FOUNDATIONS OF ALGEBRAIC GEOMETRY CLASS 5

RAVI VAKIL

CONTENTS

Last day: morphisms of (pre)sheaves; properties determined at the level of stalks; sheaves of abelian groups on X (and $\mathcal{O}_X$-modules) form an abelian category.

1. THE INVERSE IMAGE SHEAF

We next describe a notion that is rather fundamental, but is still a bit intricate. We won't need it (at least for a long while), so this may be best left for a second reading. Suppose we have a continuous map $f : X \to Y$. If $\mathcal{F}$ is a sheaf on X, we have defined the pushforward or direct image sheaf $f_*\mathcal{F}$, which is a sheaf on Y. There is also a notion of inverse image sheaf. (We won't call it the pullback sheaf, reserving that name for a later construction, involving quasicoherent sheaves.) This is a covariant functor f^{-1} from sheaves on Y to sheaves on X. If the sheaves on Y have some additional structure (e.g. group or ring), then this structure is respected by f^{-1}.

1.1. *Definition by adjoint: elegant but abstract.* Here is a categorical definition of the inverse image: f^{-1} is left-adjoint to f_*.

This isn't really a definition; we need a construction to show that the adjoint exists. (Also, for pedants, this won't determine $f^{-1}\mathcal{F}$; it will only determine it up to unique isomorphism.) Note that we then get canonical maps $f^{-1}f_*\mathcal{F} \to \mathcal{F}$ (associated to the identity in $\operatorname{Mor}_Y(f_*\mathcal{F}, f_*\mathcal{F})$) and $\mathcal{G} \to f_*f^{-1}\mathcal{G}$ (associated to the identity in $\operatorname{Mor}_X(f^{-1}\mathcal{G}, f^{-1}\mathcal{G})$).

1.2. *Construction: concrete but ugly.* Define the temporary notation $f^{-1}\mathcal{G}^{\mathrm{pre}}(U) = \varinjlim_{V \supset f(U)} \mathcal{G}(V)$. (Recall the explicit description of direct limit: sections are sections on open sets containing $f(U)$, with an equivalence relation.)

Date: Monday, October 8, 2007. Updated Oct. 29, 2007. Minor update Nov. 17.

1.A. EXERCISE. Show that this defines a presheaf on X.

Now define the *inverse image of* $\mathcal{G}$ by $f^{-1}\mathcal{G} := (f^{-1}\mathcal{G}^{\mathrm{pre}})^{\mathrm{sh}}$.

You will show that this construction satisfies the universal property in Exercise 1.F. For the exercises before that, feel free to use either the adjoint description or the construction.

1.B. EXERCISE. Show that the stalks of $f^{-1}\mathcal{G}$ are the same as the stalks of $\mathcal{G}$. More precisely, if $f(x) = y$, describe a natural isomorphism $\mathcal{G}_y \cong (f^{-1}\mathcal{G})_x$. (Possible hint: use the concrete description of the stalk, as a direct limit. Recall that stalks are preserved by sheafification.)

1.C. EXERCISE (EASY BUT USEFUL). If U is an open subset of Y, $i : U \to Y$ is the inclusion, and $\mathcal{G}$ is a sheaf on Y, show that $i^{-1}\mathcal{G}$ is naturally isomorphic to $\mathcal{G}|_U$.

1.D. EXERCISE (EASY BUT USEFUL). If $y \in Y$, $i : \{y\} \to Y$ is the inclusion, and $\mathcal{G}$ is a sheaf on Y, show that $i^{-1}(\mathcal{G})$ is naturally isomorphic to the stalk $\mathcal{G}_y$.

1.E. EXERCISE. Show that f^{-1} is an exact functor from sheaves of abelian groups on Y to sheaves of abelian groups on X. (Hint: exactness can be checked on stalks, and by Exercise 1.B, the stalks are the same.) The identical argument will show that f^{-1} is an exact functor from $\mathcal{O}_Y$-modules (on Y) to $f^{-1}\mathcal{O}_Y$-modules (on X), but don't bother writing that down. (Remark for experts: f^{-1} is a left-adjoint, hence right-exact by abstract nonsense. The left-exactness is true for "less categorical" reasons.)

1.F. IMPORTANT EXERCISE: THE CONSTRUCTION SATISFIES THE UNIVERSAL PROPERTY. If $f : X \to Y$ is a continuous map, and $\mathcal{F}$ is a sheaf on X and $\mathcal{G}$ is a sheaf on Y, describe a bijection

$$\mathrm{Mor}_X(f^{-1}\mathcal{G}, \mathcal{F}) \leftrightarrow \mathrm{Mor}_Y(\mathcal{G}, f_*\mathcal{F}).$$

Observe that your bijection is "natural" in the sense of the definition of adjoints.

1.G. EXERCISE. (a) Suppose $Z \subset Y$ is a closed subset, and $i : Z \hookrightarrow Y$ is the inclusion. If $\mathcal{F}$ is a sheaf on Z, then show that the stalk $(i_*\mathcal{F})_y$ is a one element set if $y \notin Z$, and $\mathcal{F}_y$ if $y \in Z$.
(b) *Important definition:* Define the *support* of a sheaf $\mathcal{F}$ of sets, denoted $\mathrm{Supp}\,\mathcal{F}$, as the locus where the stalks are not a one-element set:

$$\mathrm{Supp}\,\mathcal{F} := \{x \in X : |\mathcal{F}_x| \neq 1\}.$$

(More generally, if the sheaf has value in some category, the support consists of points where the stalk is not the final object. For sheaves of abelian groups, the support consists of points with non-zero stalks.) Suppose $\mathrm{Supp}\,\mathcal{F} \subset Z$ where Z is closed. Show that the natural map $\mathcal{F} \to i_*i^{-1}\mathcal{F}$ is an isomorphism. Thus a sheaf supported on a closed subset can be considered a sheaf on that closed subset.

2. RECOVERING SHEAVES FROM A "SHEAF ON A BASE"

Sheaves are natural things to want to think about, but hard to get one's hands on. We like the identity and gluability axioms, but they make proving things trickier than for presheaves. We have discussed how we can understand sheaves using stalks. We now introduce a second way of getting a hold of sheaves, by introducing the notion of a *sheaf on a base*.

First, let me define the notion of a *base of a topology*. Suppose we have a topological space X, i.e. we know which subsets of X are open $\{U_i\}$. Then a base of a topology is a subcollection of the open sets $\{B_j\} \subset \{U_i\}$, such that each U_i is a union of the B_j. There is one example that you have seen early in your mathematical life. Suppose $X = \mathbb{R}^n$. Then the way the usual topology is often first defined is by defining *open balls* $B_r(x) = \{y \in \mathbb{R}^n : |y - x| < r\}$, and declaring that any union of open balls is open. So the balls form a base of the usual topology. Equivalently, we often say that they *generate* the usual topology. As an application of how we use them, to check continuity of some map $f : X \to \mathbb{R}^n$, you need only think about the pullback of balls on $\mathbb{R}^n$.

Now suppose we have a sheaf $\mathcal{F}$ on X, and a base $\{B_i\}$ on X. Then consider the information $(\{\mathcal{F}(B_i)\}, \{\mathrm{res}_{B_i,B_j} : \mathcal{F}(B_i) \to \mathcal{F}(B_j)\})$, which is a subset of the information contained in the sheaf — we are only paying attention to the information involving elements of the base, not all open sets.

We can recover the entire sheaf from this information. Reason: we can determine the stalks from this information, and we can determine when germs are compatible.

2.A. EXERCISE. Make this precise.

This suggests a notion, that of a *sheaf on a base*. A sheaf of sets (rings etc.) on a base $\{B_i\}$ is the following. For each B_i in the base, we have a set $\mathcal{F}(B_i)$. If $B_i \subset B_j$, we have maps $\mathrm{res}_{ji} : \mathcal{F}(B_j) \to \mathcal{F}(B_i)$. (Things called B are always assumed to be in the base.) If $B_i \subset B_j \subset B_k$, then $\mathrm{res}_{B_k,B_i} = \mathrm{res}_{B_j,B_i} \circ \mathrm{res}_{B_k,B_j}$. So far we have defined a *presheaf on a base*.

We also require *base identity*: If $B = \cup B_i$, then if $f, g \in \mathcal{F}(B)$ such that $\mathrm{res}_{B,B_i} f = \mathrm{res}_{B,B_i} g$ for all i, then $f = g$.

We require *base gluability* too: If $B = \cup B_i$, and we have $f_i \in \mathcal{F}(B_i)$ such that f_i agrees with f_j on any basic open set in $B_i \cap B_j$ (i.e. $\mathrm{res}_{B_i,B_k} f_i = \mathrm{res}_{B_j,B_k} f_j$ for all $B_k \subset B_i \cap B_j$) then there exist $f \in \mathcal{F}(B)$ such that $\mathrm{res}_{B,B_i} = f_i$ for all i.

2.1. *Theorem. — Suppose* $\{B_i\}$ *is a base on* X*, and* F *is a sheaf of sets on this base. Then there is a unique sheaf* $\mathcal{F}$ *extending* F *(with isomorphisms* $\mathcal{F}(B_i) \cong F(B_i)$ *agreeing with the restriction maps).*

Proof. We will define $\mathcal{F}$ as the sheaf of compatible germs of F.

Define the *stalk* of F at $x \in X$ by

$$F_x = \varinjlim F(B_i)$$

where the colimit is over all B_i (in the base) containing x.

We'll say a family of germs in an open set U is compatible near x if there is a section s of F over some B_i containing x such that the germs over B_i are precisely the germs of s. More formally, define

$$\mathcal{F}(U) := \{(f_x \in F_x)_{x \in U} : \forall x \in U, \exists B \text{ with } x \subset B \subset U, s \in F(B) : s_y = f_y \forall y \in B\}$$

where each B is in our base.

This is a sheaf (for the same reasons as the sheaf of compatible germs was earlier).

I next claim that if U is in our base, the natural map $F(B) \to \mathcal{F}(B)$ is an isomorphism.

2.B. TRICKY EXERCISE. Describe the inverse map $\mathcal{F}(B) \to F(B)$, and verify that it is indeed inverse. □

Thus sheaves on X can be recovered from their "restriction to a base". This is a statement about *objects* in a category, so we should hope for a similar statement about *morphisms*.

2.C. IMPORTANT EXERCISE: MORPHISMS OF SHEAVES CORRESPOND TO MORPHISMS OF SHEAF ON A BASE. Suppose $\{B_i\}$ is a base for the topology of X.
(a) Verify that a morphism of sheaves is determined by the induced morphism of sheaves on the base.
(b) Show that a morphism of sheaves on the base (i.e. such that the diagram

$$\begin{array}{ccc} F(B_i) & \longrightarrow & G(B_i) \\ \downarrow & & \downarrow \\ F(B_j) & \longrightarrow & G(B_j) \end{array}$$

commutes for all $B_j \hookrightarrow B_i$) gives a morphism of the induced sheaves.

2.D. IMPORTANT EXERCISE. Suppose $X = \cup U_i$ is an open cover of X, and we have sheaves $\mathcal{F}_i$ on U_i along with isomorphisms $\phi_{ij} : \mathcal{F}_i|_{U_i \cap U_j} \to \mathcal{F}_j|_{U_i \cap U_j}$ that agree on triple overlaps (i.e. $\phi_{ij} \circ \phi_{jk} = \phi_{ij}$ on $U_i \cap U_j \cap U_k$). Show that these sheaves can be glued together into a unique sheaf $\mathcal{F}$ on X, such that $\mathcal{F}_i = \mathcal{F}|_{U_i}$, and the isomorphisms over $U_i \cap U_j$ are the obvious ones. (Thus we can "glue sheaves together", using limited patching information.) (You can use the ideas of this section to solve this problem, but you don't necessarily need to. Hint: As the base, take those open sets contained in *some* U_i.)

2.2. *Remark for experts.* This almost says that the "set" of sheaves forms a sheaf itself, but not quite. Making this precise leads one to the notion of a *stack*.

We are now ready to consider the notion of a *scheme*, which is the type of geometric space considered by algebraic geometry. We should first think through what we mean by "geometric space". You have likely seen the notion of a manifold, and we wish to abstract this notion so that it can be generalized to other settings, notably so that we can deal with non-smooth and arithmetic objects.

The key insight behind this generalization is the following: we can understand a geometric space (such as a manifold) well by understanding the functions on this space. More precisely, we will understand it through the sheaf of functions on the space. If we are interested in differentiable manifolds, we will consider differentiable functions; if we are interested in smooth manifolds, we will consider smooth functions and so on.

Thus we will define a scheme to be the following data

- *The set:* the points of the scheme
- *The topology:* the open sets of the scheme
- *The structure sheaf:* the sheaf of "algebraic functions" (a sheaf of rings) on the scheme.

Recall that a topological space with a sheaf of rings is called a *ringed space*.

We will try to draw pictures throughout, so our geometric intuition can guide the algebra development (and, eventually, vice versa). Pictures can help develop geometric intuition. Some readers will find the pictures very helpful, while others will find the opposite.

3.1. Example: Differentiable manifolds. As motivation, we return to our example of differentible manifolds, reinterpreting them in this light. We will be quite informal in this section. Suppose X is a manifold. It is a topological space, and has a *sheaf of differentiable functions* $\mathcal{O}_X$ (as described earlier). This gives X the structure of a ringed space. We have observed that evaluation at p gives a surjective map from the stalk to $\mathbb{R}$

$$\mathcal{O}_{X,p} \twoheadrightarrow \mathbb{R},$$

so the kernel, the (germs of) functions vanishing at p, is a maximal ideal $\mathfrak{m}_X$.

We could *define* a differentiable real manifold as a topological space X with a sheaf of rings such that there is a cover of X by open sets such that on each open set the ringed space is isomorphic to a ball around the origin in $\mathbb{R}^n$ with the sheaf of differentiable functions on that ball. With this definition, the ball is the basic patch, and a general manifold is obtained by gluing these patches together. (Admittedly, a great deal of geometry comes from how one chooses to patch the balls together!) In the algebraic setting, the basic patch is the notion of an *affine scheme*, which we will discuss soon.

Functions are determined by their values at points. This is an obvious statement, but won't be true for schemes in general. We will see an example in Exercise 4.A(a).

Morphisms of manifolds. How can we describe differentiable maps of manifolds $X \to Y$? They are certainly continuous maps — but which ones? We can pull back functions along continuous maps. Differentiable functions pull back to differentiable functions. More formally, we have a map $f^{-1}\mathcal{O}_Y \to \mathcal{O}_X$. (The inverse image sheaf f^{-1} was defined in §1) Inverse image is left-adjoint to pushforward, so we get a map $f^\# : \mathcal{O}_Y \to f_*\mathcal{O}_X$.

Certainly given a differentiable map of manifolds, differentiable functions pullback to differentiable functions. It is less obvious that *this is a sufficient condition for a continuous function to be differentiable.*

3.A. IMPORTANT EXERCISE FOR THOSE WITH A LITTLE EXPERIENCE WITH MANIFOLDS. Prove that a continuous function of differentiable manifolds $f : X \to Y$ is differentiable if differentiable functions pull back to differentiable functions, i.e. if pullback by f gives a map $\mathcal{O}_Y \to f_*\mathcal{O}_X$. (Hint: check this on small patches. Once you figure out what you are trying to show, you'll realize that the result is immediate.)

3.B. EXERCISE. Show that a morphism of differentiable manifolds $f : X \to Y$ with $f(p) = q$ induces a morphism of stalks $f^\# : \mathcal{O}_{Y,q} \to \mathcal{O}_{X,p}$. Show that $f^\#(\mathfrak{m}_{Y,q}) \subset \mathfrak{m}_{X,p}$. In other words, if you pull back a function that vanishes at q, you get a function that vanishes at p — not a huge surprise.

Here is a little more for experts: Notice that this induces a map on tangent spaces

$$(\mathfrak{m}_{X,p}/\mathfrak{m}_{X,p}^2)^\vee \to (\mathfrak{m}_{Y,q}/\mathfrak{m}_{Y,q}^2)^\vee.$$

This is the tangent map you would geometrically expect. Again, it is interesting that the cotangent map $\mathfrak{m}_{Y,q}/\mathfrak{m}_{Y,q}^2 \to \mathfrak{m}_{X,p}/\mathfrak{m}_{X,p}^2$ is algebraically more natural than the tangent map.

Experts are now free to try to interpret other differential-geometric information using only the map of topological spaces and map of sheaves. For example: how can one check if f is a submersion? How can one check if f is an immersion? (We will see that the algebro-geometric version of these notions are *smooth morphisms* and *locally closed immersion*.)

3.2. *Side Remark.* Manifolds are covered by disks that are all isomorphic. Schemes (or even complex algebraic varieties) will not have isomorphic open sets. (We'll see an example later.) Informally, this is because in the topology on schemes, all non-empty open sets are "huge" and have more "structure".

4. THE UNDERLYING SET OF AFFINE SCHEMES

For any ring A, we are going to define something called $\operatorname{Spec} A$, the *spectrum of* A. In this section, we will define it as a set, but we will soon endow it with a topology, and later we will define a sheaf of rings on it (the structure sheaf). Such an object is called an *affine scheme*. In the future, $\operatorname{Spec} A$ will denote the set along with the topology. (Indeed, it will often implicitly include the data of the structure sheaf.) But for now, as there is no possibility of confusion, $\operatorname{Spec} A$ will just be the set.

The set $\operatorname{Spec} A$ is the set of prime ideals of A. The point of $\operatorname{Spec} A$ corresponding to the prime ideal $\mathfrak{p}$ will be denoted $[\mathfrak{p}]$.

We now give some examples. Here are some temporary definitions to help us understand these examples. Elements $a \in A$ will be called **functions** on $\operatorname{Spec} A$, and their **value** at the point $[\mathfrak{p}]$ will be $a \pmod{\mathfrak{p}}$. "An element a of the ring lying in a prime ideal $\mathfrak{p}$" translates to "a function a that is 0 at the point $[\mathfrak{p}]$" or "a function a vanishing at the point $[\mathfrak{p}]$", and we will use these phrases interchangeably. Notice that if you add or multiply two functions, you add or multiply their values at all points; this is a translation of the fact that $A \to A/\mathfrak{p}$ is a homomorphism of rings. These translations are important — make sure you are very comfortable with them!

Example 1: $\mathbb{A}^1_{\mathbb{C}} := \operatorname{Spec} \mathbb{C}[x]$. This is known as "the affine line" or "the affine line over $\mathbb{C}$". Let's find the prime ideals. As $\mathbb{C}[x]$ is an integral domain, 0 is prime. Also, $(x - a)$ is prime, where $a \in \mathbb{C}$: it is even a maximal ideal, as the quotient by this ideal is field:

$$0 \longrightarrow (x - a) \longrightarrow \mathbb{C}[x] \xrightarrow{f \mapsto f(a)} \mathbb{C} \longrightarrow 0$$

(This exact sequence should remind you of $0 \to \mathfrak{m}_x \to \mathcal{O}_x \to \mathbb{R} \to 0$ in our motivating example of manifolds.)

We now show that there are no other prime ideals. We use the fact that $\mathbb{C}[x]$ has a division algorithm, and is a unique factorization domain. Suppose $\mathfrak{p}$ is a prime ideal. If $\mathfrak{p} \neq 0$, then suppose $f(x) \in \mathfrak{p}$ is a non-zero element of smallest degree. It is not constant, as prime ideals can't contain 1. If $f(x)$ is not linear, then factor $f(x) = g(x)h(x)$, where $g(x)$ and $h(x)$ have positive degree. Then $g(x) \in \mathfrak{p}$ or $h(x) \in \mathfrak{p}$, contradicting the minimality of the degree of f. Hence there is a linear element $x - a$ of $\mathfrak{p}$. Then I claim that $\mathfrak{p} = (x - a)$. Suppose $f(x) \in \mathfrak{p}$. Then the division algorithm would give $f(x) = g(x)(x - a) + m$ where $m \in \mathbb{C}$. Then $m = f(x) - g(x)(x - a) \in \mathfrak{p}$. If $m \neq 0$, then $1 \in \mathfrak{p}$, giving a contradiction.

Thus we have a picture of $\operatorname{Spec} \mathbb{C}[x]$ (see Figure 1). There is one point for each complex number, plus one extra point. The point $[(x - a)]$ we will reasonably associate to $a \in \mathbb{C}$. Where should we picture the point $[(0)]$? Where is it? The best way of thinking about it is somewhat zen. It is somewhere on the complex line, but nowhere in particular. Because (0) is contained in all of these primes, we will somehow associate it with this line passing through all the other points. $[(0)]$ is called the "generic point" of the line; it is "generically on the line" but you can't pin it down any further than that. We'll place it far to the right for lack of anywhere better to put it. You will notice that we sketch $\mathbb{A}^1_{\mathbb{C}}$ as one-dimensional in the real sense; this is to later remind ourselves that this will be a one-dimensional space, where dimensions are defined in an algebraic (or complex-geometric) sense.

To give you some feeling for this space, let me make some statements that are currently undefined, but suggestive. The functions on $\mathbb{A}^1_{\mathbb{C}}$ are the polynomials. So $f(x) = x^2 - 3x + 1$ is a function. What is its value at $[(x - 1)]$, which we think of as the point $1 \in \mathbb{C}$? Answer: $f(1)$! Or equivalently, we can evalute $f(x)$ modulo $x - 1$ — this is the same thing by the division algorithm. (What is its value at (0)? It is $f(x) \pmod 0$, which is just $f(x)$.)

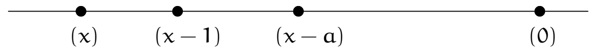

FIGURE 1. A picture of $\mathbb{A}^1_{\mathbb{C}} = \operatorname{Spec} \mathbb{C}[x]$

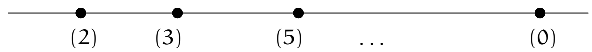

FIGURE 2. A "picture" of $\operatorname{Spec} \mathbb{Z}$, which looks suspiciously like Figure 1

Here is a more complicated example: $g(x) = (x-3)^3/(x-2)$ is a "rational function". It is defined everywhere but $x = 2$. (When we know what the structure sheaf is, we will be able to say that it is an element of the structure sheaf on the open set $\mathbb{A}^1_{\mathbb{C}} - \{2\}$.) $g(x)$ has a triple zero at 3, and a single pole at 2.

Example 2: $\mathbb{A}^1_k := \operatorname{Spec} k[x]$ *where* k *is an algebraically closed field.* This is called the affine line over k. All of our discussion in the previous example carries over without change. We will use the same picture, which is after all intended to just be a metaphor.

Example 3: $\operatorname{Spec} \mathbb{Z}$*.* One amazing fact is that from our perspective, this will look a lot like the affine line. This is another unique factorization domain, with a division algorithm. The prime ideals are: (0), and (p) where p is prime. Thus everything from Example 1 carries over without change, even the picture. Our picture of $\operatorname{Spec} \mathbb{Z}$ is shown in Figure 2.

Let's blithely carry over our discussion of functions on this space. 100 is a function on $\operatorname{Spec} \mathbb{Z}$. It's value at (3) is "$1 \pmod 3$". It's value at (2) is "$0 \pmod 2$", and in fact it has a double zero. $27/4$ is a rational function on $\operatorname{Spec} \mathbb{Z}$, defined away from (2). It has a double pole at (2), a triple zero at (3). Its value at (5) is

$$27 \times 4^{-1} \equiv 2 \times (-1) \equiv 3 \pmod 5.$$

Example 4: stupid examples. $\operatorname{Spec} k$ where k is any field is boring: only one point. $\operatorname{Spec} 0$, where 0 is the zero-ring, is the empty set, as 0 has no prime ideals.

4.A. A SMALL EXERCISE ABOUT SMALL SCHEMES. (a) Describe the set $\operatorname{Spec} k[\epsilon]/\epsilon^2$. This is called the ring of **dual numbers**, and will turn out to be quite useful. You should think of ϵ as a very small number, so small that its square is 0 (although it itself is not 0).
(b) Describe the set $\operatorname{Spec} k[x]_{(x)}$. (We will see this scheme again later.)

In Example 2, we restricted to the case of algebraically closed fields for a reason: things are more subtle if the field is not algebraically closed.

Example 5: $\mathbb{R}[x]$. Using the fact that $\mathbb{R}[x]$ is a unique factorization domain, we see that the primes are (0), $(x-a)$ where $a \in \mathbb{R}$, and (x^2+ax+b) where x^2+ax+b is an irreducible quadratic. The latter two are maximal ideals, i.e. their quotients are fields. For example: $\mathbb{R}[x]/(x-3) \cong \mathbb{R}$, $\mathbb{R}[x]/(x^2+1) \cong \mathbb{C}$.

4.B. UNIMPORTANT EXERCISE. Show that for the last type of prime, of the form $(x^2 + ax + b)$, the quotient is *always* isomorphic to $\mathbb{C}$.

So we have the points that we would normally expect to see on the real line, corresponding to real numbers; the generic point 0; and new points which we may interpret as *conjugate pairs* of complex numbers (the roots of the quadratic). This last type of point should be seen as more akin to the real numbers than to the generic point. You can picture $\mathbb{A}^1_{\mathbb{R}}$ as the complex plane, folded along the real axis. But the key point is that Galois-conjugate points are considered glued.

Let's explore functions on this space; consider the function $f(x) = x^3 - 1$. Its value at the point $[(x-2)]$ is $f(x) = 7$, or perhaps better, $7 \pmod{x-2}$. How about at (x^2+1)? We get

$$x^3 - 1 \equiv x - 1 \pmod{x^2+1},$$

which may be profitably interpreted as $i - 1$.

One moral of this example is that we can work over a non-algebraically closed field if we wish. It is more complicated, but we can recover much of the information we wanted.

4.C. EXERCISE. Describe the set $\mathbb{A}^1_{\mathbb{Q}}$. (This is harder to picture in a way analogous to $\mathbb{A}^1_{\mathbb{R}}$; but the rough cartoon of points on a line, as in Figure 1, remains a reasonable sketch.)

Example 6: $\mathbb{F}_p[x]$. As in the previous examples, this has a division algorithm, so the prime ideals are of the form (0) or $(f(x))$ where $f(x) \in \mathbb{F}_p[x]$ is an irreducible polynomials, which can be of any degree. Irreducible polynomials correspond to sets of Galois conjugates in $\overline{\mathbb{F}}_p$.

Note that $\operatorname{Spec} \mathbb{F}_p[x]$ has p points corresponding to the elements of $\mathbb{F}_p$, but also (infinitely) many more. This makes this space much richer than simply p points. For example, a polynomial $f(x)$ is not determined by its values at the p elements of $\mathbb{F}_p$, but it *is* determined by its values at the points of $\operatorname{Spec} \mathbb{F}_p$. (As we have mentioned before, this is not true for all schemes.)

You should think about this, even if you are a geometric person — this intuition will later turn up in geometric situations. Even if you think you are interested only in working over an algebraically closed field (such as $\mathbb{C}$), you will have non-algebraically closed fields (such as $\mathbb{C}(x)$) forced upon you.

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