TOPOLOGICAL QUANTUM FIELD THEORY

ARUN DEBRAY

These notes were taken in a class given by Katrin Wehrheim at UC Berkeley in Spring 2020. I live-TEXed them using `vim`, and as such there may be typos; please send questions, comments, complaints, and corrections to a.debray@math.utexas.edu.

Contents

1. TQFT: definition and Atiyah's examples: 2/19/20

We begin with the definition of a topological quantum field theory due to Atiyah, now over 30 years ago.

Definition 1.1. Fix a base field k. A d-dimensional *topological quantum field theory* (TQFT) consists of data of, for every closed, oriented, smooth d-manifold, a finitely generated k-vector space $Z(\Sigma)$, and for every compact, oriented, smooth $(d+1)$-manifold M, an element $Z(M) \in Z(\partial M)$, satisfying some axioms.

Many interrelated ideas went into this definition: Segal's mathematical formalization of two-dimensional conformal field theory, mathematical perspectives on quantum field theory (fields, Hilbert spaces, etc.).

Later, Atiyah's definition was packaged more concisely into asking for Z to be a symmetric monoidal functor

$$Z\colon \mathcal{C}ob_{n,n-1} \longrightarrow \mathcal{V}ect_k, \tag{1.2}$$

where $\mathcal{V}ect_k$ is the symmetric monoidal category of k-vector spaces with tensor product, and $\mathcal{C}ob_{n,n-1}$ is the cobordism category, whose objects are closed, oriented $(n-1)$-manifolds and whose morphisms are (diffeomorphism classes of) oriented bordisms between them. Cylinders in the cobordism category can be thought of as time evolution, but the inclusion of all other bordisms has something to do with a relativistic perspective.

Remark 1.3. Atiyah used d to denote the dimension of space, i.e. the dimension of manifolds assigned vector spaces. These days, it's more common to refer to the cobordism category using the top dimension (what we just called n), the "spacetime dimension." ◂

There are many different flavors of the cobordism category. Some of these involve technical details that we have to account for: for example, even compact 0-manifolds are too big to form a set, so to more accurately define $\mathcal{C}ob_{1,0}$ (or in any dimension) we should pick a set of representatives of oriented diffeomorphism classes of $(n-1)$-manifolds.

Remark 1.4. There are other ways to work around set-theoretic issues: for example, the topological cobordism category of Galatius, Madsen, Tillmann, and Weiss begins with the space $\mathbb{R}^\infty$ and works with manifolds and bordisms explicitly embedded in $\{t\}\times\mathbb{R}^\infty$, resp. $[t_1,t_2]\times\mathbb{R}^\infty$. Then one must quotient out by diffeomorphisms, just as in the abstract cobordism category, but now we don't just have the "internal diffeomorphisms" of an embedded M, but also "external diffeomorphisms" of the ambient space that carry M to something diffeomorphic, but embedded via a different map. We will not work with embedded bordisms, at least for now. ◂

There are several other generalizations we won't discuss today, but are worth mentioning.

- There's notions of *topological conformal field theory* (TCFT) and *homological conformal field theory* (HCFT), in which $\mathcal{C}ob_{2,1}$ is upgraded to a category where bordisms carry some additional structure (e.g. a conformal structure), and we only identify conformally equivalent bordisms.
- In *fully extended* topological quantum field theory, $\mathcal{C}ob_{n,n-1}$ becomes an (∞, n)-category, by allowing manifolds in all dimensions n and below.

In both cases, we must replace the target category $\mathcal{V}ect_k$ with something related, but different.

In these, and in any, generalizations, the overarching question is: *what kind of algebraic structure do we get from these field theories?* To address this question, we generally must first fix a target category. But there are a few "holy grail" theorems in some of these settings.

Theorem 1.5 (Cobordism hypothesis (Lurie))**.** *A fully extended topological field theory is determined by its value on the* 0*-manifold* pt_+*.*

This is more of a slogan than a theorem, but one can pin it down into a precise theorem, e.g. by making precise what kinds of TQFTs one considers. Here, by "pt_+" we might more generally mean looking at generators and relations of the appropriate bordism category.

Example 1.6. In (spacetime) dimension $n = 1$, TQFTs are vacuously fully extended (with the caveat that 1-categorical and $(\infty, 1)$-categorical TQFT aren't quite the same). Then, the theorem is that for any symmetric monoidal category $\mathcal{C}$, $\mathcal{F}un^{\otimes}(\mathcal{C}ob_{1,0}, \mathcal{C})$ is equivalent to the groupoid of *dualizable* objects in $\mathcal{C}$.[1]

Exercise 1.7. For $\mathcal{C} = \mathcal{V}ect_k$, check that dualizability is equivalent to being finite-dimensional.[2]

So fixing $\mathcal{C} = \mathcal{V}ect_k$ for now, given a one-dimensional *unoriented* (i.e. manifolds and bordisms in $\mathcal{C}ob_{1,0}$ are not oriented) TQFT Z we get a finite-dimensional vector space $V := Z(\mathrm{pt}_+)$, and a biliear pairing $e\colon V \otimes V \to k$. This pairing must be nondegenerate, as one can show via the "Z-diagram" being equivalent to an interval (which is the identity $\mathrm{pt} \to \mathrm{pt}$).

Conversely, given a finite-dimensional vector space V and an inner product $e\colon V \otimes V \to k$, we can build a TQFT $Z_{V,e}\colon \mathcal{C}ob \to \mathcal{V}ect_k$, because there aren't that many diffeomorphism classes of 1-manifolds, so we know generators and relations: the interval, regarded as a bordism from $\mathrm{pt} \to \mathrm{pt}$, is sent to id_V; the interval, regarded as a bordism from $\mathrm{pt} \amalg \mathrm{pt} \to \varnothing$, is sent to e; and the interval, regarded as a bordism $\emptyset \to \mathrm{pt} \amalg \mathrm{pt}$, is sent to the adjoint of e. ◂

Remark 1.8. Some things change in the oriented 1-dimensional case. We don't need the inner product: if you keep careful track of the orientations induced on a boundary, the interval is now a bordism between $\mathrm{pt}_+ \amalg \mathrm{pt}_-$ and $\varnothing$, and one can show that $\mathrm{pt}_- \mapsto V^*$. Then these intervals are sent to the evaluation map $V \otimes V^* \to k$ and the coevaluation map $k \to V \otimes V^*$. ◂

In dimension 1, the cobordism hypothesis feels somewhat silly. But in higher dimensions things can quickly get nontrivial, and difficult. For example, for the oriented 2-dimensional cobordism category (before we extend), this is known by the classification of surfaces: the pair of pants, regarded as a bordism $S^1 \amalg S^1 \to S^1$, and, separately, regarded as a morphism $S^1 \to S^1 \amalg S^1$; the disc, both as a bordism $S^1 \to \varnothing$ and as a bordism $\varnothing \to S^1$; and the cylinder $S^1 \to S^1$. In dimension 1 we just have the circle. But if we try to extend down to points, then discovering generators is more complicated — now we have to determine generators and relations using surfaces with corners. The surface theory isn't that bad, and this will get worse when we care about higher-dimensional manifolds.

And we do care about higher-dimensional manifolds: two key questions in this course will be:

(1) how does this (both the axiomatic structure of TQFT and tools such as the cobordism hypothesis) help build invariants for 3- and 4-manifolds, and
(2) how do geometric/PDE-based invariants of 3- and 4-manifolds yield TQFTs?

With regards to question (2) specifically, Atiyah gave a few examples in his original paper on TQFT.

[1] *A priori*, the subcategory of dualizable objects in $\mathcal{C}$ is not a groupoid, but we can make it one by throwing out the non-invertible morphisms.

[2] In higher dimensions, "dualizable" generalizes to "fully dualizable," and the fact that "fully dualizable" and "finite-dimensional" have the same initials makes for a good mnemonic.

Example 1.9. This example, built on work of Floer and Gromov, is a 2-dimensional TQFT. Fix a symplectic manifold (X, ω); the quantum field theory here will be about maps $S^1 \to X$. We begin with a "classical phase space" $\mathrm{Map}(S^1, X)$; to a closed, oriented 2-manifold Σ, we should associate the number of pseudoholomorphic maps $u\colon \Sigma \to X$. There's a lot to define here; what is a pseudoholomorphic map? Defining the number of such maps is also nontrivial; in some settings, there are infinitely many, and we must impose point constraints somehow, which makes the theory feel less topological.

The definition of a pseudoholomorphic map involves a PDE, which will be an interesting thing to dig into.

The theory also has a Lagrangian form. In the Lagrangian form, we instead look at paths in X, rather than loops, though we ask that they end on prescribed Lagrangian submanifolds of X. These are a kind of boundary condition. ◀

Atiyah doesn't go into much more detail about this theory, but Schwarz did (assuming $\omega|_{\pi_2(X)}$ vanishes), and we will discuss this example in detail. Ultimately, $Z(S^1)$ will be $H_*(X)$, and the pair-of-pants is sent to a quantum deformation of the cup product which counts pseudoholomorphic curves — Schwarz proves this with Floer theory, but it also makes contact with Gromov-Witten theory.

Example 1.10 (Chern-Simons theory)**.** There are several different flavors of this next example, a 3-dimensional theory. Pick a Lie group G, maybe compact; the classical phase space associated to a closed surface Σ is the moduli space of flat G-bundles on Σ. This isn't infinite-dimensional, because we imposed that our connections are flat, though the space of all connections is infinite-dimensional. If G is nonabelian, this is nonlinear (i.e. not a vector space).

The Lagrangian functional for this theory is the Chern-Simons functional associated to a connection. There's been plenty of work on this example, from different perspectives not just including TQFT, including work by Jones, Witten, Casson, Johnson, and Thurston. ◀

Example 1.11 (Floer theory/Donaldson theory)**.** This is a 4-dimensional example, in which the invariant assigned to a closed 4-manifold X is the Donaldson polynomials on $H_2(M)$ (a tool encoding all of the Donaldson invariants). Atiyah doesn't say what we should do with cobordisms, but fora closed 3-manifold Y, following the Hamiltonian perspective in physics, one should do Floer theory for the Chern-Simons functional on Y (for some Lie group that you have to pick — though only $G = \mathrm{SU}_2$ and $G = \mathrm{U}_2$ have really been worked out, which is Donaldson theory).

Unfortunately, this cannot be an oriented theory — Donaldson polynomials depend on more data. ◀

Awesomely, Atiyah ends with the question *why does the Chern-Simons functional appear in both the three- and four-dimensional cases?* There ought to be an answer in terms of extended TQFT: Chern-Simons theory really seems to be about dimensions 4, 3, and 2.

Example 1.12. After Atiyah's paper came out, Seiberg-Witten theory appeared, as a variant of Example 1.11, and it should fit into a TQFT framework in the same way. This is again a 4-dimensional theory. ◀

We will begin by digging into Example 1.9. Pseudoholomorphic curves are a huge subject; good references include Salamon's lecture notes and the book of Audin-Damian, which is very detailed but doesn't illustrate the analysts' perspective as well as Salamon. The big book of McDuff-Salamon is also good. The professor also has a survey paper, "Lagrangian boundary conditions for anti-self-dual instantons and the Atiyah-Floer conjecture," which is a good way to get an overview of this perspective.

Before we get into pseudoholomorphic curves, here's an important convention: *when we say "symplectic manifold," we always mean closed* (compact and without boundary).

Definition 1.13. A *symplectic manifold* (X, ω) is a manifold X and a 2-form $\omega \in \Omega^2(X)$ which is closed and *nondegenerate*, i.e. $\omega^{\wedge n}$ is a volume form.

This immediately implies $\dim X = 2n$, and is in particular even; and $[\omega] \neq 0$ in $H^2_{\mathrm{dR}}(X)$, which rules out, e.g., S^4.

You can get through a good part of the course thinking of these as even-dimensional manifolds with a particular functional on them. Let $\mathcal{L}X := \mathrm{Map}(S^1, X)$, the unbased loop space of X.

Definition 1.14. The *symplectic action functional* associated to a symplectic manifold (X, ω) is the functional $\mathcal{A}\colon \mathcal{L}X \to \mathbb{R}$ sending a loop $\gamma\colon S^1 \to X$ to the number

$$\int_{[0,1]\times S^1} u^*\omega. \tag{1.15}$$

Here $u\colon [0,1]\times S^1 \to X$ a smooth map with $u(0,-)$ a fixed reference loop u_0 and $u(1,-)=\gamma$.

Often, u_0 is constant, in which case this is choosing a disc whose boundary is γ. There are issues defining this, so the actual target is $\mathbb{R}$ modulo the possible values of ω on tori. If you want to study all of $\mathcal{L}X$, you need to fix a basepoint in each connected component (homotopy class), though often people only study the connected component containing the constant loops, as Floer did.

Given a nice functional, one should want to try gradient flow and Morse theory with it, even though $\mathcal{L}X$ is infinite-dimensional; we will see the definition of a pseudoholomorphic curve pop out naturally from this definition. We will also do Morse theory with the Chern-Simons functional. Doing Morse theory with a function valued in a circle is a bit different, but we'll be able to do it. And in fact, it's the reason we work with the Novikov ring.

2. : 2/21/20

3. 2D TFTs from symplectic manifolds: 2/26/20

We will spend the first part of class carefully setting up a precise statement to the following theorem.

Theorem 3.1 (Schwarz, Floer)**.** *Let (M,ω) be a symplectic manifold such that either $\omega|_{\pi_2(M)}=0$ or $\omega=\lambda c_1(M)$, with $\lambda>0$. Then thee is a TQFT $Z\colon \mathcal{C}ob_{\langle 2,1\rangle}\to \mathcal{V}ect_{\mathbb{F}_2}$ with $Z(S^1)\cong H_*(M)$.*

Here (TODO: I think) $c_1(M)$ is measured in any compatible almost complex structure for the symplectic form; the choice doesn't matter.

Remark 3.2.

- The algebra structure on $Z(S^1)$ is not just the usual intersection product; it's deformed by counting pseudoholomorphic curves.
- We can relax the niceness assumptions on the symplectic form, but then our target category is modules over some universal thing called the *Novikov ring*. We're not going to delve into this. ◀

Somehow the existence of this TQFT is not the entire point; instead, it leads us to interesting analytic and geometric questions.

Choose a map $H\colon S^1\times M\to\mathbb{R}$ such that the time-1 flow

$$\{(p_0,p_1)\mid \text{ there exists a }\gamma\colon[0,1]\to M, \dot\gamma(t)=X_{H_t}(\gamma(t)), \gamma(0)=0,\gamma(1)=1\}\subset M\times M \tag{3.3}$$

is transverse to the diagonal $\Delta_M \coloneqq \{(p,p)\mid p\in M\}$. Such a map induces a $\mathbb{Z}/2N$-graded complex $CF_*(H)$ generated by periodic loops of X_H, where N is the minimal positive value of $\langle c_1(TM,J),[S^2]\rangle$ over all embeddings $S^2\hookrightarrow M$ — pseudoholomorphic or not.[3]

Let $\mathcal{J}(M,\omega)$ denote the space of compatible almost complex structures on M for ω. Given the map H above, there is a comeager subset $S\subset \mathrm{Map}(S^1,\mathcal{J}(M,\omega))$ for which each $J\in S$ induces a differential $\partial\colon CF_*(H)\to CF_{*-1}(H)$; in particular, $\partial^2=0$. Fixing these two choices, we can define $Z(S^1)$ to be the homology of this chain complex.

This differential arises by an analogue to Morse theory on an infinite-dimensional Banach manifold, though instead of counting curves, we count solutions to a PDE (only counting them in the case where there's a zero-dimensional moduli space). Specifically, we consider the space of solutions

$$\{u\colon \mathbb{R}\times S^1\to M\mid u(\pm\infty,-)=\gamma_\pm, \overline{\partial}_J u = X_H\}, \tag{3.4}$$

which $\mathbb{R}$ acts on freely by time translations (just as in Morse theory); the moduli space $\mathcal{M}(\gamma_-,\gamma_+)$ is the quotient of (3.4) by this $\mathbb{R}$-action. Here $\gamma_\pm$ are choices of points in M, so that we are counting pseudoholomorphic strips in M which at infinity converge on $\gamma_\pm$.

Remark 3.5. Why is the Hamiltonian so complicated? You can build TQFTs with simpler Hamiltonians (akin to doing Morse homology with simpler Morse functions). But this example doesn't come from nowhere: Floer used these techniques to solve a piece of the Arnold conjecture. Other confusing-sounding choices sometimes also come from geometric applications. ◀

[3]The magic of symplectic geometry is how much we can do without actually knowing what the pseudoholomorphic curves actually are — or if there are any at all!

The local dimension of $\mathcal{M}(\gamma_-,\gamma_+)$ is $\deg(\gamma_-)-\deg(\gamma_+)-1$, where the degrees are in the grading of $CF_*(H)$ we discussed above; the factor of -1 arises because we quotiented by the $\mathbb{R}$-action. We will stick to the cases when this dimension is 0.

The proof that the differential squares to zero follows the same line of reasoning as in ordinary Morse theory, though it looks fancier in this setting. In Morse theory, a flow line can break into two. Here, we consider the moduli space of pseudoholomorphic strips from γ_- to γ_+, where $\deg(\gamma_-)-\deg(\gamma_+)=2$. A great deal of analysis goes into showing this is a smooth 1-manifold with ends, and some of the assumptions we made in the theorem statement eliminate some unsavory possibilities (e.g. bubbling). Anyways, we compactify, to obtain a 1-manifold with boundary, and show that it factors as

$$\coprod_{\gamma:|\gamma|=|\gamma_-|-1}\mathcal{M}(\gamma_-,\gamma)\times\mathcal{M}(\gamma,\gamma_+). \tag{3.6}$$

The high-level idea for what's happening is that the pseudoholomorphic strip breaks into two. There are many other situations in gauge theory or infinite-dimensional Morse theory where breaking (and sometimes bubbling) can happen.

Compactness is another important ingredient, and it's also fundamentally analytic. Given a sequence $\{u_n\}\subset\mathcal{M}(\gamma_-,\gamma_+)$, there is a subsequence that converges in $C^\infty_{\ell oc}$, unless the energy of the sequence behaves badly. The energy is a functional

$$E(u)=\int_{\mathbb{R}_s\times S^1}|\partial_s u|^2=\mathcal{A}(\gamma_-)-\mathcal{A}(\gamma_+). \tag{3.7}$$

The Morse-theoretic version of this is

$$\int_{\mathbb{R}}|\nabla(f(\gamma))|^2=f(p_-)-f(p_+), \tag{3.8}$$

provided $\gamma\colon\mathbb{R}\to X$ is a gradient flow line for the Morse function; in Floer theory, $\partial_s u=-\nabla\mathcal{A}_H u$.

The name "energy" is because $|\partial_s u|^2$ is always locally positive, so we can think of it as an energy density. So before we even ask about convergence of a subsequence, we can ask how the energy behaves. It can concentrate near a point in the cylinder, or it can separate, piling up near the ends of $\mathbb{R}$. Separating is good, in that it leads to breaking, which we thought about in the differential. But concentration is often trickier to deal with: it leads to bubbling, which is a PDE effect, and which we'll largely sweep under the rug.

The next big theorem we need is a gluing theorem. This will rule out pathological broken trajectories in which the ends aren't actually different (TODO: I think?). There is a gluing map

$$g\colon\coprod_\gamma\mathcal{M}(\gamma_{-1},\gamma)\times\mathcal{M}(\gamma,\gamma_+)\times(R_0,\infty)\longrightarrow\mathcal{M}(\gamma_-,\gamma_+). \tag{3.9}$$

where the disjoint union is over the γ for which these moduli spaces are zero-dimensional, and R_0 is some number. The idea is that for every broken trajectory, we embed an interval into the manifold, which chooses which end you converge to.

Theorem 3.10 (Gluing theorem (Schwarz)). *The gluing map g is an homeomorphism onto its image, and $\mathcal{M}(\gamma_-,\gamma_+)\setminus\mathrm{Im}(g)$ is compact.*

So the moduli space without these bad examples is nice. This is a key result from Schwarz's thesis.

Often, papers will only think about one of compactness or gluing, and sketch the proof of the other; this is where the mistakes creep in, so be careful.

Anyways, now we have the Floer complex. Choose $a,b,g\in\mathbb{Z}_{\ge 0}$, and consider any Riemann surface $\Sigma\coloneqq\Sigma_{g,a+b}$ (i.e. genus g, $a+b$ boundary components), regarded as a bordism with a incoming circles and b outgoing circles. This Σ induces a chain map

$$\Phi(\widetilde{X},\widetilde{J})\colon\bigotimes_{i=1}^a CF_*(H_i)\longrightarrow\bigotimes_{j=1}^b CF_*(H_j), \tag{3.11}$$

where H_i, H_j, $\widetilde{X}$, and $\widetilde{J}$ are data we haven't discussed yet, but will induces the same map when we take homology. This chain map has degree $(\dim M)/2(2g-2+a+b)$, and will be what the TQFT assigns to the bordism Σ.

How do we define this chain map? Again, it counts something, which is the number of points in a zero-dimensional moduli space of pseudoholomorphic maps $\Sigma \to M$, such that the boundary circles of Σ map to specified loops γ_i^-, γ_j^+ in M. There are again choices to make, including disjoint embeddings $\varphi_i\colon (-\infty, 0) \hookrightarrow \Sigma$ and $\varphi_j\colon (0,\infty) \hookrightarrow \Sigma$, which give us cylindrical coordinates near each boundary circle; and a choice of a complex structure on the interior of Σ which is the standard (cylinder) complex structure on the ends $\mathbb R\times S^1 \cong \mathbb C/\mathbb Z$. We also choose (H_i, J_i) on each end, and $\widetilde X, \widetilde J$ are generic (comeager subset) of a certain 1-form $\widetilde X$ with $\varphi_i^*\widetilde X = X_{H_i}\,\mathrm dt$, and $\widetilde J$ a map from the interior of Σ into $\mathcal J(M,\omega)$, such that $\varphi_i^*\widetilde J = J_i$. All of these are contractible choices, which is reassuring. Some thought has to go into ensuring this is a well-posed PDE.

Again, we want to make a count of a zero-dimensional moduli space, so the same questions come up: is this space zero-dimensional? Is it compact? And so on. Then, to verify that it's a chain map, we need to know that it commutes with the differentials. This again arises by studying the ends of the one-dimensional moduli spaces – this should correspond to breaking of pseudoholomorphic strips, caused by energy running out at the ends. Breaking can happen at each boundary circle of Σ, incoming or outgoing, and these give you the different components of the boundary of the moduli space.

This TFT involves a whole bunch of choices — we will continue with three more theorems that imply the TFT, on homology, is independent of choices. This is a common method of proof in this area, and it's difficult to avoid. Find your way of understanding how these diagrams and pictures work.

4. Schwarz theory and Gromov-Witten theory: 2/28/20

Today, we continued discussing the 2d TFT from yesterday's discussion, built from an input data of a (nice) symplectic manifold (M,ω) by counting pseudoholomorphic curves. One important point is that the reason this works is really because the pseudoholomorphic curve equation is a very nice PDE: in particular, it's elliptic, which severely constrains the places where energy can pile up, and limits the ways in which things go wrong. For example, this rules out bubbling when you define the map assigned to a bordism Σ (which, roughly, is determined by a count of pseudoholomorphic curves $\Sigma \to M$ with prescribed image of $\partial\Sigma$). Making this work involves a whole bunch of analysis.

This is not to say that nothing can go wrong! We'll see some examples in a bit.

Remark 4.1. We made a lot of choices in defining the data assigned to a surface in this TFT. One must check this is independent of choices, and it probably isn't yet clear why that should be true. ◀

When Σ is closed, we're couting pseudoholomorphic maps $u\colon \Sigma\to M$ such that $\overline\partial_{\widetilde J} u = \widetilde X$. In this case, $\widetilde X = 0$ is a reasonable choice, probably; maybe we have to move it slightly to satisfy a genericity assumption.

This looks very similar to the definition of Gromov-Witten moduli spaces, but there are a few differences. In Gromov-Witten theory, J is fixed on M, and can't vary with the domain. This means that if Σ has any holomorphic automorphisms, we can reparameterize the space of pseudoholomorphic maps $\Sigma\to M$. In Gromov-Witten theory, we quotient out by these automorphisms, so we have a space of curves in M, rather than a space of maps. In general, one also varies the (almost) complex structure j on Σ, so we begin with the space of data $(\Sigma, j, u\colon \Sigma\to M)$ and cut out the subspace with $\overline\partial_{J,j} u = 0$, then quotient out by automorphisms $\varphi\colon \Sigma'\to\Sigma$ with $\varphi^* j = j'$. In genus 0, this additional varying direction doesn't change anything, because uniformization relates all complex structures, but it can be different in higher genus.[4]

As always, when we say "count," we mean that we consider only the zero-dimensional part (so we're considering spaces of solutions, rather than moduli spaces *per se*). For example, if $\Sigma = S^2$, then the automorphism group is the Möbius group, which is six-dimensional, and therefore the difference between the dimensions of the Schwarz moduli space and the Gromov-Witten moduli space is equal to 6. In some sense, for Schwarz, symmetries are not a problem, but a computational tool: $\mathrm{Aut}(\Sigma)$ sometimes allows us to explicitly determine the map assigned to Σ in the TFT.

Example 4.2. Say $\Sigma = T^2$. In this case, the space of solutions is zero-dimensional when $c_1(u^*TM) = 0$, e.g. if $u_*[\Sigma] = 0$ in $H_2(M)$. This in particular means the energy of u is zero, so $\int \lvert \mathrm du\rvert^2 = 0$, so u is constant!

Assuming for now that all of these constructions are independent of choices (they are, but we haven't discussed why yet), choose $H\colon S^1\times M\to\mathbb R$ and $J = \widetilde J\colon S^1\times TM\to TM$, which are "regular" in that

[4] This has been a *very* quick introduction to Gromov-Witten theory; there are long books going into more detail!

the space of solutions is actually zero-dimensional (this is true for generic choices of H and J). Assume $\widetilde{X} = X_H\,\mathrm{d}t + \dots \mathrm{d}s$, where s and t are the coordinates on the torus. Now the pseudoholomorphic curve equations $\overline{\partial}_{\widetilde{J}} u = \widetilde{X}$ simplifies slightly to

$$\partial_s u + J\partial_t u = JX_H(u). \tag{4.3}$$

This space has an S^1-symmetry by reparameterizing the first coordinate. This is an example of a general truism in PDE: if the solutions of a PDE doesn't depend on some parameter in the input data, we pick up a symmetry given by reparameterizing that parameter.

So given a solution, either it's constant in the first factor, or it comes with a whole S^1-orbit of solutions. The latter case is not zero-dimensional, so $Z(T^2)$ is the number of fixed points of the S^1-action, which is precisely the number of pseudoholomorphic curves u with $\partial_s u = 0$. This also counts the periodic orbits of X_H.

Since periodic orbits generate $HF_*(H, J)$, this counts the Euler characteristic of $HF_*(H, J)$ (we don't need to worry about signs because this was defined over $\mathbb{F}_2$ from the start). ◂

This is a reassuring sanity check, because for any 2d TFT, $Z(T^2) = \chi(Z(S^1))$: factor $Z(T^2)$ as two cylinders, the first with two outgoing components and zero incoming components, and the second with two incoming components and zero outcoming components.[5]

Exercise 4.4. The axioms of TFT ensure that for the cylinder C, $Z(C) = \mathrm{id}$. Can you show that directly for this Schwarz-Floer TFT? Hint: use symmetries to compute things, both by translation in the $\mathbb{R}$ direction and by rotation in the S^1 direction.

Before we discuss why things are independent of choices, let's briefly discuss Gromov-Witten invariants, at the level of an executive summary; this is also useful towards understanding Donaldson invariants.

The Gromov-Witten moduli space $\mathcal{M}_{g,k}(M, A, J)$ has a lot of decorations (which are mostly discrete parameters). This space is a moduli space for pseudoholomorphic maps $u\colon \Sigma \to M$, where Σ has genus $g$and k marked points (and where the marked points are constrained to hit certain points of M), $A \in H_2(M)$ is $u_*[\Sigma]$, and $J \in \mathcal{J}(M, \infty)$.

Precisely, this moduli space is the space of data $(\Sigma, j, m_1, \dots, m_k, u)$, where Σ is a genus g surface with k marked points $z_1, \dots, z_k$, $u\colon \Sigma \to M$ satisfies $\overline{\partial}_{(j,j)} u = 0$, and $u_*[\Sigma] = A$. This has symmetries, so we quotient out by these symmetries; then there is a nice compactification $\overline{\mathcal{M}}_{g,k}(M, A, J)$ in which we also allow nodal curves.

This space has expected dimension

$$2c_1(A) + (\dim M - 6)(1 - g) + 2k, \tag{4.5}$$

which is to say that we hope $\overline{\mathcal{M}}_{g,k}(M, A, J)$ is an orbifold of this dimension for generic J (generic in some suitable sense), and that we get a fundamental class. The truth is more complicated in general, but this or something close to it works in a good variety of situations.

To define numerical invariants, we begin with the *evaluation maps*

$$\begin{aligned} ev_i\colon \overline{\mathcal{M}}_{g,k}(M, A, J) &\longrightarrow M \\ (\Sigma, j, z_1, \dots, z_k, u) &\longmapsto u(z_i) \end{aligned} \tag{4.6}$$

and the *forgetful map* $f\colon \overline{\mathcal{M}}_{g,k}(M, A, J) \to \overline{\mathcal{M}}_{g,k}$, which forgets the map u to M; here $\overline{\mathcal{M}}_{g,k}$ is the (compactification of the) modui space of (genus-g, k-marked) Riemannian surfaces, which has dimension $6g - 6 + 2k$. For small g, k, this space has negative expected dimension, and so f is not well-defined there.

Anyways, we can now pull back cohomology classes and define the *Gromov-Witten invariant* of $\alpha_1, \dots, \alpha_k \in H^*(M)$ and $\beta \in H^*(\overline{\mathcal{M}}_{g,k})$ to be

$$GW_{g,k}(\alpha_1, \dots, \alpha_k, \beta) \coloneqq \int_{\overline{\mathcal{M}}_{g,k}(M,A,J)} ev_1^*\alpha_1 \wedge \dots \wedge ev_k^*\alpha_k \wedge f^*\beta. \tag{4.7}$$

In order for this to be nonzero, the sum of the degrees of $\alpha_1, \dots, \alpha_k, \beta$ must equal the expected dimension (4.5) of $\overline{\mathcal{M}}_{g,k}(M, A, J)$. Moreover, in the case where it's not actually that dimension, workarounds are taken to define this integration map, and hopefully they apply to the case you care about.

[5] The Euler characteristic really only makes sense when this TFT is valued in graded vector spaces; for ungraded vector spaces, we get $Z(T^2) = \dim Z(S^1)$.

Remark 4.8. When you just care about closed surfaces, compactification is most interesting to the enumerative geometers: we need it to define counts, but it doesn't make the analysis or algebraic structure more interesting. On surfaces with boundary, both of these get considerably more interesting when we compactify. ◂

A key word you might've heard: the *quantum cup product* $H^*(M) \otimes HJ^*(M) \to H^*(M)$ sends

$$\alpha_1, \alpha_2 \longmapsto GW_{0,3}(\alpha_1, \alpha_2, -) \tag{4.9}$$

This is a functional on cohomology classes, which we identify with a cohomology class by Poincaré dualty. We can think about it as dealing with two curves with homology classes α_1 and α_2.

Hidden in this is the interesting question *how does this quantum cup product relate to what Schwarz' TFT assigns to the pair of pants?* This is known as the *PSS isomorphism.*

Finally, we should say something about why this TFT, and all the data that goes into it, does not depend on the choices we made. First, all of the choices we made are contractible (i.e. the space of possible choices is contractible). So given $a, b, g \in \mathbb{Z}_{\geq 0}$, any isotopy of data $(\widetilde{X}_\lambda, \widetilde{J}_\lambda, j_\lambda, \varphi_\lambda)$ (where φ is the data on the ends), for $0 \leq \lambda \leq 1$, will induce a chain homotopy equivalence $\Psi\colon \bigotimes_a CF_* \to \bigotimes_b CF_*$, which suffices when we pass to homology — but we also need to check that what we assign to bordisms doesn't depend on choices. Here, the statement is that the difference

$$\Phi(\widetilde{X}_1, \widetilde{J}_1, \dots) - \Phi(\widetilde{X}_0, \widetilde{J}_0, \dots) = \sum_{j=1}^{b} \partial_j \circ \Psi + \sum_{i=1}^{a} \Psi \circ \partial_i, \tag{4.10}$$

which means the induced maps on homology are equal. Actually proving this involves a lot of PDE, of course, but the idea is similar to the proofs we sketched last lecture: Ψ counts pairs (λ, u) with $\lambda \in [0,1]$ and u a solution for $(\widetilde{X}_\lambda, \widetilde{J}_\lambda, \dots)$. The boundary of this space corresponds to $\lambda = 0, 1$; one has to think about breaking. If you'd like to read the proof, you can find it in Schwarz's thesis.

There is yet more data in the TFT: the composition/gluing identity

$$Z(\Sigma \cup_{S^1} \Sigma') = Z(\Sigma) \circ Z(\Sigma'). \tag{4.11}$$

So we need to start with two sets of data of complex structure, ends, collars, ... which induce composable chain maps

$$\begin{tikzcd} \bigotimes_i CF_*(H_i) \arrow[r, "\Phi_\Sigma(\widetilde{X}, \widetilde{J})"] \arrow[rr, bend right, "\Phi_{\Sigma \cup_{S^1} \Sigma'}(\widetilde{X} \#_R \widetilde{X}, \widetilde{J} \#_R \widetilde{J}')"'] & \bigotimes_j CF_*(H_j) \arrow[r, "\Phi_{\Sigma'}(\widetilde{X}', \widetilde{J}')"] & \bigotimes_k CF_*(H_k). \end{tikzcd} \tag{4.12}$$

We would like this to be a commutative diagram, at least for R (the radius of the circles we're gluing along) sufficiently large. This is also due to Schwarz.

Here's a final fact due to Floer: if H is time-independent and small in the C^2-norm, and if J is also S^1-independent, then (CF_*, ∂) is the Morse complex of H. This in particular means that we can replace $Z(S^1) = HF_*(H_t) = HF_*(H)$, the homology of the Morse complex for H, which is the standard homology of M (we've done everything in the mod 2 case, so we get $H_*(M; \mathbb{Z}/2)$). This is not an easy theorem.

Corollary 4.13 (Weak Arnold conjecture)**.** *The number of generators of $HF_*(H_t)$ is at least the sum of the (mod 2) Betti numbers of M. Said generators are the periodic orbits of H_t.*

This is why Floer worked with S^1-time-varying Hamiltonians; you don't need that to do TFT, though, but Arnold's conjecture is a really great conjecture; important in dynamics, but also fostering the development of symplectic topology and Floer theory, leading people to investigate the rich algebraic structures in place. There's plenty more than periodic orbits — they're a relatively simple beginning subject, and it's a little curious that we bounded them in topological quantities, albeit using a massive PDE! This is why symplectic geometers are interested in algebraic structures that come from J-holomorphic curves: they give us direct geometric insights.

5. Gluing in Schwarz-Floer theory: 3/4/20

Today, we'll continue discussing gluing data in the Schwarz-Floer TQFT built from counts of pseudoholomorphic maps into a symplectic manifold (X, ω).

In TQFT, we're just looking at smooth surfaces, where we identify bordisms which are diffeomorphic rel boundary, so there's no choices to be made in gluing. But when defining $Z(\Sigma)$ for a bordism Σ, we chose a lot of additional data: a map $\widetilde{J}\colon \Sigma \to \mathcal{J}(M,\omega)$, a complex structure on Σ, and collars for it on the boundary components of Σ, and $\widetilde{X}$ – so that we count solutions to $\overline{\partial}_{\widetilde{J}} u = \widetilde{X}$. When we glue Σ and Σ' across some boundary components C, the collars help us glue $\widetilde{J}$ and $\widetilde{J}'$.

Again, this is different from Gromov-Witten theory, in which the complex structure on Σ is allowed to vary in defining the moduli space.

We need to define a complex structure on $\Sigma \cup_C \Sigma'$, and we have to make a choice. Choose $R \gg 0$, which allows us to identify

$$\Sigma \supset [0,\infty) \times S^1 \supset [0,R] \times S^1 \xrightarrow{\cong} [-R,0] \times S^1 \subset (-\infty,0] \times S^1 \subset \Sigma'. \tag{5.1}$$

Because we specified that in this cylindrical neighborhood, the complex structures on Σ and Σ' are biholomorphic to the standard ones on the cylinder glue to a single complex structure on the glued surface, and we can define the PDE. This complex structure depends on R, and we'd like to say the count of solutions to this PDE is independent of R, and this is true when R is sufficiently large, which is good enough. This field is full of constructions where one must choose a lot of data and later prove that (maybe after passing to homology) what you get is independent of choices.

Remark 5.2. We also chose a distance $R \gg 0$ when considering how paths break. That R is not the same as this R: they appear in the theory for different purposes, so don't conflate them. ◂

Great, we've glued the complex structure on $\Sigma \cup_C \Sigma'$ using a fixed large R. What about everything else? First, let's glue maps $u\colon \Sigma \to M$ and $v\colon \Sigma' \to M$. Define the *pregluing* of u and v, called $u \#_R v$, to be u on Σ minus its ends, v on Σ' minus its ends, and to interpollate between them on $[0,R] \times S^1 \sim [-R,0] \times S^1$. One can show that

$$\overline{\partial}_{\widetilde{J}''}(u \#_R v) - \widetilde{X} \approx 0, \tag{5.3}$$

and is only nonzero in the region in which we interpolated. If you're interested in analytic aspects, this is why exponential decay is so crucial. Now we use Newton's method! Given a nice enough function and an almost zero, there is an algorithm for finding an actual zero. This is a common technique for gluing in PDE, though one must prove some estimates to show it converges.[6] But in this setting it does, and does so in a unique enough manner, so we obtain a gluing map of moduli spaces

$$\varrho_R\colon \mathcal{M}(\Sigma,\dots) \times \mathcal{M}(\Sigma',\dots) \longrightarrow \mathcal{M}(\Sigma \cup_{C,R} \Sigma',\dots); \tag{5.4}$$

first compute $u \#_R v$, then apply Newton iteration to obtain a solution w_R.

Finally, we need a compactness theorem: for $R \to \infty$, w_R solves

$$\overline{\partial}_{(\widetilde{J}''_R, j''_R)} w_R = \widetilde{X}''_R, \tag{5.5}$$

or, equivalently, (in a certain sense that must be made precise) $w_R \to (u,v)$ as $R \to \infty$. Thus, for sufficiently large R, all w_R in the moduli space for $\Sigma \cup_{C,R} \Sigma'$ are in the image of ϱ_R.

The details get lengthy and involved, but when you're reading a paper providng a gluing formula you should expect to see at least these steps.

Thinking about how to extend this Schwarz-Floer theory to a point leads one to consider Lagrangian boundary conditions, which will be useful; but before we discuss this let's dig into another example: the instanton TFT.

Instantons. We will discuss how to construct a TFT on part of the 4-dimensional bordism category, which assigns to a closed 3-manifold the Floer homology of its Chern-Simons functional.

At least for now, all bundles will be trivial, so we will think of $Y \times \mathbb{C}^2$ with the standard SU_2-action.[7] This means that we can identify the space $\mathcal{A}(Y)$ of connections with $\Omega^1_Y(\mathfrak{su}_2)$.

[6] Newton's method, stated in complete generality with all the analytic detail, can be found in an appendix of McDuff-Salamon. The analytic details are important.

[7] All SU_2-bundles on a 3-manifold are trivial, but we go further and trivialize them.

Definition 5.6. Choose a connection $A \in \mathcal{A}(Y)$. The *Chern-Simons one-form* $\lambda_A \colon T_A\mathcal{A}(Y) \to \mathbb{R}$ is

$$\alpha \longmapsto \int_Y \operatorname{tr}(F_A \wedge \alpha). \tag{5.7}$$

The *Chern-Simons functional* $CS\colon \mathcal{A}(Y) \to \mathbb{R}$ is the primitive of the Chern-Simons 1-form.

The Chern-Simons functional is gauge-invariant up to multiples of $4\pi^2$; that is, under an action of $\mathcal{G}(Y) \coloneqq C^\infty(Y, G)$ on $\mathcal{A}(Y)$ by

$$(u, A) \longmapsto u^*A \coloneqq u^{-1}Au + u^{-1}\,\mathrm{d}u, \tag{5.8}$$

it only changes by multiples of $4\pi^2$.

Recall also the curvature formula

$$F_A = \mathrm{d}A + \frac{1}{2}[A \wedge A]. \tag{5.9}$$

There are alternate definitions of the Chern-Simons functional, such as

$$CS(A) = \frac{1}{2}\int_Y \operatorname{tr}\Big(A \wedge \Big(F_A - \frac{1}{6}[A \wedge A]\Big)\Big) \tag{5.10}$$

$$= \frac{1}{2}\int_X \operatorname{tr}\Big(F_{\widetilde{A}} \wedge F_{\widetilde{A}}\Big), \tag{5.11}$$

where X is a compact, oriented 4-manifold with an identification $\varphi\colon \partial X \xrightarrow{\cong} y$ and $\widetilde{A}$ is an SU_2-connection on X with $\widetilde{A}|_{\partial X} = \varphi^*A$, i.e. we've chosen a manifold and connection which Y and A bound.[8] These definitions involve some choices, most notably the bulk 4-manifold X, and these choices can change $CS(A)$, but only by multiples of $4\pi^2$. For example, if X' is a closed, oriented 4-manifold, we could imagine choosing $X \amalg X'$ as the bulk, and a complex rank-2 vector bundle $E \to X'$ with an SU_2-structure; then,

$$\int_{X'} \operatorname{tr}(F_{\widetilde{A}} \wedge F_{\widetilde{A}}) = 4\pi^2\langle c_2(E), [X]\rangle \tag{5.12}$$

and $\langle c_2(E), [X]\rangle$ is an integer.

Now let's briefly discuss instanton Floer homology; two good readable sources are Floer's and Donaldson's treatments. This is the Morse theory of the Chern-Simons functional, interpreted as a function

$$CS\colon \mathcal{A}(Y)/\mathcal{G}(Y) \longrightarrow \mathbb{R}/4\pi^2\mathbb{Z}. \tag{5.13}$$

We will obtain a chain complex CF_* generated by critical points of the Chern-Simons functional. Well, actually, we need to take a small perturbation of this functional. The differential counts gradient flow lines for connections on $\mathbb{R} \times Y$, modulo gauge.

The first thing we need to do is compute the gradient of the Chern-Simons functional. Choose a Riemannian metric on Y, and define

$$\lambda_A(\alpha) \coloneqq g_{L^2}(\nabla CS, \alpha) = \int_Y \operatorname{tr}(\nabla CS \wedge \star\alpha)), \tag{5.14}$$

where here $\star\colon \Omega^1_Y \to \Omega^2_Y$ is the Hodge star. In this degree, the Hodge star squares to 1. Now, we can also identify (5.14) with

$$\int_Y \operatorname{tr}(F_A \wedge \alpha) = \int_Y \operatorname{tr}(\alpha \wedge F_A) = \int_Y \operatorname{tr}(\alpha \wedge \star\nabla CS), \tag{5.15}$$

so $\nabla CS(A) = \star F_A \in \Omega^1_Y(\mathfrak{su}_2)$. The critical points are therefore the flat connections, and the space is flat connections modulo gauge. Flat SU_2-connections can be identified with $\operatorname{Hom}(\pi_1(Y), \mathrm{SU}_2)$, and gauge transformations act by conjugation, so the space of critical points is identified with the character variety

$$\mathcal{M}^\flat(Y, \mathrm{SU}_2) = \operatorname{Hom}(\pi_1(Y), \mathrm{SU}_2)/\mathrm{SU}_2. \tag{5.16}$$

[8] There is a question in algebraic topology here, which is: does the bordism group $\Omega_3^{\mathrm{SO}}(B\mathrm{SU}_2)$ of closed, oriented 3-manifolds with a principal SU_2-bundle vanish, so that X and $\widetilde{A}$ can always be chosen? The answer is yes; as mentioned above, all principal SU_2-bundles on 3-manifolds are trivial, so this reduces to the classical fact that all closed, oriented 3-manifolds bound.

This is generally not finite, which is why we have to perturb the equation slightly. The upshot is that CF_* is generated by some chain complex whose standard differential computes the homology of $\mathcal{M}^\flat(Y, \mathrm{SU}_2)$, but the boundary map in CF_* counts maps $B\colon \mathbb{R} \to \mathcal{A}(Y)$ such that $\partial_s B = -\star F_B$ modulo paths $\mathbb{R} \to \mathcal{G}(Y)$.

There's an equivalence

$$\{\partial_s B = -\star F_B\}/\{\mathbb{R} \to \mathcal{G}(Y)\} \cong \{F_{\widetilde{A}} + \star F_{\widetilde{A}} = 0 \mid \widetilde{A} \in \mathcal{A}(\mathbb{R} \times Y)\}/\mathcal{G}(\mathbb{R} \times Y). \tag{5.17}$$

The equation $F_{\widetilde{A}} + \star F_{\widetilde{A}} = 0$ is called the *anti-self-dual equation*, and its solutions are called *anti-self-dual instantons*. Here $F_{\widetilde{A}} \in \Omega^2$ of a 4-manifold, and the Hodge star squares to the identity in this degree. Thus even by starting with a three-dimensional question, we ended up with something interesting for arbitrary 4-manifolds, and Donaldson invariants are exactly about this PDE, but on closed 4-manifolds.

There's also the related *self-dual equation* $F_A - \star F_A = 0$. Both this and the anti-self-dual equation imply the *Yang-Mills equation* $d_A^* F_A = 0$. The spaces of solutions to the self-dual and anti-self-dual equations generally are not zero-dimensional, but Donaldson invariants are defined as certain integrals over these spaces, which are in cases of interest nice orbifolds.

6. Donaldson invariants and a little history of 4-manifolds: 3/6/20

4-manifold topology is a particularly interesting place to study the classification of manifolds: it's the lowest dimension in which the classification of topological and smooth manifolds differs. Today, all manifolds will be understood to be closed and oriented. We follow Donaldson-Kronheimer, whose exposition is concise and elegant; there are also less concise sources.

The simplest invariants one can extract from such a 4-manifold X are its homology and its *intersection form*

$$Q_X\colon H_2(X;\mathbb{Z}) \times H_2(X;\mathbb{Z}) \longrightarrow \mathbb{Z}, \tag{6.1}$$

described as follows: given two homology classes, represent them by embedded closed, oriented submanifolds A and B which are transverse, and then count $\#(A \cap B)$; that this is finite and independent of choices follows from standard theorems in differential topology. Poincaré duality implies we can equivalently describe Q_X as the map $H^2(X;\mathbb{Z}) \times H^2(X;\mathbb{Z}) \to \mathbb{Z}$ sending

$$\alpha, \beta \longmapsto \int_X \alpha \wedge \beta. \tag{6.2}$$

Here, the integral is suggesive notation for evaluating on the fundamental class of X. Using Poincaré duality, one can show this is a unimodular, symmetric bilinear form.

Sometime between Whitney and Milnor, it was shown that if X is simply connected, the homology and the intersection form of X determine its homotopy type. Conversely, Freedman showed in the early 1980s that *all* unimodular, symmetric bilinear forms appear as the intersection form of some topological 4-manifold. Both of these are large results, and the upshot is essentially a classification of simply connected topological 4-manifolds.

Only a year after Freedman's result, though, Donaldson proved a remarkable theorem restricting the possible intersection forms of smooth manifolds.

Theorem 6.3 (Donaldson)**.** *Let X be a simply connected, smooth 4-manifold whose intersection form Q_X is definite. Then, Q_X is diagonalizable.*

You think, sure, every symmetric matrix is diagonalizable, but this is about diagonalizability over $\mathbb{Z}$, which is not guaranteed. For example, the E_8 intersection form

$$\mathrm{E}_8 = \begin{pmatrix} 2 & 1 & & & & & & \\ 1 & 2 & 1 & & & & & \\ & 1 & 2 & 1 & & & & \\ & & 1 & 2 & 1 & & & \\ & & & 1 & 2 & 1 & & 1 \\ & & & & 1 & 2 & 1 & \\ & & & & & 1 & 2 & \\ & & & & 1 & & & 2 \end{pmatrix}. \tag{6.4}$$

is positive definite, but not diagonalizable, and therefore the E_8 manifold X_{E_8}, which is a topological 4-manifold whose intersection form is the E_8-form, is not smooth. This was in fact already known, using something called the *Rokhlin invariant*, but that didn't address other manifolds, such as $X_{\mathrm{E}_8} \# X_{\mathrm{E}_8}$, whose smoothability is ruled out by Donaldson's theorem — and in fact, there is no smooth 4-manifold with intersection form $\mathrm{E}_8 \oplus \mathrm{E}_8$.

The theorem statement is exciting, and so is its proof! We begin with topology, but take a turn into analysis: the proof concerns itself both with how the analysis works and how it doesn't work.

Proof sketch of Theorem 6.3. Choose a Riemannian metric on X, and a principal SU_2-bundle $P \to X$ with $\langle c_2(P), [X]\rangle = 1$. Let $A \in \mathcal{A}(P)$ and F_A denote its curvature; then A is *anti-self-dual* if $F_A + \star F_A = 0$. Let $\mathcal{M}(P)$ denote the space of anti-self-dual connections modulo the action of the gauge group $\mathcal{G}(X)$. For "regular" metrics on X (which is a generic enough condition for everything to be OK), $\mathcal{M}(P)$ is (almost!) a compact manifold of dimension

$$8\langle c_2(P), [X]\rangle - 3(1 - b_1(X) + b_2^+(X)) = 5, \tag{6.5}$$

because we've constrained $c_2(P)$ and assumed X is simply connected. The "almost!" is actually a very good thing — it means that there can be singularities in $\mathcal{M}(P)$ at reducible connections, and can be bubbling at points of X. The upshot is that if $\widetilde{\mathcal{M}}(P)$ is $\mathcal{M}(P)$ minus a small neighborhood around these bad points,

$$\partial\widetilde{\mathcal{M}}(P) \cong X \amalg \coprod_{i=1}^{k} \mathbb{CP}^2, \tag{6.6}$$

with the $\mathbb{CP}^2$s coming from the singularities and the copy of X coming from bubbling. This is a good thing: we can compactify $\widetilde{\mathcal{M}}(P)$ into a compact, oriented 5-manifold, so that it exhibits a bordism from X to some number of $\mathbb{CP}^2$s. One then can use this to show that Q_X is equal to the direct sum of the intersection forms on the $\mathbb{CP}^2$s, each of which is just $[-1]$, diagonalizing Q_X.[9] ⊠

The bubbling that was mentioned happens as follows: we know $\int_X |F_A|^2 = 8\pi^2$ because $\langle c_2(P), [X]\rangle = 1$ (using Chern-Weil theory), and for any point $x \in X$, one can find a sequence A_t of anti-self-dual connections whose curvature "piles up at x": $|F_{A_t}|^2 \to 8\pi^2\delta_x$. If one relaxes the assumption that $\langle c_2(P), [X]\rangle = 1$, there can be "bubble trees," which are trickier to deal with; probably this hasn't been fully investigated in Yang-Mills theory yet.

In general, the moduli spaces of anti-self-dual connections, or pseudoholomorphic curves, or Seiberg-Witten monopoles are for suitably generic data (almost!) finite-dimensional closed manifolds, where the dimension is the Fredholm index on a local slice to the symmetry group. The "almost!" includes that there may be isotropy, e.g. finite stabilizer subgroups of the gauge group acting on something. This is a little annoying, but we get at worst orbifold singularities, which is not the worst. But there can also be reducible connections, where the stabilizer subgroup isn't discrete. For SU_2, at least, this only happens when the connection is trivial, and in particular the associated vector bundle $P \times_{\mathrm{SU}_2} \mathbb{C}^2 \to X$ splits as a direct sum of line bundles. This is a topological fact, and an instance of a general fact about gauge theory: bad singularities can be controlled by the topology of X and P. And on a general 4-manifold, $b_2^+(X) > 0$ is a sufficient condition to preculde the existance of such reducible connections.

Remark 6.7. When you try to define the instanton Floer homology of a 3-manifold, the same problem pops up. To rule out the existence of reducible connections, we assume $H_1(Y;\mathbb{Z}) = 0$, i.e. that Y is an integer homology S^3. The critical locus of the Chern-Simons functional is $\mathcal{M}^\flat(Y) = \mathrm{Hom}(\pi_1(Y), \mathrm{SU}_2)/\mathrm{SU}_2$, as we discussed last time, and this space is highly singular at reducible connections. So if you've seen that restriction before, this is where it's coming from.

Austin-Braam suggest doing equivariant Morse theory for the Chern-Simons functional as a fix, though they do not fully implement the details. The first step would be to replace $\mathcal{A}(Y)/\mathcal{G}(Y)$ with some gauge-fixed space of connections modulo an SU_2-action, though, and maybe this gauge-fixed space of connections is $\mathcal{A}$ modulo the gauge transformations that are the identity at a chosen basepoint. We'd like this to be a Banach (infinite-dimensional) manifold in order for the analysis to work out, but it's not at all clear how to make this work with Sobolev norms, so it's only a Fréchet (infinite-dimensional) manifold, and that's not good enough.

[9] There's more to say here about the bordism, since at this point we've just determined the signature of Q_X, but one can push the argument further and learn the whole intersection form.

Kronheimer-Mrowka make use of the fact that in monopole Floer homology, the reducible connections all have automorphism group U_1 to propose a workaround: a Morse theory of something called the *Chern-Simons-Dirac functional* $CSD\colon \mathcal{P}(Y) \to \mathbb{R}$. Here $\mathcal{P}(Y)$ is a moduli space of pairs (A, ψ), where A is a spin^c connection on Y and ψ is a spinor, modulo a quotient by a gauge group. This can still have reducible connections, but you can blow up $\mathcal{P}(Y)$ at the reducibles to obtain $\widehat{\mathcal{P}}(Y)$, a Banach manifold with boundary diffeomorphic to the space of reducible connections times S^1. One then does Morse theory for (infinite-dimensional) manifolds with boundary.

If you think back to finite-dimensional manifolds with boundary, there are three different notions of homology you can consider: $H_*(M)$, $H_*(\partial M)$, and $H_*(M, \partial M)$, and there is a long exact sequence relating these three. In the same way, and for the same reasons, there are three different notions of monopole Floer homology, HF^+, HF^-, and $\widehat{HF}$, with a long exact sequence relating them.[10] Osváth and Szabó also constructed a version of this, albeit with different names.

There is a sense in which all of this data, and also Seiberg-Witten invariants of 4-manifolds, fit together into a TQFT, but it's a somewhat complicated one, assigning three chain complexes to a 3-manifold, for example.

You might think, great, now let's do this for instanton Floer homology! But the problem is that a lot more can happen with SU_2-connections, and the description of blowups gets a lot trickier. ◄

People sometimes say that if there is any justice in the world, these moduli spaces will be smooth manifolds in the end (at least, assuming some genericity conditions). To that we ask, well, is there justice in the world?

Sometimes, the "regular data," which cuts the moduli space out transversely might not exist. Unlike the previous issues of isotropy and reducibles, which are "stacky," in that they involve nontrivial stabilizer groups, this is "derived," in that it comes about from a failure of transversality.

Example 6.8. Let (M, ω) be a symplectic manifold, and consider the space $\mathcal{M}(A, J)$ of pseudoholomorphic maps $u\colon S^2 \to M$ with homology class $A \in H_2(M)$, modulo $\mathrm{Aut}(S^2)$. The expected dimension contains a factor of $2\langle c_1(u^*TM, J), [S^2]\rangle$, so if $A \in H_2(M)$ has $\langle c_1, [A]\rangle = -1$ and $\mathcal{M}(A, J) \neq \varnothing$, then $[u \circ \psi^k] \in \mathcal{M}(kA, J) \neq \emptyset$, where $\psi^k\colon S^2 \to S^2$ is a degree-k map. For sufficiently large k, the expected dimension of $\mathcal{M}(kA, J)$ is negative, so if it were carved out by a transverse section, then it would be empty. So you can't just say "generically everything works!"

You might want to fix this by choosing $J\colon S^2 \to \mathcal{J}(M, \omega)$ nicely; the additional flexibility in letting the complex structure vary might give us room to find transversality. This actually works,[11] and is what Schwarz does, but then you lose the $\mathrm{Aut}(S^2)$-action, so this cannot work for Gromov-Witten theory.

In general, equivariance and transversality are often mutually exclusive. This is a common problem. You can sometimes average over a compact group, but the Möbius group $\mathrm{Aut}(S^2)$ is noncompact. This is an issue mostly in the symplectic world, and in gauge theory things for some reason work a little better. ◄

[10] Warning: I do not guarantee that this correspondence is in order, e.g. I don't know whether HF^- corresponds to $H_*(\partial M)$, and so on.

[11] Caveat: there can be bubbling to overcome, in which energy concentrates at a point of S^2. This is related to compactification: this can be interpreted as converging in the compactified space of solutions to a nodal curve.