Lecture XIII: Pseudoholomorphic curves I

Jonathan Evans

16th December 2010

It's finally time to introduce the theory of pseudoholomorphic curves. Today we will define pseudoholomorphic curves and discuss and motivate some of their properties. My aim in this lecture is not to prove anything, rather to give you the idea of how J-holomorphic curves behave and then how to use them to prove things. If you want to see proofs, go and look at the big book by McDuff-Salamon.

Definition

Let (X, J) be an almost complex manifold. A J-holomorphic curve in X is a map $u : \Sigma \to X$ from a Riemann surface (Σ, j) to X such that

$$J \circ du = du \circ j$$

Equivalently, in local complex coordinates $z = x + iy$ on Σ and local real coordinates x^j on X

$$\frac{\partial u^j}{\partial x} + J^j{}_i \frac{\partial u^i}{\partial y} = 0$$

We restrict to a 2-dimensional source because it turns out that for a non-integrable J there are even local obstructions (coming from the Nijenhuis tensor) to the existence of higher-dimensional J-holomorphic submanifolds. In contrast there are many local J-holomorphic curves through any given point.

The differential equation governing J-holomorphic curves is "elliptic" and so they share many of the nice properties of holomorphic curves in an integrable complex manifold. For instance:

- (Unique continuation:) Suppose u and v are two J-holomorphic curves $\Sigma \to X$ for which there exists $z \in \Sigma$ at which u and v and all their derivatives agree. Then $u = v$ globally.
- The set of critical points of a compact J-holomorphic curve $u : \Sigma \to X$ is finite. The set of points of Σ which map to critical values of u is finite.
- Unless $u : \Sigma \to X$ factors through a holomorphic branched multiple cover $\Sigma \to \Sigma'$, the set of injective points (where $du \neq 0$ and $u^{-1}(u(x))$ is a single point) is nonempty, open and dense.

A curve which does not factor through a branched cover is called simple. Simple curves are our friends: they have a much better behaved deformation theory than multiple covers. The last property we listed means that a curve is somewhere injective (i.e. has an injective point) if and only if it is simple. Here's the first piece of magic:

Theorem

If A is a homology class in (X, J) and $\mathcal{M}^*(A, J)$ denotes the space of simple J-holomorphic curves u with genus g such that $u_*[\Sigma] = A$ then, for generic J, $\mathcal{M}^*(A, J)$ is a manifold of dimension

$$\dim(X)(1-g) + 2c_1(A).$$

If $\{J_t\}_{t\in[0,1]}$ is a generic path of almost complex structures (with generic endpoints) then the union

$$\bigcup_{t\in[0,1]} \mathcal{M}^*(A, J_t)$$

is also a manifold (a cobordism between the spaces at either end).

This theorem is not so useful since we do not know if any of these spaces are compact: a cobordism only tells you anything if it's compact, for instance the point (or any manifold) is cobordant to the empty set if you allow noncompact cobordisms. Compactness fails in a number of ways.

- The group of holomorphic automorphisms of Σ may be noncompact, for instance if $\Sigma = \mathbb{CP}^1$ it has $\mathrm{Aut}(\mathbb{CP}^1) = \mathbb{P}SL(2, \mathbb{C})$. Since this acts on our space we can hardly expect our space to be compact,
- A family of J-holomorphic curves might degenerate in some unspeakable way,
- For instance, a family of simple J-holomorphic curves might limit to a multiple cover.

Here's an example of some J-curves degenerating.

Example

$[x : y] \mapsto [x^2 : y^2 : \epsilon xy]$ in $\mathbb{CP}^2$ is a family of J-holomorphic maps whose images are conics which are smooth except when $\epsilon = 0, \infty$. What happens as $\epsilon \to 0$? The curve tends to a double cover of $[x : y : 0]$. What happens when $\epsilon \to \infty$? Let's reparametrise to absorb the ϵ into the x variable: $x' = \epsilon x$. This corresponds to looking at the limit of points near $\infty = [1 : 0] \in \mathbb{CP}^1$. The curve becomes $[x' : y] \mapsto [(x')^2/\epsilon^2 : y^2 : x'y]$ and as $\epsilon \to \infty$ this limits to the line $[a : b] \mapsto [0 : b : a]$. Similarly rescaling by $y' = \epsilon y$, the curve $[x : y'] \mapsto [x^2 : (y')^2/\epsilon^2 : xy']$ limits to the line $[a : b] \mapsto [a : b : 0]$, so the limit as $\epsilon \to \infty$ is a union of two lines. Actually any fixed circle $[re^{i\theta} : 1]$ limits to the point $[0 : 0 : 1]$ of intersection between these lines. To recover something sensible you have to look at the limit of a circle which gets closer and closer to 0 or ∞, i.e. $[r\epsilon^{\pm 1}e^{i\theta} : 1]$.

In fact the second piece of magic tells us that all compactness issues can be dealt with in an elegant way. First we define the energy of a pseudoholomorphic curve

$$E(u) = \frac{1}{2}\int_\Sigma |du|^2 \mathrm{dvol}$$

The choice of metric on Σ is irrelevant because E is conformally invariant: pick a metric for which j is orthogonal, rescale it pointwise by some function f and you change $|du|^2$ by f^{-1} (because du is a 1-form with values in u^*TX) and $\mathrm{dvol}_{fg} = \sqrt{\det(fg)}dx \wedge dy = f\,\mathrm{dvol}_g$. Recall that in 2-d, j determines the conformal class of the metric by setting $g = dx^2 + dy^2$ in conformal coordinates $z = x + iy$.

Suppose u_i is a sequence of maps. If we can control pointwise the norm of $|du_i|$ then by Arzela-Ascoli we can find a convergent subsequence of maps. Gromov's amazing compactness theorem tells us what happens if we can only bound the energy. Since energy controls only the L^2-norm of $|du_i|$ we don't expect to get a convergent subsequence because there could be points where $|du_i| \to \infty$ (as the previous example shows). Gromov proved that the worst that can happen is exactly what we described in the previous example!

Bubbling

Definition (Bubbling)

A J-holomorphic curve $u : \mathbb{CP}^1 \to X$ *occurs as a bubble in the limit of a sequence of maps* $u_i : \Sigma \to X$ *if there is a sequence of points* $p_i \in \Sigma$ *and a sequence of discs* $g_i : B_{R_n} \to \Sigma$ *of radius* $R_i \to \infty$ *with* $g_i(0) = p_i$ *such that* $u_i \circ g_i$ *tends to* $f|_{\mathbb{C}\subset\mathbb{CP}^1}$.

We can see that the reparametrisations in our example allowed us to embed larger and larger radius discs into $\Sigma = \mathbb{CP}^1$ centred at 0 and ∞ such that these discs recaptured different parts of the limit. Gromov's theorem tells us this is exactly what happens when $|du_i(p_i)| \to \infty$. We now assume J is ω-tame for some symplectic form ω.

Gromov compactness

Theorem (Gromov compactness, roughly speaking)

A sequence of J-holomorphic curves u_i with bounded energy has a "convergent" subsequence (after possibly reparametrising each u_i) whose "limit" is a union of J-holomorphic curves $\alpha : \Sigma \to X$ and $\beta_j : \mathbb{CP}^1 \to X$ where the curves β_j are a finite number of bubbles attached to points of the curve α. We may assume that the energies $E(u_i)$ converge to some definite E, and furthermore we know that $E = E(\alpha) + \sum E(\beta_j)$.

Note that we get a compactness statement about the *moduli space* $\mathcal{M}^*(A, J)/G$ where G is the reparametrisation group (e.g. $\mathbb{P}SL(2, \mathbb{C})$ for $\Sigma = \mathbb{CP}^1$). I won't prove this theorem as it's probably the hardest part of the basic theory of pseudoholomorphic curves. Instead, I'll talk about where the energy bound comes from.

Energy bounds

Assume J is ω-tame and use the pullback to Σ of the associated metric $g(u,v) = \frac{1}{2}(\omega(u,Jv) + \omega(v,Ju))$ (the usual prescription only works in the compatible case). The energy integrand is locally (in conformal coordinates $z = x + iy$)

$$\begin{aligned} |du|^2 &= |\partial_x u|^2 + |\partial_y u|^2 \\ &= |\partial_x u + J\partial_y u|^2 - 2g(\partial_x u, J\partial_y u) \\ &= |\overline{\partial}_J u|^2 + \omega(\partial_x u, \partial_y u) + \omega(J\partial_s u, J\partial_t u) \end{aligned}$$

When u is J-holomorphic the first term is zero and since $J\partial_x u = \partial_y u$ the two terms on the right both give

$$\omega(\partial_x u, \partial_y u) = u^*\omega(\partial_x, \partial_y)$$

Topological control

Thus the energy of a J-holomorphic curve in an ω-tame almost complex manifold is

$$E(u) = u^*\omega$$

which is topological, so a sequence of curves in the same homology class have equal (in particular bounded) energy.

We also know that if a sequence of curves Gromov converges to a union of curves α, β_j then the ω-areas of the homology classes of α and β_j add up to the ω-area of the original curve. Since J-curves are symplectic submanifolds we know that these areas must be positive, so only a finite amount of bubbling can occur and it's topologically quite controlled. This allows us to rule out bubbling in some (extremely nice) situations. For example, suppose u_i is a sequence of curves in the homology class $[H]$ of a line in $\mathbb{CP}^2$. This class has minimal ω_{FS}-area (equal to 1) and so we cannot bubble off anything (or we'd be left with a component of area zero).

Example existence theorem

We're now in a place where we can sketch how to prove an example existence theorem for holomorphic curves. This is precisely the theorem we used to prove Gromov's nonsqueezing theorem.

Theorem

Let $(X, \omega) = (\mathbb{CP}^1 \times V, \omega_{FS} \oplus \omega_V)$ be a product symplectic manifold of dimension $2n$ where the minimal area of a homology class in V is at least $\int_{\mathbb{CP}^1} \omega_{FS} = 1$. Then for any ω-compatible almost complex structure J and any point p there is a J-holomorphic sphere in the homology class $A = [S^2 \times \{\star\}]$ passing through p.

We first introduce the concept of an *evaluation map*. We take the moduli space of holomorphic spheres with a marked point

$$(\mathcal{M}^*(A, J) \times \mathbb{CP}^1)/\mathbb{P}SL(2, \mathbb{C})$$

where we divide out by simultaneous reparametrisation by $\phi \in \mathbb{P}SL(2, \mathbb{C})$

$$\phi(u, z) = (u \circ \phi, \phi^{-1}(z))$$

and observe that the evaluation map

$$[(u, z)] \mapsto u(z)$$

is well-defined.

Proof.

The moduli space $\mathcal{M}^*(A,J)/\mathbb{P}SL(2,\mathbb{C})$ is compact: there is no bubbling because A has minimal ω-area and a limit curve has to be simple because A is a primitive homology class. The expected dimension of the moduli space $\mathcal{M}^*(A,J)/\mathbb{P}SL(2,\mathbb{C})$ is $2n+2c_1(A)-6$ (6 is the dimension of the Möbius group $\mathbb{P}SL(2,\mathbb{C})$ of reparametrisations, which acts freely) but $c_1(A)=2$ so the moduli space has dimension $2n-2$ for generic J. The space of holomorphic spheres with a marked point is therefore $2n$-dimensional. For $J=J_0$ the evaluation map is clearly of degree 1. It's not hard to check in this case that J_0 is "generic" (details omitted). Now varying J gives a cobordism of moduli spaces and hence a bordism of evaluation maps. Degree is a bordism invariant of maps, so for any J the evaluation map has degree 1. But a degree 1 map is surjective. $\square$

We observe the following consequence

Theorem (Gromov)

If $(X, \omega) = (S^2 \times S^2, \omega_{FS} \oplus \omega_{FS})$ and J is an ω-compatible almost complex structure then there is a J-holomorphic sphere in each homology class $A = [S^2 \times \{\star\}]$, $B = [\{\star\} \times S^2]$ through every point. In fact there is exactly one J-holomorphic A-sphere α_p and one J-holomorphic B-sphere β_p through any given p and α_p intersects β_p exactly once, transversely.

The only part of this theorem we haven't proved is the uniqueness and intersection property. This phenomenon is very special to 4-dimensional symplectic geometry and I'll explain it in a moment. We will use the theorem next week to show that the symplectomorphism group of (X, ω) is homotopy equivalent to $SO(3) \times SO(3) \ltimes \mathbb{Z}/2$.

Positivity of intersections in 4d

Theorem (Positivity of intersections, Gromov-McDuff)

Suppose a pair of J-holomorphic curves in a 4-manifold intersect in some set of points p_i and that the intersection at p_i has multiplicity m_i (e.g. $m_i = 1$ means transverse, $m_i = 2$ means tangent but the second derivatives are linearly independent, etc.). Then $m_i > 0$. That is to say each intersection point contributes a positive amount to the homological intersection of the two curves and that contribution is 1 if and only if the intersection is transverse.

Since the curves α_p and β_p in the previous theorem had homological intersection 1 and $[\alpha_p] \cdot [\alpha_p] = [\beta_p] \cdot [\beta_p] = 0$, the unproven parts of that theorem follow from this positivity of intersections property. Positivity of intersections is not easy to prove, but it *is* easy to see that it fails in higher dimensions, even for integrable complex structures: a pair of complex lines in $\mathbb{CP}^3$ can happen to intersect, but one can easily disjoin them because they each have real codimension 4.

That's it for our whirlwind introduction to pseudoholomorphic curve theory. Next lecture we'll see how to use them to prove some more cool theorems.