathcal{O}_X$-module is a sheaf of abelian groups $\mathcal{F}$ with the following additional structure. For each U, $\mathcal{F}(U)$ is a $\mathcal{O}_X(U)$-module. Furthermore, this structure should behave well with respect to restriction maps. This means the following. If $U \subset V$, then

(3)
$$\begin{array}{ccc} \mathcal{O}_X(V) \times \mathcal{F}(V) & \xrightarrow{\text{action}} & \mathcal{F}(V) \\ \downarrow{\scriptstyle \mathrm{res}_{V,U}} & & \downarrow{\scriptstyle \mathrm{res}_{V,U}} \\ \mathcal{O}_X(U) \times \mathcal{F}(U) & \xrightarrow{\text{action}} & \mathcal{F}(U) \end{array}$$

commutes. (You should convince yourself that I haven't forgotten anything.)

Recall that the notion of A-module generalizes the notion of abelian group, because an abelian group is the same thing as a $\mathbb{Z}$-module. Similarly, the notion of $\mathcal{O}_X$-module generalizes the notion of sheaf of abelian groups, because the latter is the same thing as a $\underline{\mathbb{Z}}$-module, where $\underline{\mathbb{Z}}$ is the locally constant sheaf with values in $\mathbb{Z}$. Hence when we are proving things about $\mathcal{O}_X$-modules, we are also proving things about sheaves of abelian groups.

4.12. *For those who know about vector bundles.* The motivating example of $\mathcal{O}_X$-modules is the sheaf of sections of a vector bundle. If X is a differentiable manifold, and $\pi : V \to X$ is a vector bundle over X, then the sheaf of differentiable sections $\phi : X \to V$ is an $\mathcal{O}_X$-module. Indeed, given a section s of π over an open subset $U \subset X$, and a function f on U, we can multiply s by f to get a new section fs of π over U. Moreover, if V is a smaller subset, then we could multiply f by s and then restrict to V, or we could restrict both f and s to V and then multiply, and we would get the same answer. That is precisely the commutativity of (3).

Next day: We know about presheaves and sheaves, so we naturally ask about morphisms between presheaves and morphisms of presheaves.

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