Hilbert schemes

Avi Zeff

1. Hilbert schemes: introduction

We'll generally work in the setting of varieties over a fixed field k, although most results hold more broadly; whenever convenient we'll additionally assume that k is algebraically closed, and we'll really have $\mathbb{C}$ in mind. Let X be such a variety. We want to study (closed) subschemes of X. To apply the machinery of algebraic geometry, we need to think of this as a functor whose k-points are closed subschemes of X; we do this by associating to a test scheme T the set of T-families of subschemes of X, i.e. subschemes of $T \times_k X$ which are flat over T. Let's call this functor $\mathrm{Hilb}(X)$, i.e.

$$\mathrm{Hilb}(X)(T) = \{Z \subset T \times_k X | Z \text{ flat over } T\}.$$

(The action on morphisms sends $f : S \to T$ to $\mathrm{Hilb}(X)(f)$ sending $Z \subset T \times_k X$ to $(f \times \mathrm{id})^{-1}(Z) \subset S \times_k X$, so this is a contravariant functor.) For example, for $X = \mathbb{P}^n_k$ and a k-scheme T we have $T \times_k \mathbb{P}^n_k = \mathbb{P}^n_T$, so $\mathrm{Hilb}(\mathbb{P}^n_k)$ classifies closed subschemes of projective space.

Given a contravariant functor F from a category of schemes to sets, the usual question is whether this functor is representable: does there exist a scheme Y such that $\mathrm{Hom}(-, Y)$ is equivalent to F?

In our case where $F = \mathrm{Hilb}(X)$, it's not hard to see that no reasonable representing scheme can exist in general. For example, if $X = \mathbb{P}^1$, there is a surjection sending a subscheme of $\mathbb{P}^1$ to its support, which is either a finite set of points of $\mathbb{P}^1$ or all of $\mathbb{P}^1$. For each n the scheme classifying n-tuples of points in $\mathbb{P}^1$ is $\mathrm{Sym}^n \mathbb{P}^1 \simeq \mathbb{P}^n$, and so the scheme representing $\mathrm{Hilb}(\mathbb{P}^1)$ would have to surject onto a disjoint union of copies of each $\mathbb{P}^n$ (together with a point). This is infinite-dimensional with infinitely many connected components. More generally we should expect this functor to in general be "too big" to be representable.

We therefore want to pare it down. To see how to do this, let's look at the special case $X = \mathbb{P}^n$. Here, each subscheme $Z \subset \mathbb{P}^n$ has a finitely generated graded coordinate ring $R = \bigoplus_{i \geq 0} R_i$. Each R_i is a finite-dimensional k-vector space, and so we can define a function $i \mapsto \dim_k R_i$. For i sufficiently large this is a polynomial in i with integer values (but not necessarily coefficients), called the Hilbert polynomial of Z.

More generally, for a quasi-coherent sheaf $\mathcal{F}$ on a variety X over k we can consider the function

$$i \mapsto \dim_k H^0(X, \mathcal{F}(i))$$

where $\mathcal{F}(i)$ is the Serre twist. By the Serre vanishing theorem the higher cohomology of $\mathcal{F}(i)$ vanishes for sufficiently large i, and so the above function is equal to

$$i \mapsto \chi(\mathcal{F}(i))$$

for i sufficiently large; this latter turns out to be a polynomial in i, which we call the Hilbert polynomial of $\mathcal{F}$. In particular, a subscheme Z of X gives rise to a quasi-coherent ideal sheaf $\mathcal{O}_Z$ (or more accurately its pushforward along the inclusion), by which we can define the Hilbert polynomial p_Z of Z.

For example, suppose that Z is a zero-dimensional subscheme of a projective variety X, so that its function ring has finite dimension n. We say that n is the length of Z, and that the dimension of the local ring at each point p in the support of Z is the length of Z at p. In the special case we considered above where $X = \mathbb{P}^1$, each n gives a component $\mathbb{P}^n$; the remaining component, which for each T gives a unique point corresponding to $\mathbb{P}^n_T$, classifies the subschemes with Hilbert polynomial

$$i \mapsto \chi(\mathcal{O}_{\mathbb{P}^1}(i)) = i + 1.$$

In other words the Hilbert functor decomposes, in this case, into components corresponding to the Hilbert polynomials.

This suggests the following definition. For each integer-valued polynomial f, set

$$\mathrm{Hilb}^f(X)(T) = \{Z \subset T \times_k X | Z \text{ flat over } T, \forall t \in T : p_{Z_t} = f\},$$

where Z_t denotes the fiber of Z over $t \in T$. It turns out that this is the right definition: for X projective, these functors will be representable by projective (and thus finite type) schemes over k.

(Remark: there's an implicit use here of a lemma stating that the Hilbert polynomial is locally constant in (flat) families, and that we can reduce to the connected case; we'll ignore this sort of thing. Note also that although the Hilbert scheme is independent of the embedding of X into projective space, the Hilbert polynomial is not since it depends on a choice of $\mathcal{O}_X(1)$, and so a different choice will relabel the components.)

We will mostly be concerned with the simplest case of zero-dimensional subschemes, for which the Hilbert polynomial is constant. For each integer n, we abbreviate

$$X^{[n]} = \mathrm{Hilb}^n(X).$$

These can be equivalently thought of as ideals I of $\mathcal{O}_X$ such that $\dim \mathcal{O}_X/I = n$. Let's examine some cases.

First, suppose that $n = 0$. This is the most trivial situation: the only subscheme with Hilbert polynomial identically zero is the empty subscheme, so $\mathrm{Hilb}^n(X)$ is trivially represented by a point. Henceforth we assume $n > 0$.

If $n = 1$, things are mildly more interesting, but not much: length 1 subschemes correspond to ideals I with $\dim \mathcal{O}_X/I = 1$, i.e. maximal ideals, or closed points of X. Thus $X^{[1]}$ is just the scheme classifying points of X, namely X itself. We can also see this more formally as functors: $X^{[1]}(T)$ is the set of subvarieties Z of $T \times_k X$ such that each Z_t is a point of X, and therefore defines a map $T \to X$ sending $t \mapsto Z_t$. (The condition that Z be flat over T is essentially just the condition that this map be a morphism.) Therefore as functors $X^{[1]}(T) = \mathrm{Hom}(T, X) = X(T)$.

The case $n = 2$ is already interesting enough to be difficult to deal with. A point in $X^{[2]}$ is an ideal I of $\mathcal{O}_X$ such that $\mathcal{O}_X/I$ has length 2. There are two possibilities for this quotient: either it (equivalently the corresponding subscheme) is supported at two distinct points x_1, x_2, or it is supported at only one point. In the first case, at each point the length is 1, and so as in the case $n = 1$ we're just classifying points of X. Since we require the points to be distinct and don't care about their order, it follows that this part of $X^{[2]}$ looks like

$(X \times X - \Delta)/S_2$, where Δ is the diagonal and S_2 is the symmetric group on two elements. If $\mathcal{O}_X/I$ is supported at a unique point x, so that we have inclusions $\mathfrak{m}_x^2 \subseteq I \subset \mathfrak{m}_x \subset \mathcal{O}_X$. Such ideals are in natural bijection with maps $\phi : \mathfrak{m}_x/\mathfrak{m}_x^2 \to \mathcal{O}_X/\mathfrak{m}_x$ up to scaling; in other words this part of $X^{[2]}$ looks like $\mathbb{P}(\operatorname{Tan} X)$. Thus $X^{[2]} = (X \times X - \Delta)/S_2 \sqcup \mathbb{P}(\operatorname{Tan} X)$.

For general n, we can do a similar decomposition. We have a morphism $\pi : X^{[n]} \to \operatorname{Sym}^n X$, sending Z to its support in X, with each point p counted by multiplicity given by the length of Z at p. This is called the Hilbert-Chow morphism. Although it is not an isomorphism in general, it is once restricted to subschemes Z which are supported at n distinct points; this is the big open piece of $X^{[n]}$, given in the case $n = 2$ by $(X \times X - \Delta)/S_2$. The opposite end of the spectrum is the case where Z is supported at a single point of X, which in the case $n = 2$ is given by $\mathbb{P}(\operatorname{Tan} X)$. More broadly, we have a stratification of $\operatorname{Sym}^n X$ by partitions $\lambda = (\lambda_1, \lambda_2, \ldots)$ of n,

$$\operatorname{Sym}^n X = \bigsqcup_{\lambda} \operatorname{Sym}^\lambda X,$$

where $\operatorname{Sym}^\lambda X$ classifies unordered n-tuples of points $(x_1, \ldots, x_n)$ where $x_1 = x_2 = \cdots = x_{\lambda_1}$, $x_{\lambda_1+1} = \cdots = x_{\lambda_1+\lambda_2}$, and so on. This stratification pulls back to $X^{[n]}$ via $X_\lambda^{[n]} = \pi^{-1}(\operatorname{Sym}^\lambda X)$.

In particular, for $\lambda = (n)$, $X_\lambda^{[n]}$, the small stratum consisting of subschemes supported at a single point of X, fibers over X through the Hilbert-Chow morphism $\pi : X_\lambda^{[n]} \to \operatorname{Sym}^\lambda X = \operatorname{Sym}^1 X = X$. For each $p \in X$, write $X_p^{[n]}$ for the fiber over p, i.e. the scheme classifying subschemes of X of length n supported at p, or equivalently ideal sheaves I of $\mathcal{O}_X$ such that $\mathcal{O}_X/I$ is supported at p with length n.

For X a smooth curve, the Hilbert-Chow morphism $\pi : X^{[n]} \to \operatorname{Sym}^n X$ is an isomorphism: the local ring at each point p is generated by a uniformizer, and so the only information that the ideal at p contains is the length at p. Thus each subscheme is fully determined by its support.

This is generally speaking the only situation where we should expect this to hold. In general, the Hilbert scheme is not even smooth, and can be arbitrarily horrible. In the case of smooth surfaces, though, it will turn out that the Hilbert scheme actually is smooth; let's say something about this case.

2. Hilbert schemes of smooth surfaces

2.1 Dimension and smoothness

Let X be a smooth surface, and Z be a closed subscheme corresponding to an ideal I of length n. Suppose that Z is supported at a single point p, so that it corresponds to a point of $X_p^{[n]}$. As in the case $n = 2$ we examined before, we have $\dim \mathcal{O}_{X,p}/I = n$ and therefore (after completing at p) $\mathfrak{m}_p^n \subset I \subset \mathfrak{m}_p^{n-1} \subset \mathcal{O}_{X,p}$. In particular I is completely determined by $\bar{I} = I/\mathfrak{m}_p^n \subset \mathcal{O}_{X,p}/\mathfrak{m}_p^n$. Therefore the problem of studying $X_p^{[n]}$ reduces to that of studying ideals of $\mathcal{O}_{X,p}/\mathfrak{m}_p^n$. But the (completed) local ring and its maximal ideal are independent of the point p and indeed of X by smoothness: abstractly, $\mathcal{O}_{X,p} \simeq k[[T_1, T_2]]$. Therefore we can just pick our favorite smooth surface and study the local Hilbert scheme there; we'll take the affine plane $\mathbb{A}^2$, and set p to be the origin 0.

For $n = 2$, we already computed that the fibers of $X^{[2]}$ supported at a single point p are given by the projectivization of the tangent space at p. In particular for $\mathbb{A}^2$ it follows that $(\mathbb{A}^2)_0^{[2]} \simeq \mathbb{P}^1$. In general, it turns out that $(\mathbb{A}^2)_0^{[n]}$ is irreducible of dimension $n-1$, and therefore the same holds for any smooth surface and point.

Assuming this, let's think about the dimension of the whole Hilbert scheme $(\mathbb{A}^2)^{[n]}$, or more generally $X^{[n]}$ for a smooth surface X. At distinct points we can choose components independently, so if $\lambda = (\lambda_1, \ldots, \lambda_r)$ is a partition of n and $x = (x_1, \ldots, x_1, x_2, \ldots, x_2, x_3, \ldots, x_r, \ldots, x_r)$ is a point in $\operatorname{Sym}^\lambda X$, then the preimage of x under the Hilbert-Chow morphism consists of choices of length λ_i ideals supported at x_i for each i, i.e. elements of $(\mathbb{A}^2)^{[\lambda_i]}$. Since each has dimension $\lambda_i - 1$, it follows that the preimage of x has dimension $\lambda_1 + \cdots + \lambda_r - r = n - r$. Via the map $\operatorname{Sym}^\lambda X \to X^r$ sending $x \mapsto (x_1, \ldots, x_r)$, each stratum $\operatorname{Sym}^\lambda$ has dimension $2r$ since X has dimension 2, and so the total dimension of $X_\lambda^{[n]}$ is $n + r$. In particular the unique top-dimensional stratum comes from $\lambda = (1, \ldots, 1)$, the big open piece, and it has dimension $2n$. Thus $\dim X^{[n]} = 2n$.

To check that $X^{[n]}$ is smooth, we then need to compute the dimension of the tangent space at a given k-point Z of $X^{[n]}$. A not-too-hard exercise shows that this tangent space is isomorphic to $\operatorname{Hom}(I, \mathcal{O}_X/I)$ where I is the defining ideal of Z. For homological algebra reasons plus Hirzebruch-Riemann-Roch plus the fact that $\dim H^0(\mathcal{O}_X/I) = n$ by the definition of I we can conclude that the tangent space has dimension $2n$, so $X^{[n]}$ is smooth. (We can give a slightly less sketchy proof in a bit.)

For general n, the symmetric product $\operatorname{Sym}^n X$ is highly singular even for smooth surfaces X. In this case the Hilbert-Chow morphism is a resolution of singularities.

Let's specialize back to our special case $X = \mathbb{A}^2$ for a moment. Notice that here we have an action of the torus $T^2 = k \times k$ via scaling the coordinates. This action lifts to the symmetric product and thus to the Hilbert scheme; its fixed points correspond to homogeneous ideals spanned by monomials, which correspond via Young tableauxs to partitions of n. The singular locus of $(\mathbb{A}^2)^{[n]}$ is stable under the torus action and closed, and so must contain a fixed point if it is nonempty; therefore it suffices to check smoothness at the fixed points, which can be done explicitly via some combinatorics. But since all this is local, this actually gives a proof of smoothness for all smooth surfaces.

2.2 Cohomology and representation theory

The key idea when looking at cohomology is that we shouldn't look at the cohomology of a single Hilbert scheme of points $X^{[n]}$, but rather all of them together. Explicitly, we have a formula

$$\sum_{n=0}^{\infty} \sum_{i=0}^{4n} b_i(X^{[n]}) t^i q^n = \prod_{m=1}^{\infty} \prod_{j=0}^{4} (1 - (-1)^j t^{2m+j-2} q^m)^{-(-1)^j b_j(X)}$$

where b_i is the ith Betti number. Setting $t=-1$, this gives the generating function of the Euler characteristics as

$$\sum_{n=0}^{\infty} \chi(X^{[n]}) q^n = \prod_{j=0}^{4} \left(\prod_{m=1}^{\infty} (1-q^m) \right)^{-(-1)^j b_j(X)} ;$$

in the special case $X = \mathbb{A}^2$, we have $b_i(X) = 0$ for $i > 0$ and so the right-hand side of the first formula is just

$$\prod_{m=1}^{\infty} \frac{1}{1 - t^{2m-2} q^m}.$$

Taylor expanding in q, we can compute the Poincaré polynomials of the first few $(\mathbb{A}^2)^{[n]}$:

$$\begin{aligned}
P_{(\mathbb{A}^2)^{[0]}}(t) &= 1 \\
P_{(\mathbb{A}^2)^{[1]}}(t) &= 1 \\
P_{(\mathbb{A}^2)^{[2]}}(t) &= 1 + t^2 \\
P_{(\mathbb{A}^2)^{[3]}}(t) &= 1 + t^2 + t^4 \\
P_{(\mathbb{A}^2)^{[4]}}(t) &= 1 + t^2 + 2t^4 + t^6 \\
P_{(\mathbb{A}^2)^{[5]}}(t) &= 1 + t^2 + 2t^4 + 2t^6 + t^8,
\end{aligned}$$

and more generally (combinatorically) $b_{2i}((\mathbb{A}^2)^{[n]})$ is the number of partitions of n into $n-i+1$ parts, and all odd Betti numbers are zero.

The "correct" proof of this fact is by viewing the total cohomology as a representation of the Heisenberg algebra. I won't go into much detail here, but vaguely the way that action is built is via certain correspondences on Hilbert schemes. In particular, let $X^{[n,\ell]} \subset X^{[n]} \times X^{[n+\ell]}$ be the scheme classifying subschemes Z_1, Z_2 of lengths n and $n+\ell$ respectively such that $Z_1 \subset Z_2$, and there is a unique point x at which Z_1 and Z_2 vary. Then we have projections to $X^{[n]}$, $X^{[n+\ell]}$, and X. We can use these to build operators on cohomology

$$H^*(X^{[n]}) \to H^*(X^{[n+\ell]})$$

by pulling back, taking the cup product with the pullback of an element of the cohomology of X, and pushing forward; putting a Heisenberg action on $H^*(X)$, this gives the desired action.

3. Singular curves and knots

Our main remaining goal is to say something about the Oblomkov-Shende conjecture and its associates, which Alvaro mentioned last time.

Let C be an integral curve over $\mathbb{C}$, with singularities at worst planar. The example we'll have in mind is $y^k = x^n$.

If there are no singularities, i.e. C is smooth, we know that the Hilbert scheme of points is also smooth, and indeed isomorphic to $\operatorname{Sym}^n C$ with dimension n. Combining these schemes for all n, Macdonald's formula tells us that

$$\sum_{n=0}^{\infty}\sum_{i=0}^{2n} b_i(\operatorname{Sym}^n C)t^i q^n = \frac{(1+qt)^{2g}}{(1-q)(1-t^2q)}$$

where g is the genus of C, and setting $t=-1$ gives

$$\sum_{n=0}^{\infty}\chi(C^{[n]})q^n = \sum_{n=0}^{\infty}\chi(\operatorname{Sym}^n C)q^n = (1-q)^{2g-2} = (1-q)^{-\chi(C)}.$$

More generally if C is singular at points p_i with smooth locus C_{sm}, the Hilbert scheme decomposes as $C^{[n]} = \bigsqcup_{a+b_1+\cdots+b_r=n} C_{\mathrm{sm}}^{[a]} \times C_{p_1}^{[b_1]} \times \cdots \times C_{p_r}^{[b_r]}$, and so

$$\sum_{n=0}^{\infty}\chi(C^{[n]})q^n = \left(\sum_{n=0}^{\infty}\chi(C_{\mathrm{sm}}^{[n]})q^n\right)\prod_i \sum_{n=0}^{\infty}\chi(C_{p_i}^{[n]})q^n.$$

We can treat the first factor as above; the tricky thing is to evaluate the local factors at the singularities.

At each singularity p of C, which is locally embedded in the plane $\mathbb{C}^2$, we can consider a small 3-sphere around p in $\mathbb{C}^2$. The intersection of this sphere with C is a one-dimensional submanifold of S^3, and therefore a link; we call this the link of C at p, and write L_p for it. In a certain sense this captures the topology of C near p. In particular we might hope that link invariants of L_p can describe the sum over Euler characteristics of $\chi(C_p^{[n]})$.

Conjecturally, this is possible. The invariant we need is the HOMFLY polynomial, which we discussed last time, defined by the rules

$$tP(\quad) + t^{-1}P(\quad) = (q-q^{-1})P(\quad)$$

and

$$P(\quad) = 1.$$

Then the desired connection is

$$\sum_{n=0}^{\infty}\chi(C_p^{[n]})q^{2n} = \lim_{t\to 0}(q/t)^{\mu}P(L_p)$$

where μ is the Milnor number of the singularity at p. Thus in all we conjecture

$$\sum_{n=0}^{\infty}\chi(C^{[n]})q^{2n} = (1-q^2)^{-\chi(C_{\mathrm{sm}})}\prod_i \lim_{t\to 0}(q/t)^{\mu_i}P(L_{p_i}).$$

This is the Oblomkov-Shende conjecture (or one of them).

More generally, if we want to recover the whole HOMFLY polynomial, we write $C_{p,m}^{[n]}$ for the subscheme of the Hilbert scheme classifying ideals whose minimal number of generators is m. Then we have

$$P(L_p) = (t/q)^{\mu}(1-q^2)\sum_{n,m}\chi(C_{p,m}^{[n]})(1-t^2)^{m-1}q^{2n}.$$

Let's verify these conjectures for the curve $y^2 = x^n$ for $n = 1, 2, 3$. The link of the singularity at the origin is the $(2, n)$ torus knot $T_{2,n}$ [draw: unknot, linked circles, trefoil], whose HOMFLY polynomials satisfy

$$P(T_{2,n}) = -t(q - q^{-1})P(T_{2,n-1}) + t^2 P(T_{2,n-2}),$$

and for $n = 1$ the curve is smooth and $T_{2,1}$ is the unknot. Therefore $P(T_{2,1}) = 1$, $P(T_{2,2}) = -t(q - q^{-1}) + \frac{t^3-t}{q-q^{-1}}$, and $P(T_{2,3}) = t^2(q^2 + q^{-2}) - t^4$.

For $n = 1$, the local ring at the origin is $\mathbb{C}[[x]]$ with ideals (x^n), and so $C^{[n]}_{p,m}$ is empty unless $m = 1$, in which case it is a point. Therefore the right-hand side is

$$(t/q)^0(1 - q^2)\sum_{n\geq 0} q^{2n} = 1,$$

verifying the conjecture in this case.

If $n = 2$, the local ring is $\mathbb{C}[[x, y]]/(x^2-y^2) = \mathbb{C}[[X, Y]]/(XY)$ where $X = x+y$, $Y = x-y$. The finite colength ideals are either the whole ring, the principle ideals $(X^i + \lambda Y^j)$ of length $i + j$ for $\lambda \in \mathbb{C}^\times$, and (X^i, Y^j) of length $i + j - 1$. Therefore $C^{[n]}_{p,m}$ is trivial for $m > 2$. For $m = 0$, there is only the length 0 ideal (1); for $m = 1$, there are $\mathbb{C}^\times$'s worth of ideals of each length n, and therefore the Euler characteristic is 0; and for $m = 2$ there are n, and so the Euler characteristic is n. Therefore the right-hand side is

$$\frac{t}{q}(1 - q^2)\left(1 + \sum_{n\geq 0} n(1 - t^2)q^{2n}\right) = -t(q - q^{-1}) + t(t^2 - 1)\frac{q}{q^2 - 1},$$

which is the desired formula.

For $n = 3$, we have $\mu = 2$ and the local ring is $\mathbb{C}[[x^2, x^3]]$, whose ideals are (1), $(x^i+\lambda x^{i+1})$ for $\lambda \in \mathbb{C}$, and (x^{i+1}, x^{i+2}), the latter two with length i. Each of the middle type has Euler characteristic 1 and the last type are points, and so the right-hand side is

$$\frac{t^2}{q^2}(1 - q^2)\left(1 + \sum_{n\geq 0} q^{2n} + \sum_{n\geq 0}(1 - t^2)q^{2n}\right) = t^2(q^2 - q^{-2}) - t^4$$

as desired.

These conjectures are known in a variety of special cases, including the limit as $t \to -1$ where we recover the Alexander polynomial, but not in general.