A Short Introduction to Schemes

Brian Lawrence

October 28, 2013

Abstract

A while back a topologist friend asked me why the set of prime ideals of a ring deserved to be singled out for study. Another friend, a symplectic geometer, asked why anyone would want to put a scheme structure on the paramaters that classify a type of object (i.e. construct a moduli space). Algebraic geometry does have a reputation for being somewhat impenetrable (and I won't say it's undeserved), but there are natural reasons to ask these questions. In this talk I introduce schemes and show how they relate to classical varieties. In particular, I discuss the value of allowing rings with nilpotents, one of the major innovations of scheme theory.

These are notes from a talk given at Stanford's KIDDIE Colloquium for first-year grads, Oct. 28, 2013. (I do not know whether KIDDIE stands for anything.)

1 From Varieties to Schemes

Schemes are a generalization of classical algebraic varieties. My goal today is show how basic geometric intuition can be translated into the language of schemes (so nothing is lost in going from varieties to schemes) and how the generality of scheme theory allows one to introduce useful new techniques (so something is gained).

Since this talk is an intuitive introduction to schemes, I will feel free to give impressionistic definitions and state imprecise theorems.

We begin by reviewing the classical notion of a variety. A variety is, loosely speaking, a sort of "algebraic manifold," but note that varieties need not be smooth.

"Definition". Given a field K and polynomials $f_1, \ldots, f_k$ in $K[x_1, \ldots, x_n]$, define

$$V(f_1, \ldots, f_k) = \{x \in K^n | f_i(x) = 0, 1 \leq i \leq k\}.$$

An *affine variety* over K is a set of the form $V(f_1, \ldots, f_k)$ for some n and $f_1, \ldots, f_k$. A morphism of affine varieties is a map given by polynomials in the coefficients. A *variety* is something that can be covered by finitely many affine varieties. A morphism of varieties is a map that, when restricted to affine varieties, gives a morphism of affine varieties.

Examples. If $k = 0$ (no equations) we recover the set K^n, called *affine n-space.*

In the plane, the equation $x_2 = 0$ cuts out a line. The equation $x_2^2 = 0$ cuts out the same line. There is no notion of "multiplicity."

The circle $x^2 + y^2 = 1$ and the line $y = 1$ intersect in a single point. Again, no multiplicity.

The equation $xy = 0$ cuts out a variety which is the union of two lines and has a singular point at the origin (whatever that means). The equation $y^2 = x^2(x-1)$ cuts out a variety that also has a singular point. Even over a field like $\mathbb{R}$ or $\mathbb{C}$ these varieties are not manifolds.

One defines n-dimensional projective space by gluing affine spaces appropriately. It is possible to describe subvarieties of projective space as subsets cut out by equations, loosely speaking.

Now we need to make a subtle shift in perspective. Instead of the set of points of an affine variety, we focus on the ring of polynomial functions on it.

Examples. The ring of polynomial functions on K^n is $K[x_1, \ldots, x_n]$.

The ring of polynomial functions on the circle $x^2 + y^2 = 1$ (defined over $\mathbb{C}$, say) is $\mathbb{C}[x, y]/(x^2 + y^2 - 1)$.

In many cases the ring of polynomials on the variety in K^n cut out by $f_1, \ldots, f_k$ is exactly the ring $K[x_1, \ldots, x_n]/(f_1, \ldots, f_k)$.

However, the ring of polynomials on the variety in K^2 cut out by $y^2 = 0$ is $K[x, y]/(y)$, not $K[x, y]/(y^2)$.

In scheme theory, the ring of functions is fundamental; the set of points and topological structure are recovered from the ring.

"Definition". An *affine scheme*, denoted Spec A, is the data of a ring A. A morphism of affine schemes Spec $A \to$ Spec B, is a morphism of rings $B \to A$, with the arrow pointing in the opposite direction. An affine scheme *over a field k* is a scheme Spec A where A is equipped with a k-algebra structure. A morphism of affine schemes over k is a ring morphism respecting the k-algebra structure.

Why is a morphism defined this way? Given a map of classical varieties, functions defined on the target pull back to the source. In scheme theory we take this as the definition of a morphism.

Where did the points go? The points of a classical variety can be recovered as the maximal ideals in its ring of functions. This is the content of the following theorem, a form of Hilbert's Nullstellensatz.

Theorem. *Every maximal ideal of the ring $A = K[x_1, \ldots, x_n]/(f_1, \ldots, f_k)$ is of the form $(x_1 - a_1, \ldots, x_n - a_n)$, where $(a_1, \ldots, a_n)$ is a common zero of the polynomials f_i. This gives a bijective correspondence between maximal ideals of A and points of the classical variety.*

This suggests the definition: a point of an affine scheme is a maximal ideal of its ring. This would not be the right definition, as we will see soon, but it is a step in the right direction.

A morphism of affine schemes should give a map of points. How do we construct that map? In other words, given a ring map $B \to A$, what objects ("points") on A pull back to the same kind of object on B? The following result is easy.

Lemma. *Given a map of rings (commutative with identity) $f : B \to A$, for every prime ideal p of A, the set $f^{-1}(p)$ is a prime ideal of B.*

It is not true, in general, than the inverse image of a maximal ideal is maximal. (But see below.) Hence we are led to define points as prime, not maximal, ideals.

Definition. A *point* of an affine scheme Spec A is a prime ideal p of A. Hence, a morphism of affine schemes induces a (set-theoretic) map on points, via the above Lemma.

A point corresponding to a maximal ideal is sometimes called a "classical" or "closed" point. For most purposes it is enough to consider only classical points. But we will see that occasionally it is useful to also work with generic points.

Theorem. *Under some mild hypotheses, a map of affine schemes takes classical points to classical points.*

Aside: the point set of an affine scheme comes with a topological structure, called the Zariski topology. Also, it is possible to define a "scheme" in general, by gluing affine schemes, and to give a definition for a "morphism" of schemes. Thus, for example, we can define projective space $\mathbb{P}^n_R$ as a scheme.

We have seen that schemes allow us to recover at least the point set of a classical variety. At this point it should not come as a surprise that any of the classical algebro-geometric constructions can be carried over into schemes. But why bother with schemes? Schemes generalize varieties in several ways. First, schemes have generic (non-classical) points, as we have seen. Second, schemes can be defined over arbitrary rings: Spec $\mathbb{Z}$ is a scheme, but it is not a variety over any field. Third, schemes can have nilpotent elements: Spec $k[x, y]/(y)$ and Spec $k[x, y]/(y^2)$ are nonisomorphic affine schemes (since their rings are not isomorphic), but there is no variety corresponding to Spec $k[x, y]/(y^2)$. For the remainder of this talk we will discuss some applications of these three generalizations.

2 A Word about Generic Points

Let A be the ring $\mathbb{C}[x, y]$. As we have seen, Spec A is supposed to be a scheme-theoretic analogue of the complex plane $\mathbb{C}^2$. (It is traditional in algebraic geometry to use complex, not real, dimension. So a "curve" over $\mathbb{C}$ is a scheme of complex dimension 1 and real dimension 2; a "surface" is of complex dimension 2 and real dimension 4, and so forth. The one exception is the phrase "Riemann surface," so called for historical reasons. Thus a Riemann surface is a curve, not a surface. In this talk I will draw $\mathbb{C}^2$ as a plane $\mathbb{R}^2$, and curves in $\mathbb{C}^2$ as real curves in the real plane.)

What are the prime ideals of A? We have seen that the maximal ideals are exactly those of the form $(x-a, y-b)$; these correspond to points of the plane. Any irreducible polynomial generates a prime ideal: so, for instance, (y) and $(y-x^2)$ are prime ideals. Such an ideal naturally corresponds to its vanishing locus, which is a curve in the plane. Finally, since A is an integral domain, the zero ideal (0) is also a prime ideal. This corresponds to the whole plane. One can show that these are all the prime ideals of $\mathbb{C}[x, y]$.

We can visualize the generic point associated to a prime ideal, say $(y-x^2)$, as being located somewhere on the curve, but nowhere in particular. Does $x = 0$ at this generic point? No. Algebraically, evaluation at a point is given by the quotient map $A \to A/(y-x^2)$, and x does not map to zero. Geometrically, x is nonzero at a "general" point of the curve, so we should say x is nonzero at the generic point. (This is probably why they are called generic points.) On the other hand, $y = x^2$ at the generic point, and at any closed point as well, of this curve.

Topologically, the closure of this generic point contains the point itself plus every closed point of the curve.

Similarly, the generic point associated to (0) is somewhere in the plane, but not at any particular point and not on any particular curve. Its closure is the plane itself.

The following argument gives a typical application of generic points. (It is not a complete proof – I've just included it to give a vague idea of how such a proof might go.)

"Theorem". *A reasonably nice affine scheme X over a field, whose ring has no nilpotent elements, is smooth on a dense open set.*

Proof. The set of smooth points of X is open, essentially by the Jacobi criterion. Let X_i be the irreducible components of X, and let p_i be the generic point of X_i. Using the fact that the ring has no nilpotent elements, one shows that X_i is smooth at p_i. But then the smooth locus of X_i, being an open set containing the generic point, must be dense; and patching the X_i's together one finds that the smooth locus of X is itself dense. □

3 A Word about Arbitrary Rings

It is useful in number theory to view a scheme "over $\mathbb{Z}$" as a family of schemes, one over each of the finite fields $\mathbb{F}_p$ and one over $\mathbb{Q}$. The scheme-theoretic picture provides geometric intuition, and many scheme-theoretic results, initially of a geometric nature, have arithmetic meaning in this context. Due to time constraints I will not discuss this.

4 Some Applications of Nilpotents

Let X be the complex plane Spec $A = \text{Spec } \mathbb{C}[x, y]$. As we have seen, for any point $p = (a, b)$ of X, evaluation at p gives a map $A \to \mathbb{C}$; this map is exactly the quotient map of A by the maximal ideal $(x - a, y - b)$ of X.

Now suppose we have a tangent vector v to the plane at p. For any $f \in A$, we can ask for the derivative of f along v at p; the result is a complex number. Thus we get some sort of map $A \to \mathbb{C}$. But it's not a map of $\mathbb{C}$-algebras, or even a map of rings; instead it satisfies a Leibniz rule. This isn't very convenient.

The following algebraic procedure encodes the same information in a more elegant way. Given f, write formally

$$f(p + \epsilon v) = f(p) + D_v(f)\epsilon + O(\epsilon^2).$$

Then the Leibniz rule for differentiation tells us exactly that

$$(fg)(p) + D_v(fg)\epsilon = (f(p) + D_v(f)\epsilon)\,(f(p) + D_v(f)\epsilon) + O(\epsilon^2).$$

In other words, the data of p and v give a map of rings

$$A \to \mathbb{C}[\epsilon]/\epsilon^2.$$

The value of f at p is given by the constant term in the image; the derivative is given by the ϵ term. In other words, composing with the map

$$\mathbb{C}[\epsilon]/\epsilon^2 \to \mathbb{C}[\epsilon]/\epsilon = \mathbb{C}$$

gives evaluation at p.

But we said earlier that a map of rings is by definition a morphism of affine schemes! The first map, evaluation at p, gives a map from Spec $\mathbb{C}$ into X. Clearly Spec $\mathbb{C}$ is a single point, it gets mapped to the point p of X, everything tautologically works and there is nothing more to say. What's going on with the second map? We can see that $\Lambda = \text{Spec } \mathbb{C}[\epsilon]/\epsilon^2$ is again a single point, and as expected it also maps to the point p of X. But the map carries more information than that: it also somehow "remembers" the vector v. So it's probably better to thing of Λ as somehow coming with a tangent vector, even though Λ is just a single point. To specify a map of Λ into a scheme S we have to specify both the point p that is the image and the tangent vector to p in S which is the image of $d/d\epsilon$, so to speak.

The composition of ring maps above can be seen geometrically as follows: the point Spec $\mathbb{C}$ maps into the point with tangent vector Λ, which maps into X.

Of course we could construct even "larger" points. A map from Spec $\mathbb{C}[\epsilon]/\epsilon^3$ would carry second-order information as well as a tangent vector; a map from Spec $\mathbb{C}[\epsilon]/\epsilon^n$ would carry even higher-order information. Adjoining nilpotents ϵ_2, ϵ_3, and so forth allows for a multidimensional tangent space. The rings whose spectra we are considering here are called Artin local rings and they are very useful in studying the local properties of schemes, as we will see.

We say a ring is "reduced" if it has no nonzero nilpotents; we also say that the corresponding scheme is reduced.

4.1 Multiplicity and Intersection Theory

Let's go back to the question of multiplicity, which arose when we discussed classical varieties.

In the plane, the equation $y^2 = 0$ cuts out a non-reduced line, while the equation $y = 0$ cuts out a reduced line. Contrast this with the situation of varieties, where both equations cut out the same variety.

Similarly, the scheme-theoretic intersection of the line $y = 1$ with the circle $x^2+y^2 = 1$ is given by the two equations $y = 1$ and $x^2 = 0$. This scheme is a non-reduced point. (In fact, it is isomorphic to the scheme Λ above.) Furthermore, if we define the degree of an affine scheme over k to be the dimension of its coordinate ring as a k-vector space, then the intersection has degree 2. In fact, the intersection of any line with any conic (in the projective plane) has degree 2. Thus non-reduced schemes give an elegant definition of multiplicity, in which Bezout's Theorem holds.

Theorem. *Suppose C_1 and C_2 are two curves in the projective plane with no common component, defined by equations of degrees d_1 and d_2, respectively. Then their intersection is a scheme of degree $d_1 d_2$.*

Similarly, consider n homogenous polynomials on $\mathbb{P}^n$ over a field, of degrees $d_1, \ldots, d_n$. If the scheme they cut out has no components of dimension greater than zero, then its degree is $d_1 \cdots d_n$.

4.2 Deformation Theory

In algebraic geometry as elsewhere, there are some natural classification questions which arise. For instance: Classify elliptic curves up to isomorphism. Classify curves of genus g and degree d in $\mathbb{P}^3$. Classify line bundles L on a given elliptic curve E.

What sort of answer should we hope for? Consider the problem of classifying elliptic curves. Over an algebraically closed field k, elliptic curves are classified by their j-invariant. In other words, every elliptic curve has a well-defined j-invariant; two elliptic curves have the same j-invariant if and only if they are isomorphic; and one knows exactly which values the j-invariant can take on. So, we're done, right? This seems like a satisfactory result.

But now consider the following setup. Suppose we have a map of schemes $X \to S$, such that the fiber over every point of S is an elliptic curve. (Feel free to assume the map is flat – a technical condition that we won't worry about here. Don't worry about generic points either, but it's also reasonable to request that the fiber over each generic point be an elliptic curve as well.) For every $s \in S$ we can compute the j-invariant of the elliptic curve over s, so we get a function on the points of S. (Specifically, every closed point of S gets mapped to some value in k.) Now we ask, is this function algebraic? In other words, does it come from a map of schemes from S to the j-line? It seems reasonable to hope for this, but it certainly does not follow from the "pointwise" classification result cited above.

Thus we need a stronger result. We would like to have a scheme structure on the j-line, so every (flat) family of elliptic curves over S gives a scheme-theoretic map to the j-line. Ideally, we would also like to have a universal elliptic curve E_{univ} over the j-line J, so that any family of elliptic curves over S is the pullback of E_{univ} through a unique map $S \to J$. In fact for elliptic curves this is not possible: if E is an elliptic curve with nontrivial automorphisms, then it is possible to produce a map $X \to S$ such that each fiber is isomorphic to E, but (loosely speaking) "going around a loop" on S takes the fiber through an automorphism. That is, there are nontrivial E-bundles on S.

This issue doesn't arise for subschemes of a given scheme, or for vector bundles on a given scheme – in these cases, there is always a well-defined moduli space. Even when there is not a moduli space with all the functorial properties we want it's still possible to weaken some hypotheses and make inroads on the problem. But this is beyond the scope of today's talk.

So suppose for simplicity that we're in a situation where there is a well-defined moduli space. For instance, suppose we want to classify degree-zero line bundles on a curve C of genus g. (I have not defined line bundles.) It is well-known that a moduli space exists. What is its dimension? Is it smooth?

The moduli space $\mathrm{Pic}^0 C$ satisfies a property which should be seen as analogous to the desired universal property above – there is a classifying space $\mathrm{Pic}^0 C$, and then any family of degree-zero line bundles on C, parametrized by S, induces a map $S \to \mathrm{Pic}^0 C$. Somewhat more precisely: there is a universal line bundle P on $\mathrm{Pic}^0 C \times C$, and for any scheme S and any (reasonable) line bundle L on $S \times C$ (which should be thought of as a "line bundle on each fiber C_s, for $s \in S$) there is a unique map $S \to \mathrm{Pic}^0 C$ such that L is more-or-less isomorphic to the pull-back of P to $S \times C$.

The point I would like to make is the following. Take Λ be the spectrum of the Artin local ring $k[\epsilon]/\epsilon^2$, as before. Then "line bundles on C over Λ" (I am being deliberately imprecise) are in bijection with maps from Λ to $\mathrm{Pic}^0 C$. But such maps correspond to tangent vectors in $\mathrm{Pic}^0 C$, so by using nonreduced schemes we can get our hands on the tangent space to $\mathrm{Pic}^0 C$ at any point.

Specifically, a point p of $\mathrm{Pic}^0 C$ corresponds to a degree-zero line bundle L on C. A tangent vector at p corresponds to a line bundle over Λ, which pulls back to L over the reduced point in Λ. But (if you know all the machinery) it is very easy to classify these. They are classified (more or less) by the cohomology group $H^1(O_C)$. It does not matter what this means – it is a fundamental quantity which is known to have dimension g as a vector space over k. In short, we find by an easy calculation that the tangent space to $\mathrm{Pic}^0 C$ at any point has dimension g. We can do this sort of calculation in general, not just for this particular problem; thus it is in general relatively easy to find the dimension of the tangent space to any moduli space.

Our problem is not yet solved. Schemes can have singular points, and at those points the tangent space as we have defined it has higher dimension than the dimension of the variety itself. For example, the scheme cut out by $xy = 0$ in the plane has a tangent space of dimension 2 at the origin. Worse, a nonreduced scheme can have extra dimensions in its tangent space at every point. The tangent space to the scheme Λ has dimension 1; the tangent space to the double line at any point has dimension 2. In the case of our moduli spaces, we need to determine whether all the tangent directions we have found can be extended into the space $\mathrm{Pic}^0 C$.

In general one might attack this problem by classifying line bundles over a larger base than Λ, say $\Lambda_n = \mathrm{Spec}\ k[\epsilon]/\epsilon^n$, for $n = 3, 4, \ldots$. Hopefully, we can determine which first-order deformations (i.e. objects over Λ) extend to second-order deformations (over Λ_3), and which of these extend further, and so on. In the limit, we would have objects over a power series ring, and we might hope to get information about the moduli space in the large. This is the subject of deformation theory.

In our particular case one can show by other means that $\mathrm{Pic}^0 C$ must have dimension g. Since the tangent space at any point has dimension g, we deduce that $\mathrm{Pic}^0 C$ is smooth at every point.

In fact, building up from Artin local rings is a common method of proof. Typically, one wishes to prove some result globally on a variety or scheme. One shows that it is true for Artin local rings, perhaps by induction on the dimension, from the case of a field or the ring Λ. Taking a limit one obtains the result in some "formal" sense, over a power series ring. For many results the global result follows from the result over a power series ring. (For example, a function on an irreducible variety whose power series expansion at a point is 0 must be identically 0.)

5 References

Ravi's notes are an excellent exposition of scheme theory. In particular my discussion of schemes and their points drew from Chapter 3 and Section 4.2 . Bezout's Theorem in the plane is Exercise 18.6K; the generalization to projective n-space follows from the discussion in Section 20.1.

On a less elementary level:

Hartshorne's book *Deformation Theory* is a detailed introduction to the subject named in its title.

Mumford's *Lectures on Curves on an Algebraic Surface* has a construction of the Picard scheme $\mathrm{Pic}^0 X$.

The introduction to Grothendieck's EGA 1 has an exposition of the technique of proof mentioned at the end of the talk. See the section from pages 8 to 9, starting from the second paragraph: "La technique generale...."

Mumford's *Abelian Varieties* has some proofs that run along those lines, but are written more concretely. (For example, there is no mention of "faithfully flat descent.") The scheme-theoretic proof of the Theorem of the Cube in Chapter 10 and the proof of the Lemma on irreducible fibers in Chapter 18 are good examples.