18.726 Algebraic Geometry
Spring 2009

18.726: Algebraic Geometry (K.S. Kedlaya, MIT, Spring 2009)
More on abelian sheaves

We now specialize the discussion of sheaves to the situation where the target category consists of abelian groups. At the end, I'll explain how to generalize to the case of a target which is an *abelian category*.

1 Abelian groups

Assume until I say otherwise that $\mathcal{C} = \underline{\mathrm{Ab}}$. (At the end, we'll generalize to the case where $\mathcal{C}$ can be any abelian category.) Let me first set some notation and terminology about morphisms of abelian groups themselves.

For $f : A \to B$ a morphism of abelian groups,

$$\begin{aligned} \ker(f) &= \{x \in A : f(x) = 0\} \\ \mathrm{im}(f) &= \{f(x) : x \in A\} \\ \mathrm{coker}(f) &= A/\mathrm{im}(f) = \{y + \mathrm{im}(f) : y \in B\}. \end{aligned}$$

A (finite or infinite) sequence

$$\cdots \to A_{i-1} \to A_i \to A_{i+1} \to \cdots$$

in $\underline{\mathrm{Ab}}$ is *exact* if for each i,

$$\mathrm{im}(A_{i-1} \to A_i) = \ker(A_{i+1} \to A_i).$$

If we only have the weaker assertion that

$$\mathrm{im}(A_{i-1} \to A_i) \subseteq \ker(A_{i+1} \to A_i)$$

(i.e., the composition $A_{i-1} \to A_i \to A_{i+1}$ is zero) for each i, we say that the sequence is a *complex*.

Here are some useful facts about exact sequences; their proofs are fun exercises in what is sometimes called *diagram chasing*. Remember that in $\underline{\mathrm{Ab}}$, monomorphism equals injective and epimorphism equals surjective (so mono plus epi equals iso, which is not true in an arbitrary category).

Lemma (Five lemma)**.** *Let*

$$\begin{array}{ccccccccc} A_0 & \longrightarrow & A_1 & \longrightarrow & A_2 & \longrightarrow & A_3 & \longrightarrow & A_4 \\ \downarrow{\scriptstyle f_0} & & \downarrow{\scriptstyle f_1} & & \downarrow{\scriptstyle f_2} & & \downarrow{\scriptstyle f_3} & & \downarrow{\scriptstyle f_4} \\ B_0 & \longrightarrow & B_1 & \longrightarrow & B_2 & \longrightarrow & B_3 & \longrightarrow & B_4 \end{array}$$

be a commuting diagram in $\mathcal{C}$ with exact rows.

(a) If f_1 and f_3 are monomorphisms and f_0 is an epimorphism, then f_2 is a monomorphism.

(b) If f_1 and f_3 are epimorphisms and f_4 is a monomorphism, then f_2 is an epimorphism.

Proof. Exercise. □

Lemma (Snake lemma)**.** *Let*

$$\begin{array}{ccccccccc} 0 & \longrightarrow & A_1 & \longrightarrow & A_2 & \longrightarrow & A_3 & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle f_1} & & \downarrow{\scriptstyle f_2} & & \downarrow{\scriptstyle f_3} & & \\ 0 & \longrightarrow & B_1 & \longrightarrow & B_2 & \longrightarrow & B_3 & \longrightarrow & 0 \end{array}$$

be a short exact sequence. Then there exists a canonical homomorphism $\delta : \ker(f_3) \to \operatorname{coker}(f_1)$ (the connecting homomorphism*) such that*

$$0 \to \ker(f_1) \to \ker(f_2) \to \ker(f_3) \xrightarrow{\delta} \operatorname{coker}(f_1) \to \operatorname{coker}(f_2) \to \operatorname{coker}(f_3) \to 0$$

is exact, where all the maps other than δ are the obvious ones induced by the diagram.

Proof. Here is what δ is supposed to be: given $a_3 \in \ker(f_3)$, lift it to $a_2 \in A_2$, then apply f_2 to get $b_2 \in B_2$. Since the diagram commutes, b_2 must map to zero in B_3, so it lifts to b_1 in B_1. Declare $\delta(a_3) = b_1$.

It remains to show that δ is well-defined and is a homomorphism, and that the claimed sequence is exact. These are left as exercises. □

Corollary (Short five lemma)**.** *Let*

$$\begin{array}{ccccccccc} 0 & \longrightarrow & A_1 & \longrightarrow & A_2 & \longrightarrow & A_3 & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle f_1} & & \downarrow{\scriptstyle f_2} & & \downarrow{\scriptstyle f_3} & & \\ 0 & \longrightarrow & B_1 & \longrightarrow & B_2 & \longrightarrow & B_3 & \longrightarrow & 0 \end{array}$$

be a commuting diagram in $\mathcal{C}$ with exact rows. Then f_2 is a monomorphism/epimorphism if and only if f_1 and f_3 both are.

2 Exact functors

For $\mathcal{C}_1 = \mathcal{C}_2 = \underline{\mathrm{Ab}}$, a covariant functor $F : \mathcal{C}_1 \to \mathcal{C}_2$ is *additive* if it commutes with addition of morphisms. Any additive functor sends complexes to complexes (because the property of the composition of two maps being zero is preserved), but not necessarily exact sequences to exact sequences. Hence the following definitions.

We say F is *left exact* if for any exact sequence

$$0 \to A_1 \to A_2 \to A_3$$

the sequence

$$0 \to F(A_1) \to F(A_2) \to F(A_3)$$

is exact. We say F is *right exact* if for any exact sequence

$$A_1 \to A_2 \to A_3 \to 0$$

the sequence

$$F(A_1) \to F(A_2) \to F(A_3) \to 0$$

is exact. We say F is *exact* if it is both left exact and right exact; equivalently, for any exact sequence

$$0 \to A_1 \to A_2 \to A_3 \to 0$$

the sequence

$$0 \to F(A_1) \to F(A_2) \to F(A_3) \to 0$$

It in turn implies that any exact sequence of any length goes into another exact sequence under F. (I'll try avoid using these notions for contravariant functors, since there is a left/right ambiguity.)

Examples:

- For any given $X \in \mathcal{C}$, the covariant functor $\mathrm{Hom}(X, \cdot)$ is left exact.
- For any given $X \in \mathcal{C}$, the covariant functor $X \otimes \cdot$ is right exact.

Many left/right exact functors arise from the following proposition.

Proposition. *Suppose the covariant functors $f^* : \mathcal{C}_1 \to \mathcal{C}_2$ and $f_* : \mathcal{C}_2 \to \mathcal{C}_1$ form an adjoint pair. Then f^* is right exact and f_* is left exact.*

Proof. Exercise. □

3 Abelian sheaves

Let $\mathcal{F}$ be a sheaf on a topological space X with values in $\mathcal{C} = \underline{\mathrm{Ab}}$. A *subsheaf* of $\mathcal{F}$ is what you think: take a subset of the sections on each open so that you still have a sheaf. The *quotient* of $\mathcal{F}$ by a subsheaf $\mathcal{G}$ is a bit trickier: take the presheaf $U \mapsto \mathcal{F}(U)/\mathcal{G}(U)$, then sheafify. Note that the stalk at x is indeed $\mathcal{F}_x/\mathcal{G}_x$.

Given a morphism $\phi : \mathcal{F} \to \mathcal{G}$ of sheaves, the presheaf $U \mapsto \ker(\phi(U))$ is a sheaf; we call it the *kernel* of ϕ. The presheaves $U \mapsto \mathrm{im}(\phi(U))$ and $U \mapsto \mathrm{coker}(\phi(U))$ are not in general sheaves; their sheafifications are called the *image* and *cokernel* of ϕ.

Proposition. *For $x \in X$, we have $\ker(\phi)_x = \ker(\phi_x)$, $\mathrm{im}(\phi)_x = \mathrm{im}(\phi_x)$, and $\mathrm{coker}(\phi)_x = \mathrm{coker}(\phi_x)$. Consequently,*

$$\mathrm{im}(\phi) \cong \mathcal{F}/\ker(\phi), \qquad \mathrm{coker}(\phi) \cong \mathcal{G}/\mathrm{im}(\phi).$$

Proof. Exercise. □

Using these, we extend the notion of exactness to a sequence of sheaves; it's equivalent to define it using sheaves or stalks, but *not* using sections.

Let $\underline{\mathrm{Sh}}_{\mathcal{C}}(X)$ be the category of sheaves on X with values in $\mathcal{C}$. We define the *global sections functor* $\Gamma(\cdot, X) : \underline{\mathrm{Sh}}_{\mathcal{C}}(X) \to \mathcal{C}$ by the formula

$$\Gamma(\mathcal{F}, X) = \mathcal{F}(X).$$

(No set-theoretic difficulties here: $\underline{X}$ is a small category, so sheaves on X with values in $\mathcal{C}$ do form a class.)

Proposition. *The global sections functor is left exact.*

Proof. Exercise. □

The failure of the global sections functor to be right exact will give rise to the notion of *sheaf cohomology* later.

4 Abelian categories

Everything I defined above can be generalized to the case where $\mathcal{C}$ is what is called an *abelian category*, i.e., a category which captures the useful properties of abelian groups.

First, let me give an *ad hoc* definition which will suffice for our purposes. A *nice abelian category* is an additive category in which all limits and colimits exist, together with a forgetful functor to $\underline{\mathrm{Ab}}$ which preserves limits and colimits.

Next, let's figure out what the correct abstract definition shoul be. We first write down the definition of an *preadditive category* (which I called an *additive category* by mistake on Problem Set 1). That is a category $\mathcal{C}$ equipped with the structure of an abelian group on each set $\mathrm{Hom}(X, Y)$, over which composition is distributive.

We next define an *additive category*. The key notion is that direct sum and direct product coincide for a finite collection of abelian groups. We should thus require the existence of *biproducts*: that is, for any $X_1, \dots, X_n \in \mathrm{Obj}(\mathcal{C})$, there must exist an object X equipped with maps $\pi_i : X \to X_i$ and $\iota_i : X_i \to X$, such that X is both a product (using the π_i) and a coproduct (using the ι_i), and the sum $\iota_1 \circ \pi_1 + \cdots + \iota_n \circ \pi_n$ is the identity on X. (Exercise: this exists as soon as you have finite products.)

Since the empty biproduct exists, an additive category has a terminal (initial and final) object, which we call the *zero object* and label 0. In an additive category, we can define a *kernel* of the morphism $f : X \to Y$ to be a limit of the diagram

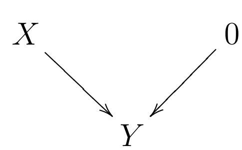

i.e., an object W plus a morphism $g : W \to X$ such that $f \circ g = 0$, and any other morphism $h : V \to X$ for which $f \circ h = 0$ factors uniquely through g. Similarly, a *cokernel* of f is a colimit of

$$\begin{array}{ccc} & X & \\ \swarrow & & \searrow \\ Y & & 0 \end{array}$$

To get a *preabelian category*, we insist that every morphism admit a kernel and cokernel (which as usual are only unique up to unique isomorphism). To get an *abelian category*, we insist that every monomorphism be the kernel of its cokernel, and every epimorphism be the cokernel of its kernel.

The *Freyd-Mitchell embedding theorem* asserts that at least for every *small* abelian category $\mathcal{C}$, we can construct an exact and fully faithful functor $F : \mathcal{C} \to \underline{\mathrm{Mod}}_R$ for a *not necessarily commutative* ring R (where $\underline{\mathrm{Mod}}_R$ now means *left* modules). This lets you prove theorems about abelian categories by reducing to situations where objects really do have elements.

The main difference between my nice abelian categories and true abelian categories is that I want *all* limits and colimits to exist. This is a bit strong for some purposes, but since I need limits anyway to work with sheaves, it's not so strange.

Anyway, the point here is that if you start with a (nice) abelian category $\mathcal{C}$, for any topological space X, the category $\underline{\mathrm{Sh}}_{\mathcal{C}}(X)$ is again a (nice) abelian category. This follows by assembling various homework exercises.