Mini-Course on Moduli Spaces

Emily Clader
June 2011

1 What is a Moduli Space?

1.1 What should a moduli space do?

- Suppose that we want to classify some kind of object, for example:
  - Curves of genus g,
  - One-dimensional subspaces of $\mathbb{R}^n$,
  - Finite sets.
- "Classify" means that we should have a notion of equivalence, which we are free to choose.
- Typically, such objects come in families. (This is a way of making precise the idea that two such objects are "near" one another.) Once again, the exact notion of family is up to us to choose, but in general a family over a base variety B will be a flat morphism

  $$X \to B$$

  whose fibers are all objects of the type we are trying to classify.
- Having made all these decisions, a moduli space should be a variety (or manifold, or orbifold, or ...) whose points are in bijection with equivalence classes of objects, and whose algebraic structure reflects how the objects vary in families.
- As an example of the latter notion, a curve in the moduli space should trace out a one-parameter family of objects.

1.2 Example: quadruples of points in $\mathbb{P}^1$

- A *quadruple* of points in $\mathbb{P}^1$ is an ordered set of four distinct points (p_1, p_2, p_3, p_4) in $\mathbb{P}^1$.
- A *family* of quadruples over a base variety B is a diagram

  $$\begin{array}{c} B \times \mathbb{P}^1 \\ \pi \downarrow \;\; \uparrow\uparrow\uparrow\uparrow \sigma_i \\ B \end{array}$$

in which π is the projection and the four sections σ_i are disjoint.

- The moduli space of quadruples is $Q := (\mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1) \setminus$ diagonals.
- It is clear that the points of Q are in bijection with quadruples. But is it true that a curve in Q traces out a one-parameter family of quadruples, and more generally that the algebraic structure of Q captures how quadruples vary in families?
- Notice that Q has a *universal family*:

$$\begin{array}{c} Q \times \mathbb{P}^1 \\ \pi \downarrow \; \uparrow\uparrow\uparrow\uparrow \sigma_i \\ Q \end{array}$$

where $\sigma_i(\mathbf{p}) = (\mathbf{p}, p_i)$.

- This family is "tautological" in the sense that the fiber over a quadruple $\mathbf{p}$ is exactly the quadruple $\mathbf{p}$ itself!
- But moreover, it is "universal" in the sense that for any family $B \times \mathbb{P}^1 \to B$, there is a unique morphism $\varphi : B \to Q$ such that the original family is obtained from the universal family by pullback along φ.
- (What does pullback mean? Given a family $B \times \mathbb{P}^1 \to B$ with sections $x \mapsto (x, \sigma_i(x))$ and a morphism $\varphi : B' \to B$, the pullback family is the family $B' \times \mathbb{P}^1 \to B'$ whose sections are $x \mapsto (x, \sigma_i(\varphi(x)))$.)
- This universality is not hard to prove: the map $B \to Q$ is the map $b \mapsto$ point of Q corresponding to the quadruple over b.
- The existence of such a universal family is exactly what is needed to ensure that statements like "curves in the moduli space trace out one-parameter families" hold. For example, a curve in the moduli space might be a map $\mathbb{A}^1 \to Q$, and there is a bijection between such maps and families over $\mathbb{A}^1$.
- We have arrived at the definition of a fine moduli space:

1.3 Definition of a moduli space

- Suppose we have a *moduli problem*– that is:
  - A notion of object and equivalence of objects,
  - A notion of family over a base scheme B and equivalence of families,
  - A notion of pullback of families compatible with equivalence.

- Then a *fine moduli space* for this moduli problem is a scheme M that admits a "universal family" X. This is a family over M with the property that every other family over a scheme B is obtained, up to equivalence, by pulling back X via a unique morphism $\kappa : B \to M$.
- Thus, there is a bijection:

$$\{\text{equivalence classes of families over } B\} \leftrightarrow \{\text{morphisms } B \to M\}.$$

- Some things to notice:
  - An object is simply a family over a point, so it is a consequence of this definition that the points of M are in bijection with equivalence classes of objects.
  - The universal family is always tautological. Indeed, for $x \in M$, consider the inclusion morphism $\kappa : \{x\} \to M$. Suppose that x corresponds to an object C_x; by definition, this means that $\kappa^* X \cong C_x$, where we consider C_x as a family over the point. But $\kappa^* X = \pi^{-1}(x)$, where $\pi : X \to M$ is the universal family. Thus, $\pi^{-1}(x) \cong C_x$.

1.4 Example 2: quadruples up to projective equivalence

- Let's explore this idea a bit further by elaborating on our previous example.
- Two quadruples are *projectively equivalent* if there is an automorphism of $\mathbb{P}^1$ taking one quadruple (in order) to the other.
  - It may be helpful to remember that the automorphisms of $\mathbb{P}^1$ are precisely the Mobius transformations $x \mapsto \frac{ax+b}{cx+d}$, and that there exists a unique automorphism sending any three fixed points to any other.
- It is easy to extend the notion of projective equivalence to families.
- Recall the *cross ratio* from complex analysis: given a quadruple $\mathbf{p} = (p_1, p_2, p_3, p_4)$, let $\lambda(\mathbf{p}) \in \mathbb{P}^1 \setminus \{0, 1, \infty\}$ be the image of p_4 under the unique automorphism sending p_1, p_2, p_3 (in order) to $0, 1, \infty$.
- **Fact**: Two quadruples are projectively equivalent if and only if they have the same cross ratio.
- Therefore, the points of $\mathcal{M}_{0,4} := \mathbb{P}^1 \setminus \{0, 1, \infty\}$ are in bijection with the equivalence classes of quadruples.
- To be a fine moduli space, we need a universal family. It is the family $\mathcal{M}_{0,4} \times \mathbb{P}^1 \to \mathcal{M}_{0,4}$ with sections

$$\begin{aligned}
\sigma_1(p) &= 0 \\
\sigma_2(p) &= 1 \\
\sigma_3(p) &= \infty \\
\sigma_4(p) &= p.
\end{aligned}$$

- It is easy to see that this family is tautological: the fiber over p is a quadruple with cross-ratio p. What's a bit harder is to see that it's universal... but it is.
- The same proof shows that there is a fine moduli space

$$\mathcal{M}_{0,n} = (\mathcal{M}_{0,4} \times \cdots \times \mathcal{M}_{0,4}) \setminus \text{diagonals}$$

  ($n-3$ factors) of n-tuples up to projective equivalence for any $n \geq 4$

1.5 Other examples

- $\mathbb{P}^n$ is a fine moduli space for the moduli problem of lines through the origin in $\mathbb{R}^{n+1}$.
- $\mathbb{N}$ is a fine moduli space for the moduli problem of finite sets up to bijection.

1.6 When fine moduli spaces don't exist

- What would you expect to be the moduli space for one-dimensional vector spaces up to isomorphism?
- You'd expect it to be a point, since there's only one.
- But this can't be a fine moduli space. Indeed, a family of one-dimensional vector spaces is a line bundle, and over the base $B = S^1$, there are two non-isomorphic line bundles: the cylinder and the Mobius band. Since the moduli space is a point, these must induce the same map to the moduli space, despite being non-isomorphic.
- This is an example of a general phenomenon: when the objects we are parameterizing have nontrivial automorphisms, a fine moduli space generally does not exist. This is quite a problem; many objects admit automorphisms, and therefore fine moduli spaces almost never exist!
- There are three basic solutions:
  - Modify the moduli problem, adding extra structure to kill automorphisms.
  - Look for a moduli space that is not fine but "coarse".
  - Look for a moduli stack/orbifold rather than a moduli space.

  I will briefly discuss each of these strategies.

1.7 Modify the moduli problem

- For example, in the case of one-dimensional vector spaces, we can add the datum of $1 \in V$. Thus, an isomorphism of vector spaces would be required to preserve this datum, and this would imply that there are no nontrivial automorphisms.

- A family, then, would be a line bundle together with a nowhere-vanishing section. In particular, the Mobius band would no longer be an admissible family, so the problem observed above would disappear.
- Indeed, every family over any base B would be trivial. This implies that the point is a fine moduli space in this case.

1.8 Coarse moduli spaces

- In order to explain the definition of a coarse moduli space, we need a digression into category theory (which, arguably, is the right language to talk about moduli spaces in, anyway).
- One can define a moduli problem as a contravariant functor

$$F : \mathrm{Sch} \to \mathrm{Set},$$

  where $F(B)$ is the set of equivalence classes of families over B, and $F(\varphi)$ is the pullback map on families.
- To say that a fine moduli space exists, in this language, is to say that F is representable.
- What does this mean?
- Any scheme M gives rise to a *functor of points*, a contravariant functor

$$h_M : \mathrm{Sch} \to \mathrm{Set}$$

  given by

$$B \mapsto \mathrm{Hom}(B, M),$$
$$(\varphi : B' \to B) \mapsto (\beta \mapsto \beta \circ \varphi).$$

- We say that F is representable if F is naturally isomorphic to h_M for some scheme M.
- (If you know Yoneda's lemma, you know that the isomorphism $U : h_M \to F$ can be viewed as an object in the set $F(M) = \{$equivalence classes of families over $M\}$. Which equivalence class of families is it? It's the universal family!)
- This makes sense: it says that for any scheme B, equivalence classes of families over B are in bijection with morphisms $B \to M$.
- A *coarse moduli space* for a moduli functor F is a pair (M, V), where M is a scheme and $V : F \to h_M$ is a natural transformation (*not* necessarily an isomorphism) satisfying:
  1. (M, V) is initial among all such pairs,
  2. The set map $V_{\mathrm{point}} : F(\mathrm{point}) \to \mathrm{Hom}(\mathrm{point}, M)$ is a bijection.

- The second of these conditions says that points of M are in bijection with objects we are trying to classify, while the first says that every natural transformation $F \to h_{M'}$ factors uniquely through V, so that h_M is the representable functor "closest" to F.

- Exercise: The point is a coarse moduli space for one-dimensional vector spaces.

1.9 Moduli stacks/orbifolds

- An *orbifold* is a generalization of a manifold; rather than looking locally like $\mathbb{R}^n$, an orbifold looks locally like $\mathbb{R}^n/G$ for some finite group G, called the *isotropy group* of the point around which we are looking.

- Orbifolds are naturally encoded in the language of category theory via groupoids.

- A *groupoid* is a category in which every morphism is an isomorphism.

- We can view a quotient X/G (where X is a manifold and G is a finite group) as a groupoid: the objects are X, and there is an arrow between two objects for every element of G sending one object to the other.

- By gluing together this local picture, it is possible to view any orbifold as a groupoid.

- A stack is an even further generalization, but I won't go into any details.

- Notice that orbifolds appear to be exactly the right sort of object to salvage a fine moduli space when the existence of one was prohibited by the presence of nontrivial automorphisms, since every point in an orbifold has an isotropy group that can encode the "automorphisms" of that point. More precisely...

- Consider a quotient orbifold $[X/G]$. (Remember, this is a category with objects X and morphisms $X \times G$.) Given a scheme B, the definition of orbifold morphism $B \to [X/G]$ is cooked up in such a way that such morphisms are in bijection with pairs consisting of:

  1. A principal G-bundle $p : P \to B$,
  2. A G-equivariant map $\phi : P \to X$.

- For example, the easiest quotient orbifold is [point/G]. Then orbifold morphisms $B \to$ [point/G] are in bijection with principal G-bundles over B. (Thus, [point/G] is the "classifying stack" for principal G-bundles.)

- Remember, a fine moduli space is defined by the fact that maps into the moduli space correspond to equivalence classes of families. The above definition of morphism is constructed precisely so that is corresponds to a family of objects having automorphism group G.

- For example, the moduli stack of one-dimensional vector spaces is [point/$\mathbb{C}^*$]. Why? A line bundle over a base scheme B is the same as a principal $\mathbb{C}^*$ bundle, with the correspondence sending the line bundle L to $L \setminus \{\text{zero section}\}$.

2 The Moduli Space of Curves

2.1 The construction

- We would like to construct a fine moduli space/orbifold $\mathcal{M}_g$ for curves (that is, compact connected Riemann surfaces) of genus g.
- We will need to adjust this goal in two (related) ways:
  - As it stands, the moduli problem does not have a compact moduli space. We will need to allow certain nodal curves in order to compactify it.
  - Considering nodal curves naturally leads one to the idea of "marked points", since a nodal curve is specified by a number of smooth curves together with a number of points at which they are glued. Thus, the right objects to parameterize are actually n-pointed, genus-g curves, whose moduli space is denoted $\mathcal{M}_{g,n}$.
- Let's focus on $\mathcal{M}_{0,n}$ first.
- An *n-pointed smooth rational curve* is a projective smooth rational curve (that is, a $\mathbb{P}^1$) with n distinct marked points.
- An *isomorphism* of such is an isomorphism of curves that respects the marked points, in order.
- A *family* of such is a flat, proper morphism $\pi : X \to B$ equipped with n disjoint sections $\sigma_i : B \to X$ such that each geometric fiber $\pi^{-1}(b)$ is a projective smooth rational curve.
- One can easily extend the notion of isomorphism to families.
- **Fact**: Given any family $X \to B$ of n-pointed smooth rational curves, there is a unique isomorphism between this family and the trivial family $B \times \mathbb{P}^1 \to B$, such that the first three sections are mapped to the constant sections at 0, 1, and ∞.
- Thus, a family is completely determined by where the remaining sections are sent under the above isomorphism. It follows that the moduli problem of classifying n-pointed smooth rational curves is equivalent to the moduli problem of classifying n-tuples of points in $\mathbb{P}^1$. We saw for quadruples (and it's easy to extend this to all n-tuples) that this moduli problem admits a fine moduli space.
- **Example**: $\mathcal{M}_{0,3} =$ point.
- **Example**: $\mathcal{M}_{0,4} = \mathbb{P}^1 \setminus \{0, 1, \infty\}$.
- But there's a problem: this moduli space is in general noncompact. We'd eventually like to do intersection theory on it, so we really need to compactify. But we should do this in a geometrically meaningful way, so that the addtional points that we add correspond to some objects that are a slight generalization of n-pointed smooth rational curves.

2.2 The compactification

- Notice that the noncompactness can be understood easily from a geometric perspective: marked points are necessarily distinct, so a sequence of curves in which two marked points get closer and closer has no limit in the moduli space.

- This leads to a natural guess as to how to compactify: allow the marked points to coincide.

- Unfortunately, this doesn't work. For example, consider the family

  $$(\mathbb{P}^1; 0, 1, \infty, t),$$

  where $t \in \mathbb{P}^1 \setminus \{0, 1, \infty\}$. As t tends to 0, we would expect this family to tend to $(\mathbb{P}^1; 0, 1, 0)$. But this family is isomorphic to the family

  $$(\mathbb{P}^1; 0, 1/t, \infty, 1),$$

  and in this case letting t tend to 0 gives $(\mathbb{P}^1; 0, \infty, \infty, 1)$. Thus, allowing coincident marked points gives a non-separated moduli space.

- There's a remarkably elegant answer to this problem: since there is no way to decide between the above two limits, keep both of them! More precisely, the limit of this family is a nodal curve with one component in which x_2 and x_3 coincide and another in which x_1 and x_4 coincide.

- We like to think of this as saying that when two marked points approach one another, in the limit the curve sprouts an extra component to receive them.

- Now, we should only add as many new objects into our moduli space as we need to in order to compactify. It turns out that, rather than allowing any old nodal curve, we can get away with only the following:

- A *stable n-pointed curve* is a curve with only ordinary double points as singularities, and with n distinct marked points that are smooth points of the curve, such that there are only finitely many automorphisms as a pointed curve.

- (The notion of genus for such a thing is arithmetic genus, given by $h^0(\omega)$. Intuitively, this means that a "hole" can come either from a hole in an irreducible component or a circuit in the dual graph.)

- In genus 0, this last condition is equivalent to the requirement that every irreducible component has at least three special points.

- Indeed, this ensures that there are no nontrivial automorphisms at all in genus 0, since each component is isomorphism to $\mathbb{P}^1$ and a map from $\mathbb{P}^1$ to itself that fixes three points must be the identity.

- **Fact**: For each $n \geq 3$, there is a smooth projective variety $\overline{\mathcal{M}}_{0,n}$ that is a fine moduli space for stable n-pointed rational curves. It contains $\mathcal{M}_{0,n}$ as a dense open subset.

- **Example**: $\overline{\mathcal{M}}_{0,4} = \mathbb{P}^1$. This had better be the right answer, since it's a compactification of $\mathcal{M}_{0,4} = \mathbb{P}^1 \setminus \{0, 1, \infty\}$, but we can also see explicitly what the three additional points are:

  Moreover, we can see explicitly what the universal family is:

- For general genus, the automorphism condition in the definition of a stable curve is satisfied as long as every genus-0 component has at least three special points and every genus-1 component has at least one. However, there will in general be automorphisms, just finitely many. Thus, we should expect:

- **Fact**: For $(g, n) \notin \{(0,0), (0,1), (0,2), (1,0)\}$, there is a compact *orbifold* $\overline{\mathcal{M}}_{g,n}$ that is a fine moduli space for stable n-pointed genus-g curves.

- There are a number of beautiful properties of these moduli spaces that make them extermely useful.

2.3 Attaching and forgetful maps

- There are attaching maps

$$\overline{\mathcal{M}}_{g_1,n_1+1} \times \overline{\mathcal{M}}_{g_2,n_2+1} \to \overline{\mathcal{M}}_{g_1+g_2,n_1+n_2}$$

  and

$$\overline{\mathcal{M}}_{g-1,n+2} \to \overline{\mathcal{M}}_{g,n}.$$

- There is also a forgetful map
$$\overline{\mathcal{M}}_{g,n+m} \to \overline{\mathcal{M}}_{g,n}$$
that deletes some marked points and collapses unstable components.

- The attaching maps give a nice recursive structure to the "boundary" $\overline{\mathcal{M}}_{g,n} \setminus \mathcal{M}_{g,n}$. In particular, notice that the boundary is the union of the images of the attaching maps. In genus zero, the attaching maps are injective, so this implies that the boundary of $\overline{\mathcal{M}}_{0,n}$ is isomorphic to a product of moduli spaces $\overline{\mathcal{M}}_{0,m}$ for $m \leq n$. However, in higher genus, this is no longer the case; nevertheless, the boundary can be obtained by taking products and quotients of smaller moduli spaces.

2.4 Kapranov's recursive construction

- There is a stronger sense in which the moduli spaces of curves are recursive.

- Recall that since $\overline{\mathcal{M}}_{g,n}$ is a fine moduli space (or at least a fine moduli orbifold), it carries a universal family $U_{g,n}$ with the property that every family of genus-g, n-pointed stable curves over a base B is the pullback of $U_{g,n}$ under a unique morphism $B \to \overline{\mathcal{M}}_{g,n}$.

- **Claim**: $U_{g,n} \cong \overline{\mathcal{M}}_{g,n+1}$, and under this isomorphism, the universal family $U_{g,n} \to \overline{\mathcal{M}}_{g,n}$ is the forgetful morphism.

- **Sketch of proof**: The map
$$\phi : U_{g,n} \to \overline{\mathcal{M}}_{g,n+1}$$
is as follows. Choose $q \in U_{g,n}$, and suppose that $\pi(q) = C \in \overline{\mathcal{M}}_{g,n}$. Then, since the universal family is tautological, we know that $\pi^{-1}(C)$ is C itself, so $q \in C$. We would like to say $\phi(q)$ is the $(n+1)$-pointed curve obtained by adding q to C as a marked point. However, this may not be stable, since q might be a node, or it might already be a marked point. There is a canonical stabilization, though, obtained by adding an extra component in these cases. This stabilization is $\phi(q)$.

  It is not too difficult to convince oneself that this gives a set-theoretic bijection. Is it an isomorphism? Well, in fact, one uses this bijection to *define* $\overline{\mathcal{M}}_{g,n+1}$, giving it orbifold structure in such a way that ϕ is forced to be an isomorphism. All that needs to be checked is that this actually does define a fine moduli space; one does this by explicitly constructing a universal family via a fiber product.

3 Further Topics

3.1 Deformation theory

- By definition, a *deformation* of a smooth variety X over a base (pointed) scheme (Y, y_0) is a proper flat morphism $\varphi : \mathcal{X} \to Y$ together with an isomorphism $\psi : X \to \varphi^{-1}(y_0)$.
- A *first-order* or *infinitesimal deformation* is a deformation over the base $(\mathrm{Spec}(\mathbb{C}[\epsilon]/\epsilon^2), *)$, where $*$ denotes the closed point (ϵ).
- **Claim**: The space of infinitesimal deformations of C is isomorphic to $T_C\mathcal{M}_g$.
- **Proof**: Recall that the tangent space can be defined as $(\mathfrak{m}_C/\mathfrak{m}_C^2)^\vee$, where $\mathfrak{m}_C$ is the maximal ideal of the local ring of $\mathcal{M}_g$ at C. Thus, a tangent vector in $T_C\mathcal{M}_g$ is the same as a map $\mathrm{Spec}(\mathbb{C}[\epsilon]/(\epsilon^2)) \to \mathcal{M}_g$ sending the closed point to C. By the definition of a fine moduli space, this is the same as a flat family over $\mathrm{Spec}(\mathbb{C}[\epsilon]/(\epsilon^2))$, i.e. an infinitesimal deformation.
- **Claim**: The space of infinitesimal deformations of C is isomorphic to $H^1(C, T_C)$.
- Thus, the vector spaces $H^1(C, T_C)$ fit together into a vector bundle $\mathrm{Def}(C)$ over the moduli space $\mathcal{M}_g$, and this bundle is nothing but the tangent bundle.
- One great upshot of this is that we can compute the dimension of $\mathcal{M}_g$.
- Indeed, the dimension of $\mathcal{M}_g$ (as an orbifold) is the same as the dimension of $T_C\mathcal{M}_g$ (as a vector space) for any curve C, which is $\dim(H^1(C, T_C))$. By Serre duality:

  $$H^1(C, T_C) \cong H^0(C, T_C^\vee \otimes \omega_C)^\vee \cong H^0(C, \omega_C^{\otimes 2}),$$

  since $T_C^\vee = \omega_C$. And by Riemann-Roch:

  $$\dim(H^0(C, \omega^{\otimes 2})) - \dim(H^1(C, \omega^{\otimes 2})) = \deg(\omega^{\otimes 2}) + g - 1 = 2g - 2 + g - 1 = 3g - 3.$$

  Since $H^1(C, \omega^{\otimes 2}) \cong H^0(C, \omega^\vee)^\vee$ and $\deg(\omega^{vee}) = 2-2g < 0$, it follows that $H^1(C, \omega^{\otimes 2}) = 0$. Thus,

  $$\dim(\mathcal{M}_g) = \dim(H^1(C, T_C)) = 3g - 3.$$

3.2 Gromov-Witten theory

- For a smooth projective variety X and a homology class $\beta \in H_2(X; \mathbb{Z})$, let

  $$\overline{\mathcal{M}}_{g,n}(X, \beta) = \{f : C \to X \mid C \text{ an } n\text{-pointed curve of genus } g,\ f_*[C] = \beta,\ f \text{ stable}\}/\sim.$$

- Here, a map is "stable" if it has finitely-many automorphisms. This is equivalent to the condition that every irreducible component of C that is mapped to a point by f is stable as a genus-g curve.

- These are parameterized by a fine moduli stack $\overline{\mathcal{M}}_{g,n}(X,\beta)$.
- This is not a particularly nice object. It isn't smooth, and it has components of different dimensions. Thus, it doesn't carry a fundamental class, which we really need because we want to integrate over the moduli space.
- Fortunately, however, it has a "virtual fundamental class". This has a number of nice properties:
  - The virtual fundamental class of one moduli space pulls back to the virtual fundamental class of the other under the forgetful morphism.
  - In cases where $\overline{\mathcal{M}}_{g,n}(X,\beta)$ actually is a smooth orbifold with no components of excessive dimension, the virtual fundamental class is the ordinary fundamental class.
  - The virtual fundamental class always lives in the "expected dimension", which is

$$\dim(\mathrm{Def}(C,f)) - \dim(\mathrm{Ob}(C,f)).$$

- One other nice thing about the moduli stack $\overline{\mathcal{M}}_{g,n}(X,\beta)$ is that it has evaluation maps

$$\mathrm{ev}_i : \overline{\mathcal{M}}_{g,n}(X,\beta) \to X$$

given by $f \mapsto f(x_i)$, where x_i is the i^{th} marked point.

- Given a nonsingular projective variety X and an n-tuple of cohomology classes $T_1, \ldots, T_n \in H^*(X)$, the corresponding *Gromov-Witten invariant* is defined as

$$\langle T_1, \ldots, T_n \rangle_{g,\beta}^X = \int_{[\overline{\mathcal{M}}_{g,n}(X,\beta)]^{\mathrm{vir}}} \mathrm{ev}_1^*(T_1) \cup \cdots \cup \mathrm{ev}_n^*(T_n).$$

- What does this measure?
- Suppose that $T_1, \ldots, T_n$ are Poincaré dual to subvarieties $H_1, \ldots, H_n$ of X. Then $\mathrm{ev}_i^*(H_i)$ is Poincaré dual to the subvariety of $\mathcal{M}_{g,n}(X,\beta)$ consisting of maps sending the i^{th} marked point to H_i, and integrating the product of these classes amounts to intersecting these n subvarieties. Thus, the Gromov-Witten invariant gives a count of the number of stable maps of class β that send the n marked points to the n chosen subvarieties.... at least in good cases.
- **Example**: Let $X = \mathbb{P}^2$, let $g = 0$, and let $\beta = d[\text{hyperplane}]$. One can check that

$$\dim(\overline{\mathcal{M}}_{0,n}(\mathbb{P}^2, d)) = 3d - 1 + n.$$

In particular, if H is the class of a point, then

$$\langle H \cdots H \rangle_{0,d}^{\mathbb{P}^2} = \int_{[\mathcal{M}_{0,3d-1}(\mathbb{P}^2,d)]} \mathrm{ev}_1^*(H) \cup \cdots \cup \mathrm{ev}_n^*(H)$$

has the right degrees to be nonzero. It counts the number of degree d rational curves that pass through $3d - 1$ fixed (generic) points. Using the recursive structure of the moduli space, one can obtain a beautiful equation that determines all of these numbers.