

18.726: Algebraic Geometry (K.S. Kedlaya, MIT, Spring 2009)
Sheaves of modules (updated 27 Feb 09)

Having discussed sheaves of sets, abelian groups, and rings, we now consider sheaves of modules over a locally ringed space, with an emphasis on the situation for schemes. We'll then use what we know to talk about closed immersions and separated morphisms. (Many more properties of morphisms to follow in the next lectures!) References: Hartshorne II.4, II.5; EGA I 0.4, I 0.5, I 1.1.

1 Sheaves of modules

Let (X, \mathcal{O}_X) be a ringed space. A *sheaf of \mathcal{O}_X -modules* is a sheaf of sets \mathcal{F} such that for each $U \subseteq X$ open, the group of sections $\mathcal{F}(U)$ is equipped with a structure of a $\mathcal{O}_X(U)$ -module, in a fashion compatible with restriction. We sometimes call such a thing simply an \mathcal{O}_X -module. A *morphism* of two \mathcal{O}_X -modules $\mathcal{F} \rightarrow \mathcal{G}$ is a morphism of sheaves which gives a $\mathcal{O}_X(U)$ -module homomorphism $\mathcal{F}(U) \rightarrow \mathcal{G}(U)$ for each $U \subseteq X$ open.

Fun fact: for \mathcal{F} a sheaf of \mathcal{O}_X -modules, there is a natural bijection between

$$\mathrm{Hom}(\mathcal{O}_X, \mathcal{F}) \rightarrow \mathcal{F}(X).$$

If \mathcal{F} is a sheaf of \mathcal{O}_X -modules, then for each $x \in X$, the stalk \mathcal{F}_x inherits a structure of a $\mathcal{O}_{X,x}$ -module. Using that, we can talk about submodules, quotient modules, kernels, cokernels, images, and the like; in fact, these agree with the corresponding notions in the category of sheaves of abelian groups.

An \mathcal{O}_X -module is *free* if it is isomorphic to a direct sum \mathcal{O}_X^I of copies of \mathcal{O}_X . This property cannot be checked locally; if \mathcal{O}_X is only locally isomorphic to a free \mathcal{O}_X -module, we say it is *locally free*. More on locally free sheaves later.

One new operation is the tensor product, but it requires some care. If \mathcal{F}, \mathcal{G} are two \mathcal{O}_X -modules, the presheaf

$$U \mapsto \mathcal{F}(U) \otimes_{\mathcal{O}_X} \mathcal{G}(U)$$

may not be a sheaf. Form its sheafification, and call that $\mathcal{F} \otimes_{\mathcal{O}_X} \mathcal{G}$; it has the expected arrow-theoretic behavior. (I might forget the \mathcal{O}_X sometimes when it is not ambiguous.)

2 Direct and inverse image

If $f : (X, \mathcal{O}_X) \rightarrow (Y, \mathcal{O}_Y)$ is a morphism of ringed spaces, and \mathcal{F} is a \mathcal{O}_X -module, then $f_*\mathcal{F}$ may naturally be viewed as a $f_*\mathcal{O}_X$ -module. Using the map $f^\sharp : \mathcal{O}_Y \rightarrow f_*\mathcal{O}_X$, we may also give $f_*\mathcal{F}$ the structure of an \mathcal{O}_Y -module. We call this again the *direct image* of \mathcal{F} .

On the other hand, if \mathcal{G} is a sheaf of \mathcal{O}_Y -modules, then $f^{-1}\mathcal{G}$ is an $f^{-1}\mathcal{O}_Y$ -module. Adjointness turns f^\sharp into a homomorphism $f^{-1}\mathcal{O}_Y \rightarrow \mathcal{O}_X$, so we can form the tensor product

$$\mathcal{G} \otimes_{f^{-1}\mathcal{O}_Y} \mathcal{O}_X$$

and get a \mathcal{O}_X -module. We notate this $f^*\mathcal{G}$ and call it the (*module-theoretic*) *inverse image* of \mathcal{G} under f .

Again, f^* and f_* are adjoint in the obvious fashion. Statement and proof left to the reader.

3 Quasicoherent sheaves of modules

Since thinking about affine schemes is supposed to be equivalent to thinking about rings (after all, the two categories are equivalent!), we would like our thinking about sheaves of modules on affine schemes to be equivalent to thinking about modules over rings. We even know what the functors realizing this equivalence should be: on an affine scheme $X = \text{Spec}(R)$, we should go from R -modules to \mathcal{O}_X -modules via $M \mapsto \tilde{M}$ (which we already proved is a sheaf!), and back via $\mathcal{F} \mapsto \Gamma(X, \mathcal{F})$.

While it is clear that the composition one way is naturally isomorphic to the identity functor on R -modules, the other way fails because there are “too many” \mathcal{O}_X -modules. For instance, for $R = k[x]$ with k an algebraically closed field, we can make an \mathcal{O}_X -module \mathcal{F} by declaring

$$\mathcal{F}(U) = \begin{cases} 0 & (0) \notin U \\ k[x] & (0) \in U \end{cases}$$

and putting in the obvious restriction maps. This sheaf has the same global sections as \mathcal{O}_X itself but is clearly not the same!

The fix for this is a bit heavy-handed: we simply declare that we only want sheaves which locally come from a module over a ring. For schemes, it is clear what this means: we want \mathcal{F} to have the property that for each $x \in X$, there is an open affine neighborhood $U = \text{Spec}(A)$ of x in X such that $\mathcal{F}|_U = \tilde{M}$ for some A -module M . For locally ringed spaces, we must be a bit more careful: we want \mathcal{F} to locally be given as the cokernel of some module homomorphism $\mathcal{O}_X^I|_U \rightarrow \mathcal{O}_X^J|_U$ between free \mathcal{O}_X -modules.

It still remains to check that this gives what we expect for affine schemes. I call this the *third fundamental theorem of affine schemes*.

Theorem 1. *Let \mathcal{F} be a quasicoherent sheaf of \mathcal{O}_X -modules for $X = \text{Spec}(R)$, and put $M = \Gamma(X, \mathcal{F})$. Then the natural homomorphism $\tilde{M} \cong \mathcal{F}$ of \mathcal{O}_X -modules is in fact an isomorphism. In other words, the category of quasicoherent \mathcal{O}_X -modules on $\text{Spec}(R)$ is equivalent to the category of modules on R .*

Proof. The claim is equivalent to the fact that for each prime ideal \mathfrak{p} of R , the natural map $M_{\mathfrak{p}} \cong \tilde{M}_{\mathfrak{p}} \rightarrow \mathcal{F}_{\mathfrak{p}}$ is a bijection. Since \mathcal{F} is quasicompact, we can find a distinguished open $D(f)$ of X on which $\mathcal{F} \cong \tilde{N}$ for some R_f -module N . The map $\tilde{M} \rightarrow \mathcal{F}$ induces a map $\tilde{M}_f \cong \tilde{M}|_{D(f)} \rightarrow \mathcal{F}|_{D(f)} \cong \tilde{N}$. Taking global sections gives a homomorphism $M_f \cong N$ of R_f -modules.

We check injectivity of $M_f \rightarrow N$. Suppose $m/f^h \in M_f$ maps to zero in N . For each prime ideal \mathfrak{q} of R , we can find a distinguished open neighborhood $D(g)$ of \mathfrak{q} in X such that

$\mathcal{F}|_{D(g)} \cong \tilde{P}$ for some R_g -module P . Now $P_f \cong N_g \cong M_{fg}$ as R_{fg} -modules since they all give rise to the same sheaf. Hence the image of m/f^h in P_f is zero, so the image of m in $P = \Gamma(D(g), \mathcal{F})$ is killed by some power of f . We conclude that for some nonnegative integer j , $f^j m$ restricts to the zero section of \mathcal{F} on $D(g)$. Since this holds for each \mathfrak{q} , and we only need finitely many $D(g)$ to cover X (since X is quasicompact), there exists a nonnegative integer j such that $f^j m = 0$ in $\Gamma(X, \mathcal{F}) = M$. Hence m/f^h represents the zero element in M_f .

We check surjectivity of $M_f \rightarrow N$. Let $n \in N$ be any class. Cover X with finitely many $D(g_i)$ on each of which $\mathcal{F}|_{D(g_i)}$ is represented by a module P_i over R_{g_i} . Then $\mathcal{F}_{D(fg_i)}$ is represented by $(P_i)_{g_i}$. Hence for some j , $f^j n$ is the restriction to $D(fg_i)$ of a section s_i of \mathcal{F} over $D(g_i)$. We may enlarge j so that it works for all i ; then $s_i - s_j$ represents the zero section of \mathcal{F} over $D(fg_i g_j)$. By the previous paragraph, we can find some k such that $f^k(s_i - s_j)$ is the zero section of \mathcal{F} over $D(g_i g_j)$. Ergo, for some very large j , the $f^j n$ give sections s_i of \mathcal{F} over $D(g_i)$ which glue to give a section of \mathcal{F} on X itself.

We now have $M_f \cong N$, so $M_{\mathfrak{p}} \cong N_{\mathfrak{p}} \cong \mathcal{F}_{\mathfrak{p}}$. Since $\mathfrak{p} \in X$ was arbitrary, this proves $\tilde{M} \cong \mathcal{F}$. \square

4 Relative Spec, ideal sheaves, and closed immersions

Let X be a scheme, and let \mathcal{F} be a quasicoherent sheaf of \mathcal{O}_X -modules which also carries an \mathcal{O}_X -algebra structure (or for short, a *quasicoherent \mathcal{O}_X -algebra*). Then for each open affine subscheme U of X , we can form the scheme $\text{Spec } \Gamma(U, \mathcal{F})$, and these glue. We thus obtain a scheme $Y = \mathbf{Spec } \mathcal{F}$ which comes with a morphism $Y \rightarrow X$. This is called the *relative spectrum* of \mathcal{F} .

One important class of examples of relative spectra are closed immersions. (For more examples, see exercises.) An *ideal sheaf* on a locally ringed space X is a quasicoherent subsheaf of \mathcal{O}_X . For \mathcal{I} an ideal sheaf, the quotient $\mathcal{O}_X/\mathcal{I}$ is a quasicoherent \mathcal{O}_X -algebra. Let Z be its relative Spec; any map arising as such a map $f : Z \rightarrow X$ is called a *closed immersion*.

Let us see why this name is fitting in the case of schemes. Say $X = \text{Spec } R$. Then an ideal sheaf corresponds to an ideal I of R , and $Z = \text{Spec}(R/I)$. (Notice that this means that any closed immersion is an affine and finite morphism!) The points of Z are in bijection with the points $\mathfrak{p} \in X$ where the stalks $R_{\mathfrak{p}}$ and $I_{\mathfrak{p}}$ differ, which is precisely the vanishing set $V(I)$. But there can be many different closed immersions with the same underlying set! For instance, in $\text{Spec } k[x, y]$, the ideals (x) and (x^2) define the same closed set but not the same closed immersion.

Beware that algebraic geometers have the habit of calling the source of a closed immersion a *closed subscheme* of X even though it's not really a subscheme of X in any precise sense. But this isn't *so* misleading because the map $Z \rightarrow X$ is indeed a monomorphism. And there is something comforting in the thought that (looking at our previous example) the x, y -plane "contains" a doubled line defined by $x^2 = 0$, which in turn "contains" the undoubled line defined by $x = 0$.

Finally, note that if $f : Z \rightarrow X$ is a closed immersion, the defining ideal sheaf can be recovered as the kernel of $f^\# : \mathcal{O}_X \rightarrow f_*\mathcal{O}_Z$. In fact, following Hartshorne, you may define a closed immersion as a map of schemes $f : Z \rightarrow X$ which induces a homeomorphism of Z with a closed subset of X , such that $f^\# : \mathcal{O}_X \rightarrow f_*\mathcal{O}_Z$ is surjective. This has the advantage that it is clear that the property of being a closed immersion is local on the target (though you can check it the other way too, since the ideal sheaf is uniquely determined).

5 Separated schemes and morphisms

We are now ready to introduce the analogue of the Hausdorff property for schemes. However, it is more natural to introduce it for morphisms of schemes.

A morphism $f : Y \rightarrow X$ is *separated* if the *diagonal morphism* $\Delta : Y \rightarrow Y \times_X Y$ is a closed immersion. Since the formation of Δ , and the property of it being a closed immersion, are local on the target, so is the notion of being separated.

For instance, if $X = \text{Spec } k$ for k a field, and Y consists of two copies of $\text{Spec } k[x]$ glued along $\text{Spec } k[x, x^{-1}]$, then $Y \times_X Y$ is the affine plane with *two* copies of each axis, and *four* copies of the origin. The image of Δ contains two copies of the origin, but its closure contains all four. Hence $Y \rightarrow X$ is not separated.

Lemma. *A morphism $f : Y \rightarrow X$ is separated if and only if the image of the diagonal $\Delta : Y \rightarrow Y \times_X Y$ is a closed subset of $Y \times_X Y$.*

Proof. See Hartshorne, Corollary II.4.2. It relies on another useful (but easy) fact: any *affine* morphism is separated. \square

We say a scheme itself is *separated* if its unique morphism to $\text{Spec}(\mathbb{Z})$ is separated. This means that by fiat $\text{Spec}(\mathbb{Z})$ is itself separated, but this seems reasonable enough, especially in light of the following.

Theorem 2. *Let X be a separated scheme. Then the intersection of any two open affine subschemes of X is again affine.*

Proof. Exercise. (Beware that the converse is false. The nonseparated example above also satisfies this condition.) \square

There is a valuative criterion for separatedness, but it is hardly ever useful (because it involves arbitrary valuation rings, which can be rather nasty); see Hartshorne Theorem II.4.3.

6 Separatedness and base change

Theorem 3. *Separatedness is stable under base change.*

Proof. Let $f : Y \rightarrow X$ be a separated morphism, and let $Z \rightarrow X$ be an arbitrary morphism. We are supposed to check that the diagonal

$$Y \times_X Z \rightarrow (Y \times_X Z) \times_Z (Y \times_X Z)$$

is a closed immersion. On the other hand, since closed immersions are stable under base change, we already know that

$$Y \times_X Z \rightarrow (Y \times_X Y) \times_X Z$$

is a closed immersion. It thus suffices to identify the two right sides in a way commuting with the arrows.

The way to see this is to use the functor of points: given $W \rightarrow X$, we must identify the maps into the two right hand sides in a natural way. The maps to $(Y \times_X Y) \times_X Z$ give pairs of maps $W \rightarrow Y \times_X Y$ and $W \rightarrow Z$ which agree on X , then (by splitting the first fibre product) gives triples of maps $W \rightarrow Y$, $W \rightarrow Y$, and $W \rightarrow Z$ which all agree on X .

The maps to $(Y \times_X Z) \times_Z (Y \times_X Z)$ similarly give pairs of maps $W \rightarrow (Y \times_X Z)$, $W \rightarrow Y \times_X Z$ which agree on Z . Splitting again, we get quadruples $W \rightarrow Y$, $W \rightarrow Z$, $W \rightarrow Y$, $W \rightarrow Z$ which agree on X , but moreover the two maps $W \rightarrow Z$ must be the *same map*. We thus identify with the previous description. □