18.726: Algebraic Geometry (K.S. Kedlaya, MIT, Spring 2009)
Schemes

We next introduce locally ringed spaces, affine schemes, and general schemes. References: Hartshorne II.2, Eisenbud-Harris I.1, EGA 1.1.

1 Ringed and locally ringed spaces

A *ringed space* is a topological space X equipped with a sheaf $\mathcal{O}_X$ on X with values in $\underline{\mathrm{Ring}}$ (called the *structure sheaf*). This definition isn't so useful because it doesn't force the topology to have much to do with the ring structure; for instance, any ring can be viewed as a ringed space on a one-element topological space.

A more useful notion is that of a *locally ringed space*. This is a ringed space in which for each $x \in X$, the stalk $\mathcal{O}_{X,x}$ of $\mathcal{O}_X$ at x is a *local ring*, i.e., a ring with a unique maximal ideal $\mathfrak{m}_{X,x}$. (The zero ring is not a local ring!)

For example, suppose X is a manifold and let $\mathcal{O}_X$ be the sheaf of real-valued continuous functions. We check that $(X, \mathcal{O}_X)$ forms a locally ringed space. Given $x \in X$, let $\mathfrak{m}_{X,x}$ be the ideal of $\mathcal{O}_{X,x}$ consisting of germs of functions taking the value 0 at x. This is clearly an ideal, and the quotient $\mathcal{O}_{X,x}/\mathfrak{m}_{X,x}$ is certainly contained in $\mathbb{R}$. Since X is a manifold, the quotient is nonzero, so $\mathfrak{m}_{X,x}$ is indeed a maximal ideal of $\mathcal{O}_{X,x}$. To check that it is the unique maximal ideal, it suffices to check that any $f \in \mathcal{O}_{X,x}$ not contained in $\mathfrak{m}_{X,x}$ is a unit in $\mathcal{O}_{X,x}$. For such an f, $f(x)$ is some nonzero real number, so we can find an open subinterval $I \subseteq \mathbb{R}$ such that $f(x)$ belongs to I but 0 does not. Represent f by a continuous function on some open subset U of X containing x, which I'll also call f. The key point is that by continuity, $V = f^{-1}(I)$ is again an open subset of X containing x, and f takes nonzero values everywhere on V. Hence there exists a multiplicative inverse g of f on V, which is necessarily continuous.

Similarly, a smooth manifold, complex manifold, or abstract algebraic variety equipped with the obvious sheaf is a locally ringed space.

For any $x \in X$, the quotient $\mathcal{O}_{X,x}/\mathfrak{m}_{X,x}$ is a field. We denote it by $\kappa(x)$ and call it the *residue field* of x. In the aforementioned examples, the residue fields of all of the points of x are the same (either $\mathbb{R}$, $\mathbb{C}$, or a prescribed algebraically closed field), but that will not be the case for schemes!

I'll talk about morphisms of (locally) ringed spaces later. For the moment, let me at least point out that an *isomorphism* of (locally) ringed spaces is what you think: a homeomorphism of topological spaces and corresponding bijections of sections which commute with restriction.

2 The prime spectrum of a ring

The notion of a locally ringed space is a sufficiently broad generalization of manifolds that it admits a meaningful functor from the category of *arbitrary* (commutative unital) rings.

This gives rise to the concept of an affine scheme; to define this, we must first recall the construction of the prime spectrum of a ring. See the exercises for lots of examples.

Let R be an arbitrary ring. Following Zariski, we define the *prime spectrum* of R, denoted $\mathrm{Spec}(R)$, to be the set of prime ideals of R. (An ideal $\mathfrak{p}$ of R is *prime* if $R/\mathfrak{p}$ is an integral domain. The zero ring is not an integral domain, so the trivial ideal is not prime.) For a general ring, this is a better idea than using only maximal ideals because a ring homomorphism $\phi : R \to S$ induces a map $\mathrm{Spec}(S) \to \mathrm{Spec}(R)$ taking $\mathfrak{p} \subseteq S$ to $\phi^{-1}(\mathfrak{p})$. (The latter is prime because ϕ induces an *injective* map $R/\phi^{-1}(\mathfrak{p}) \to S/\mathfrak{p}$, so the source is an integral domain.) By contrast, ϕ may not carry maximal ideals of S to maximal ideals of R; for instance, consider $\phi : \mathbb{Z} \to \mathbb{Q}$.

Again following Zariski, we equip the set $\mathrm{Spec}(R)$ with the *Zariski topology*, in which the closed sets have the form

$$V(I) = \{\mathfrak{p} \in \mathrm{Spec}(R) : I \subseteq \mathfrak{p}\}$$

for I an ideal of R. This is indeed a topology because

$$V(I) \cup V(J) = V(I \cap J) = V(IJ)$$

$$\bigcap_i V(I_i) = V\left(\sum_i I_i\right).$$

We will use a special basis of open sets for this topology: the *distinguished open sets*, of the form

$$D(f) = \{\mathfrak{p} \in \mathrm{Spec}(R) : f \notin \mathfrak{p}\}$$

for f an element of R. Note that this basis is *nice* in the sense that the intersection of any two distinguished opens $D(f)$ and $D(g)$ is again a distinguished open, namely $D(fg)$. Note also that for $\phi : R \to S$ a homomorphism, the induced map $\mathrm{Spec}(S) \to \mathrm{Spec}(R)$ is continuous because the inverse image of $D(f)$ is $D(\phi(f))$.

Lemma. *Any distinguished open $D(f)$ of $\mathrm{Spec}(R)$ is quasicompact for the Zariski topology. In particular, $\mathrm{Spec}(R) = D(1)$ itself is quasicompact.*

Proof. It is enough to prove that any covering of $D(f)$ by distinguished open subsets admits a finite subcover. If the sets $D(f_i)$ cover $D(f)$ (for i running over some arbitrary index set), then the radical of (f) is contained in the radical of the ideal generated by the f_i. In particular, some power of f is in the ideal generated by the f_i. But that means that we can write f as a *finite* R-linear combination of the f_i, so those $D(f_i)$ already cover $D(f)$. □

For example, if k is an algebraically closed field, then $\mathrm{Spec}\, k[x]$ consists of one point of the form $(x - a)$ for each $a \in k$, plus a point corresponding to the prime ideal (0). The latter is an example of a *generic point*, a point whose closure is equal to the entire space in question. For the analogous picture of $\mathrm{Spec}\, k[x, y]$, see Hartshorne Example 2.3.4.

3 A presheaf of rings

We now specify a presheaf of rings on $X = \operatorname{Spec}(R)$, but only on the distinguished open subsets. To do this, we must do a bit of work to clean up their description, to account for the fact that prime ideals don't see the difference between an element of a ring and a power of that element.

Lemma. *For $f, g \in R$, we have $D(f) \subseteq D(g)$ if and only if some power of f is a multiple of g.*

Proof. Note that $D(f) = D(f^n)$ for any positive integer n. Hence if f^n is a multiple of g for some n, then $D(f) = D(f^n)$ is contained in $D(g)$. Conversely, suppose $D(f) \subseteq D(g)$, or in other words, $V(g) \subseteq V(f)$. Recall that the radical of the ideal (g) is the intersection of the prime ideals containing (g). Since $V(g) \subseteq V(f)$, it follows that the radical of (f) is contained in the radical of (g), so in particular f belongs to the radical of (g). That is, some power of f is a multiple of g, as desired. $\square$

A *multiplicative* subset of R is a subset closed under multiplication. For example, $S_f = \{1, f, f^2, f^3, \dots\}$ is a multiplicative subset. A multiplicative subset S is *saturated* if for any $x \in R$ such that some power of x equals an element of S times a unit, we have in fact $x \in S$. For any multiplicative subset S of R, there is a unique saturated multiplicative subset $\tilde{S}$ containing it, formed in the obvious fashon. By the previous lemma, we now have the following.

Corollary. *For $f, g \in R$, we have $D(f) = D(g)$ if and only if $\tilde{S}_f = \tilde{S}_g$.*

Given any multiplicative subset S of R, there is a unique initial object among the R-algebras in which each element of S has a multiplicative inverse. It is called the *localization* of R at S, denoted $S^{-1}R$. We can construct it as the polynomial ring in one variable x_f for each $f \in S$, modulo the relations $x_f f - 1$. Note that there is a canonical isomorphism $\tilde{S}^{-1}R \cong S^{-1}R$ since they both satisfy the same universal property. In particular, we can write

$$\tilde{S}_f^{-1}R \cong R[x]/(xf - 1).$$

From now on, write R_f instead of $\tilde{S}_f^{-1}R$.

Let D be the set of distinguished open subsets of $X = \operatorname{Spec} R$. Define a presheaf of rings $\mathcal{O}_X$ on X specified on D as follows. First put

$$\mathcal{O}_X(D(f)) = R_f;$$

this is well-defined by the previous corollary. Then note that given an inclusion $D(g) \subseteq D(f)$, we have $R_f \subseteq R_g$, so the universal property of localization gives a canonical homomorphism $R_f \to R_g$. If you want to write this more concretely (but less canonically), apply the lemma above to write $f^n = gh$ for some positive integer n, identify $\mathcal{O}_X(D(f)) = R[x]/(xf-1)$ and $\mathcal{O}_X(D(g)) = R[y]/(yg-1)$, and take the R-algebra homomorphism

$$R[x]/(xf-1) \to R[y]/(yg-1), \qquad x \mapsto f^{n-1}hy.$$

4 The fundamental theorem of affine schemes

We are now ready to prove what I call the *fundamental theorem of affine schemes*. I don't know whether its appearance in EGA 1 is its first.

Theorem 1. *The presheaf $\mathcal{O}_X$ on $X = \operatorname{Spec} R$ specified on D satisfies the sheaf axiom for coverings of distinguished opens by other distinguished opens. Consequently, it extends uniquely to a sheaf of rings on* $\operatorname{Spec} R$*.*

While we're at it, though, we may as well prove something stronger which we will need later. This proof is basically the same one used to compute the regular functions on an affine algebraic variety. It may also be thought of as an enhancement of the Chinese remainder theorem; indeed, the latter is an immediate corollary (exercise).

Theorem 2. *Let M be an R-module. Define a presheaf $\tilde{M}$ of abelian groups on X specified on D by the formula $D(f) \mapsto M \otimes_R R_f$. Then $\tilde{M}$ satisfies the sheaf axiom for coverings of distinguished opens by other distinguished opens. Consequently, it extends uniquely to a sheaf on* $\operatorname{Spec} R$*.*

Proof. By replacing R with R_f, we may reduce to checking the sheaf axiom for a cover of X itself by some distinguished open subsets $D(f_i)$. We first verify that the map $M \to \prod_i M \otimes_R R_{f_i}$ is injective, as follows. Suppose $m \in M$ belongs to the kernel of this map. Then the annihilator of m is an ideal of R which cannot be contained in any prime ideal $\mathfrak{p}$ of R, or else we would have $\mathfrak{p} \in D(f_i)$ for some i, and the image of m in $M \otimes_R R_{f_i}$ would be nonzero. Thus $1 \cdot m = 0$, so $m = 0$.

This proves the first half of the sheaf axiom; we must now check the glueing property. For this, we remember that X is quasicompact, so we may reduce to checking for a *finite* cover. Say $D(f_1), \dots, D(f_n)$ cover X. Suppose that some $D(f_i)$ cover $D(f)$, and that we are given elements $m_i/f_i^{h_i} \in M \otimes_R R_{f_i}$ such that $m_i/f_i^{h_i}$ and $m_j/f_j^{h_j}$ have the same image in $R_{f_i f_j}$. Since there are only finitely many f_i, we may take the nonnegative integers h_i to be equal to a common value h. For each i, j, we then have

$$(f_i f_j)^{g_{ij}} (f_i^h m_j - f_j^h m_i) = 0$$

for some nonnegative integers g_{ij}. By rechoosing the m_i, we can force $g_{ij} = 0$ for all i, j, that is, we now have literal equalities

$$f_i^h m_j = f_j^h m_i.$$

Since $D(f_i^h) = D(f_i)$, the $D(f_i^h)$ again cover X, so the f_i^h generate the unit ideal. We may now pick $g_1, \dots, g_n \in R$ such that $g_1 f_1^h + \cdots + g_n f_n^h = 1$. Put

$$m = g_1 m_1 + \cdots + g_n m_n.$$

We then have

$$f_i^h m = \sum_j f_i^h g_j m_j = \sum_j f_j^h g_j m_i = m_i,$$

so m is an element of M restricting to m_i/f_i^h for each i. This completes the proof of the glueing property, so we are done. $\square$

5 Schemes

From now on, we view $X = \mathrm{Spec}(R)$ as a ringed space with structure sheaf $\mathcal{O}_X$ as defined above. Note that for any prime ideal $\mathfrak{p}$ of R, the stalk $\mathcal{O}_{X,\mathfrak{p}}$ is canonically isomorphic to the local ring $R_\mathfrak{p}$ (the localization of R at the multiplicative set $R \setminus \mathfrak{p}$). Hence $\mathrm{Spec}(R)$ is in fact a *locally* ringed space.

At this point, we make schemes out of prime spectra by glueing, just as we would make manifolds out of open subspaces of $\mathbb{R}^n$. We define an *affine scheme* to be any locally ringed space X isomorphic to $\mathrm{Spec}(R)$ for some ring R; note that the ring R is uniquely determined by the fact that

$$\Gamma(\mathrm{Spec}(R), \mathcal{O}_{\mathrm{Spec}(R)}) = R$$

(from the previous theorem). A *scheme* is a locally ringed space in which each point has an open neighborhood isomorphic to an affine scheme.

Warning: if $X = \mathrm{Spec}(R)$ is an affine scheme, each distinguished open subset $D(f)$ is an affine scheme, namely $\mathrm{Spec}(R_f)$ (exercise). By construction, these form a basis of open sets. However, it is possible for there to be an open subset U of X such that $(U, \mathcal{O}_X|_U)$ is isomorphic to an affine scheme but U is not distinguished. (Counterexample to appear as an exercise.)

6 Schemes by glueing

We often specify nonaffine schemes using glueing data. For instance, if X_1 and X_2 are two schemes admitting open subsets U_1, U_2 which are isomorphic as locally ringed spaces, we can glue along this isomorphism to get a third scheme X. For more than two schemes, though, we must add a cocycle condition to keep the glueing maps consistent. Here is how that works.

Let us first specify glueing data for *sets*. Let $(X_i)_{i \in I}$ be a collection of sets. For each pair $(i,j) \in I \times I$, let U_{ij} be an open subset of X_i, and suppose that $U_{ii} = X_i$. Let $\phi_{ij} : U_{ij} \to U_{ji}$ be an isomorphism, and suppose that $\phi_{ii} = \mathrm{id}_{X_i}$. Suppose also that for $i,j,k \in I$, ϕ_{ij} restricts to an isomorphism of $U_{ij} \cap U_{ik}$ with $U_{ji} \cap U_{jk}$, and the cocycle condition

$$\phi_{ik} = \phi_{jk} \circ \phi_{ij}$$

holds on $U_{ij} \cap U_{ik}$. (In particular, $\phi_{ji} = \phi_{ij}^{-1}$.)

We would like to identify the X_i with subsets of a single set X in such a way that U_{ij} identifies with $X_i \cap X_j$ and ϕ_{ij} identifies with the identity map on $X_i \cap X_j$. To do this, first form the disjoint union $X' = \coprod_{i \in I} X_i$. Then define a binary relation on X' as follows: for $x_i \in X_i$ and $x_j \in X_j$, we say that $x_i \sim x_j$ if $x_i \in U_{ij}$, $x_j \in U_{ji}$, and $\phi_{ij}(x_i) = x_j$. The glueing conditions guarantee that this is an equivalence relation, so we may form the quotient X of X' by $\sim$; this gives the desired glueing. (Exercise: reformulate this definition in terms of a limit construction.)

We next specify glueing data for *topological spaces*. Set notation as above, except that each U_{ij} must be an *open* subset of X_i, and each ϕ_{ij} must be a *homeomorphism*. Using the glueing construction for sets, identify the X_i with subsets of a single set X. We may then use the topologies on the X_i as a basis for a topology on X; in particular, X_i is open in X.

We must still check, however, that the given topology on X_i coincides with the subspace topology from X (it is only obvious that the subspace topology is finer). Suppose $x_i \in X_i$ and V is an open neighborhood of x_i in X. There then exists some j such that $x_i \in X_j$ and V contains an open neighborhood of x_i for the topology on X_j. Since $x_i \in X_i \cap X_j = U_{ji}$ and the latter is open in X_j, $V \cap U_{ji}$ also contains an open neighborhood of x_i for the topology on X_j. Since ϕ_{ij} is a homeomorphism, $V \cap U_{ji} = V \cap U_{ij}$ contains an open neighborhood of x_i for the topology on X_i. This proves the claim.

We next specify glueing data for *(locally) ringed spaces*. Set notation as above, except that each X_i now carries a structure sheaf $\mathcal{O}_{X_i}$, and each ϕ_{ij} is an isomorphism of (locally) ringed spaces. Using the glueing construction for topological spaces, identify the X_i with open subsets of a single topological space X. By the glueing property for sheaves, we now obtain a sheaf of rings $\mathcal{O}_X$, so X may be viewed as a ringed space. Moreover, for $x \in X_i$, we have a canonical identification of $\mathcal{O}_{X,x}$ with $\mathcal{O}_{X_i,x}$; hence if each X_i is a locally ringed space, so is X.

We finally specify glueing data for schemes. This is the easy part: set notation as above, except that each X_i is a scheme. Then it is evident that X is also locally isomorphic to an affine scheme, so X is a scheme! (This part also works for manifolds and the like.)

7 Examples of glueing

Glueing can be a force for both good and evil. Let's start with good.

Start with any ring R. For $i = 0, \dots, n$, put

$$X_i = \operatorname{Spec} R[x_0/x_i, \dots, x_{i-1}/x_i, x_{i+1}/x_i, \dots, x_n/x_i].$$

Define the distinguished open subset

$$U_{ij} = D(x_j/x_i) \subset X_i.$$

Then there is an obvious isomorphism of U_{ij} with U_{ji} given by identifying x_k/x_i with $(x_k/x_j)(x_j/x_i)$. It is easy to check that the cocycle condition is satisfied, so we get a scheme $\mathbb{P}^n_R$, the *projective space* over R. (An alternate construction of projective space uses graded rings. More on this later.)

Now for the evil. Let k be an algebraically closed field. Let X_1 and X_2 be two copies of $\operatorname{Spec} k[x]$. We may glue these on the open sets obtained by removing the point $x = 0$ (i.e., the distinguished opens $D(x)$) to get a rather unpleasant object; it is a *line with a doubled point*.

We would like to formulate a condition that rules out such pathologies. In topology, the Hausdorff condition does the job, but that won't work for schemes. We need a more category-theoretical notion, which will be provided once we define *separatedness*.