Gauge theoretic invariants

Shing Tak Lam

October 7, 2025

1 Introduction

We want to understand 4-manifolds. From algebraic topology, we have seen that various *invariants* give us a way to distinguish between different 4-manifolds. For example, we can consider the fundamental group π_1, the homology groups H_k, the cohomology groups H^k, and the intersection form $H^2 \otimes H^2 \to \mathbb{Z}$.

One reason why we want to study gauge theory is that it gives us a way to construct new invariants of 4-manifolds. The first result to come out of gauge theory is the following:

Theorem 1.1 (Donaldson)**.** *Let X be a (simply connected) compact oriented smooth 4-manifold, Q_X its intersection form. Then Q_X is diagonalisable.*

Contrasting this with Freedman's result that any unimodular symmetric bilinear form is the intersection form of some closed oriented topological 4-manifold, we can construct topological 4-manifolds with no smooth structure. Similar techniques can be used to construct exotic $\mathbb{R}^4$.

Plan for today:

- Review content on differential geometry of vector and principal bundles,
- Overview on gauge theory and Seiberg–Witten invariants.

Throughout, all manifolds will be assumed to be smooth. Apologies for the bias towards 4-manifolds over 3-manifolds.

2 Bundles

2.1 Fibre, vector and principal bundles

A smooth map $\pi : E \to X$ between smooth manifolds is called a *locally trivial fibration* with fibre F if there exists an open cover $\{U_\alpha\}$ of X, and diffeomorphisms

$$\Phi_\alpha : \pi^{-1}(U_\alpha) \to U_\alpha \times F$$

such that

$$\begin{tikzcd}
\pi^{-1}(U_\alpha) \arrow[rr, "\Phi_\alpha"] \arrow[dr, "\pi"'] & & U_\alpha \times F \arrow[dl, "\mathrm{pr}_1"] \\
& U_\alpha &
\end{tikzcd}$$

commutes. The Φ_α are called *local trivialisations* of the fibre bundle.

The *transition maps* of the fibre bundle are the maps

$$g_{\beta\alpha} : U_\alpha \cap U_\beta \to \mathrm{Diff}(F)$$

given by

$$\Phi_\beta \circ \Phi_\alpha^{-1}(x, v) = (x, g_{\beta\alpha}(x)(v)).$$

The transition maps satisfy the *cocycle conditions*

$$g_{\alpha\alpha} = \mathrm{id},$$
$$g_{\gamma\beta} g_{\beta\alpha} = g_{\gamma\alpha}.$$

Conversely, given an open cover $\{U_\alpha\}$ of X and maps $g_{\beta\alpha} : U_\alpha \cap U_\beta \to \mathrm{Diff}(F)$ satisfying the cocycle conditions, we can construct a locally trivial fibration with fibre F and transition maps $g_{\beta\alpha}$.

A *section* of the fibre bundle $\pi : E \to X$ is a smooth map $s : X \to E$ such that $\pi \circ s = \mathrm{id}_X$. The space of all sections is denoted by $C^\infty(X, E)$.

Let $G \subseteq \mathrm{Diff}(F)$ be a Lie group, which acts on F via a **left** action. We say that the fibration $\pi : E \to X$ has *structure group* G if there exists local trivialisations $\{\Phi_\alpha\}$ such that the transition maps $g_{\beta\alpha}$ take values in G. In this case, we call E a G-*bundle*.

An *automorphism* or a *gauge transformation* of a G-bundle $\pi : E \to X$ is a smooth map $u : E \to E$ such that $\pi \circ u = \pi$, and the maps

$$u_\alpha(x) = \Phi_\alpha|_x \circ u \circ \Phi_\alpha|_x^{-1} : U_\alpha \to \mathrm{Diff}(F)$$

take values in G. The u_α satisfy

$$u_\beta = g_{\beta\alpha} u_\alpha g_{\beta\alpha}^{-1}.$$

The set of all gauge transformations is denoted by $\mathcal{G}(E)$, and this forms a group, called the *gauge group* of E.

Example 2.1. If $F = \mathbb{R}^n$, and

- $G = GL(n, \mathbb{R})$ we obtain *(real) vector bundles*,
- $G = SL(n, \mathbb{R})$ we obtain *oriented vector bundles*,
- $G = O(n)$ we obtain *vector bundle with inner product*,

and so on... If we use $\mathbb{C}^n$ then we obtain *complex* vector bundles.

Finally, we need to introduce the notion of a principal bundle. Let G be a Lie group, a *principal* G-*bundle* $\pi : P \to X$ is a locally trivial fibre bundle with a smooth **right** G-action

$$P \times G \to P \qquad (p, g) \mapsto pg$$

such that $\pi(pg) = \pi(g)$, and we have an atlas of **equivariant** local trivialisations

$$\Phi_\alpha : \pi^{-1}(U_\alpha) \to U_\alpha \times G.$$

In this case, the transition maps are of the form

$$\Phi_\beta \circ \Phi_\alpha^{-1}(x, \gamma) = (x, g_{\beta\alpha}(x)\gamma).$$

2.2 Associated and frame bundles

Let $E \to X$ be a fibre bundle with structure group G. G acts on the model fibre F, but to write down the action of G on E_x, we must choose an identification of F with E_x.

A G-*frame* at $x \in X$ is a diffeomorphism $\psi : F \to E_x$, such that $\Phi_\alpha(x, \cdot) \circ \psi \in G$ for any α with $x \in U_\alpha$. The collection of all such forms a bundle, called the G-*frame bundle*

$$\mathcal{F}(E) = \{(x, \psi) \mid x \in X, \psi : F \to E_x \text{ a } G\text{-frame}\}$$

and this has a right G-action.

Conversely, given a principal bundle $\pi : P \to X$, and a homomorphism $\rho : G \to \mathrm{Diff}(F)$, we can form a locally trivial G-bundle

$$E = P \times_\rho F = \frac{P \times F}{G}$$

where G acts on $P \times F$ by $(p, v)g = (pg, \rho(g^{-1})v)$. This is called the *associated bundle*.

One example of this is that of the adjoint bundle. A Lie group G acts on its Lie algebra $\mathfrak{g}$ via the adjoint representation. For matrix Lie groups, this is exactly conjugation. The associated bundle is denoted by

$$\operatorname{Ad} P = P \times_{\operatorname{Ad}} \mathfrak{g}.$$

Sections of this bundle are maps $\xi : P \to \mathfrak{g}$ satisfying

$$\xi(pg) = g^{-1}\xi(p)g,$$

which forms a Lie algebra, which we denote by

$$\Omega^0(X, \mathfrak{g}_P) = \Omega^0_{\mathrm{ad}}(P, \mathfrak{g}).$$

The notation on the right means that we are considering G-equivariant maps $P \to \mathfrak{g}$.

2.3 Bundle valued forms

Using cocycles, given vector bundles E, F, we can form bundles $E^*, E \otimes F, E \oplus F, S^n E, \Lambda^n E$. In particular, given smooth manifold X, we have the tangent bundle TX. The dual is the *cotangent bundle* T^*X.

The space of k-forms on X is

$$\Omega^k(X) = C^\infty(X, \Lambda^k T^*X).$$

Given any vector bundle $E \to X$, the space of k-forms on X with values in E is

$$\Omega^k(X, E) = C^\infty(X, \Lambda^k T^*X \otimes E).$$

3 Connections

3.1 Connections on vector bundles

Let $E \to X$ be a vector bundle. A *covariant derivative* on E is an $\mathbb{R}$-linear map $\nabla : C^\infty(X, E) \to \Omega^1(X, E)$, such that

$$\nabla(fs) = f\nabla s + df \otimes s$$

for $f : X \to \mathbb{R}$ and $s \in C^\infty(X, E)$. Within a trivialisation Φ_α,

$$(\nabla s)_\alpha = ds_\alpha + A_\alpha s_\alpha$$

for a matrix valued 1-form $A_\alpha \in \Omega^1(U_\alpha, \mathfrak{gl}(n, \mathbb{R}))$. If E has structure group $G \subseteq GL(n, \mathbb{R})$, and $\{\Phi_\alpha\}$ is a collection of compatible trivialisations, a covariant derivative ∇ is called a G-*covariant derivative* if each $A_\alpha \in \Omega^1(U_\alpha, \mathfrak{g})$.

The space of covariant derivatives on E is naturally an affine space modelled on $\Omega^1(X, \operatorname{End}(E))$. On this space, we have a $\mathcal{G}(E)$-action as follows:

$$u^*\nabla = u^{-1} \circ \nabla \circ u.$$

3.2 Ehresmann connections

Let $\pi : E \to X$ be a fibre bundle. Then for each $p \in E$, we have a map

$$d\pi_p : T_p E \to T_x X.$$

We will write

$$V_p = \ker(d\pi_p)$$

for the *vertical subspace*. A subspace H_p of $T_p E$ is *horizontal* if

$$T_p E = V_p \oplus H_p.$$

An *Ehresmann connection* on E is $\{H_p\}_{p \in E}$ horizontal subspaces, which depend smoothly on p. When E has additional structure, we can ask for the Ehresmann connection to be compatible with this additional structure.

If E is a vector bundle, let $s_\lambda : E \to E$ denote the map defined by scalar multiplication by $\lambda \in \mathbb{R}$. Then we would ask for

$$H_{\lambda p} = ds_\lambda(H_p).$$

Note that for a vector bundle, a covariant derivative determines an Ehresmann connection, and an Ehresmann connection determins a covariant derivative. Thus, I'll be sloppy and call everything a connection.

3.3 Parallel transport

Given an Ehresmann connection, we have a natural notion of parallel transport. Since $d\pi_p : T_pE \to T_xX$ is surjective, we have an isomorphism

$$d\pi_p : H_p \to T_xX.$$

Thus, given any tangent vector $v \in T_xX$, there exists a unique *horizontal lift* $\tilde{v} \in H_p$.

Given a curve $\gamma : \mathbb{R} \to X$, there exists a *parallel transport* map

$$P_\gamma : E_{\gamma(0)} \to E_{\gamma(1)}$$

which allows us to identify fibres of the fibre bundle. Note that if E has structure group G, and the connection is compatible with the structure group, then P_γ is "compatible" with the G-structure. For example, if E is an $O(n)$-bundle, then P_γ is an isometry.

Suppose E is a vector bundle. The *holonomy group* of the connection is the group

$$\mathrm{Hol}(E, \nabla, x) = \{P_\gamma : \text{closed loops } \gamma \text{ based at } x\}.$$

3.4 Connections on principal bundles

Let $P \to X$ be a principal G-bundle. For $g \in G$, let $r_g : P \to P$ denote the right multiplication map. Then an Ehresmann connection on P is a smooth family of horizontal subspaces $\{H_p\}$, such that

$$H_{pg} = dr_g(H_p).$$

Any such connection is given by the kernel of a one-form $A \in \Omega^1(P, \mathfrak{g})$, which satisfies

$$\begin{aligned} A_{pg}(vg) &= g^{-1}A_p(v)g \\ A_p(p\xi) &= \xi \end{aligned}$$

for $v \in T_pP$, $g \in G$, $\xi \in \mathfrak{g}$. We call A a *connection* 1*-form*, and the set of all such is denoted by $\mathcal{A}(P)$. Note that $\mathcal{A}(P)$ is an affine space, modelled on $\Omega^1(X, \mathfrak{g}_P)$. The group $\mathcal{G}(P)$ of gauge transformations acts on $\mathcal{A}(P)$ by

$$u^*A = u^{-1}du + u^{-1}Au.$$

4 Curvature

4.1 Curvature on vector bundles

Given a covariant derivative ∇ on a vector bundle E, we get associated operators

$$\nabla : \Omega^k(X, E) \to \Omega^{k+1}(X, E).$$

A natural question is then: is $\nabla^2 = 0$? In general, the answer is **no**.

We define for v, w vector fields on X, s a section of E,

$$F_\nabla(v, w)s = \nabla_v\nabla_w s - \nabla_w\nabla_v s + \nabla_{[v,w]}s.$$

Note that $F_\nabla(v, w) = F_\nabla(w, v)$, and that F_∇ is $C^\infty(X)$-linear. Thus, we get an element $F_\nabla \in \Omega^2(X, \mathrm{End}(E))$.

The *Bianchi identity* states that

$$\nabla F_\nabla = 0.$$

Locally,

$$F_\nabla = dA_\alpha + A_\alpha \wedge A_\alpha.$$

The connection ∇ is *flat* if $F_\nabla = 0$.

Note that if E has structure group G, and ∇ is a G-connection, then $F_\nabla \in \Omega^2(X, \mathrm{End}_{\mathfrak{g}}(E))$, where

$$\mathrm{End}_{\mathfrak{g}}(E) = \{A : E_x \to E_x \mid \psi^{-1} \circ A \circ \psi \in \mathfrak{g}\}$$

for every G-frame $\psi : \mathbb{R}^n \to E_x$. For example, if we choose an inner product on E, and suppose ∇ is an $O(n)$-connection, then F_∇ is skew-symmetric.

4.2 Curvature on principal bundles

Recall that a connection on a principal G-bundle $P \to X$ is given as the kernel of a connection 1-form $A \in \Omega^1(P, \mathfrak{g})$. The *curvature* of A is the 2-form $F_A \in \Omega^2_{\mathrm{ad}}(P, \mathfrak{g})$ defined by

$$F_A(u, v) = dA_p(u, v)$$

if $u, v \in H_p$, and zero otherwise.

Suppose we have a representation $\rho : G \to GL(V)$, and consider the associated bundle $E = P \times_\rho V$. The connection A on P induces a connection ∇ on E, and

$$F_\nabla(s) = d\rho(F_A) \cdot s.$$

4.3 Chern classes and Chern–Weil theory

Let $E \to X$ be a complex vector bundle over a compact manifold. Suppose E has rank k. Define the total Chern class of E as

$$c(E) = c_0(E) + \cdots + c_k(E),$$

where $c_i(E) \in H^{2i}(E)$ is the i-*th Chern class*. Axiomatically, we have a functor

$$c : (E \to X) \mapsto c(E) \in H^*(X),$$

which satisfies:

(i) naturality - isomorphic vector bundles have the same Chern class,

(ii) functoriality - $c(f^*E) = f^*c(E)$,

(iii) direct sum - $c(E \oplus F) = c(E)c(F)$

(iv) trivial - $c(\text{trivial}) = 1$,

(v) $c_1(\mathcal{O}_{\mathbb{CP}^1}(1)) = [H]$.

Note that c_k is the Euler class. In fact, the Chern class (as a functor) is determined by the Chern classes of line bundles.

Chern-Weil theory says that the Chern class (modulo torsion) of a complex vector bundle can be computed using the curvature of any connection. For example,

$$c_1(E) = \frac{i}{2\pi}\operatorname{tr}(F_\nabla) \qquad c_2(E) = \frac{1}{8\pi^2}(\operatorname{tr}(F_h^2) - \operatorname{tr}(F_h)^2)$$

and so on. In particular, the Chern classes, which are topological quantities, can be computed using analytic methods.

5 Riemannian geometry

5.1 Levi-Civita connection

Let (X, g) be a Riemannian manifold of dimension n. The metric g gives TM an $O(n)$-structure. However, $O(n)$-connections are not unique. The *Levi-Civita connection* of (X, g) is the unique $O(n)$-connection such that

$$[v, w] = \nabla_v w - \nabla_w v$$

for all vector fields v, w on X.

5.2 Riemann curvature tensor

The *Riemann curvature tensor* is the curvature of the Levi-Civita connection ∇, which we will (now) denote by R. This is defined by

$$R(u,v)w = \nabla_u \nabla_v w - \nabla_v \nabla_u w - \nabla_{[u,v]} w.$$

There are various symmetry properties which this satisfies, but I won't cover them today.

Let $e_1, \dots, e_n$ be an orthonormal frame (i.e. local sections) for TX. The *Ricci curvature* is

$$\mathrm{Ric}(u,v) = \sum_{i=1}^{n} \langle R(e_i, u)v, e_i \rangle,$$

which is a symmetric bilinear form. The *scalar curvature* is

$$\mathrm{Scal} = \sum_{i=1}^{n} \mathrm{Ric}(e_i, e_i).$$

Warning: Different people have different conventions on which indices to sum over, and (potentially) factors of 2 appearing.

When $n = 2$, R and Ric are determined by Scal, which is double the Gaussian curvature.

When $n = 3$, R is determined by Ric.

5.3 Hodge theory

Let V be a vector space of dimension n, and let $\Lambda^k V$ denote the k-th exterior power of V. Recall that

$$\dim(\Lambda^k V) = \binom{n}{k}.$$

In particular, $\dim(\Lambda^n V) = 1$. An inner product on V induces an inner product on $\Lambda^k V$ for all V. Thus, there are exactly two elements of $\Lambda^n V$ of norm 1, which are the *orientations* of V.

Choosing an orientation gives us a fixed isomorphism $\Lambda^n V \cong \mathbb{R}$. We now get a pairing

$$\begin{aligned} \Lambda^k V \times \Lambda^{n-k} V &\to \Lambda^n V \cong \mathbb{R} \\ (\alpha, \beta) &\mapsto \alpha \wedge \beta. \end{aligned}$$

This pairing is non-degenerate, and so we get a map $* : \Lambda^k V \to \Lambda^{n-k} V$, which satisfies

$$\langle \alpha, \beta \rangle = \alpha \wedge *\beta.$$

Suppose X is compact and oriented. Applying the above to $V = T_x^* X$, we get Hodge star operators

$$* : \Omega^k(X) \to \Omega^{n-k}(X).$$

We can use this to define the *codifferential*

$$d^* = (-1)^{n(k+1)+1} * d * : \Omega^k(X) \to \Omega^{k-1}(X),$$

which is the adjoint with respect to the L^2-inner product to the exterior derivative $d : \Omega^{k-1}(X) \to \Omega^k(X)$.

The *Laplace–Beltrami operator* is

$$\Delta = dd^* + d^*d : \Omega^k(X) \to \Omega^k(X).$$

Warning: On $\mathbb{R}^n$ with the usual inner product,

$$\Delta = -\sum_i \frac{\partial^2}{\partial x_i^2}.$$

We say that a differential form α is *harmonic* if $\Delta\alpha = 0$. Equivalently, $d\alpha = 0$ and $d^*\alpha = 0$. Let $\mathcal{H}^k(X)$ denote the space of harmonic k-forms on X.

Theorem 5.1 (Hodge)**.**

$$\mathcal{H}^k(X) \cong H^k(X; \mathbb{R}).$$

The left hand side is defined using elliptic PDE theory, whereas the right hand side is a purely topological quantity.

6 Gauge theory

6.1 Maxwell's equations

Recall that the (vacuum) Maxwell equations are that

$$\begin{aligned}\operatorname{div}(E) &= 0\\ \operatorname{div}(B) &= 0\\ \frac{\partial B}{\partial t} + \operatorname{curl}(E) &= 0\\ \frac{\partial E}{\partial t} - \operatorname{curl}(B) &= 0,\end{aligned}$$

where $E = (E_x, E_y, E_z)$ is the electric field and $B = (B_x, B_y, B_z)$ is the magnetic field. We can write

$$F = E_x dx \wedge dt + E_y dy \wedge dt + E_z dz \wedge dt + B_x dy \wedge dz + B_y dz \wedge dx + B_z dx \wedge dy$$

for the electromagnetic tensor, which we can think of as a 2-form on $\mathbb{R}^4$. Then Maxwell's equations become

$$dF = 0 \qquad d^*F = 0.$$

Since $H^2(\mathbb{R}^4) = 0$, we can find some A such that $F = dA$, which is called the electromagnetic potential. Then since $d^2 = 0$, all we need is that $d^*dA = 0$.

Note that A is not unique, it is only unique up to addition of $a \in \Omega^1(\mathbb{R}^4)$ such that $da = 0$. Physically, this represents a "redundancy" of the system. One way of ensuring A is unique is to ask for $d^*A = 0$, which is a "gauge fixing" condition. Then Maxwell's equations becomes

$$(dd^* + d^*d)A = 0.$$

But the left hand side is just the Laplacian! So if we consider Maxwell's equations on a compact oriented 4-manifold X, then Hodge's theorem tells us that the space of solutions is exactly $H^1(X; \mathbb{R})$.

We can think of A as defining a connection on the trivial bundle $X \times \mathbb{R}^2$, which we can think of as an $SO(2)$-bundle. The non-uniqueness corresponds exactly to the $SO(2)$-symmetry of the equation, or redundancies in the system. The electromagnetic field F is then given by the curvature of this connection.

6.2 Yang–Mills theory

The group $U(1) = SO(2)$ is abelian, which considerably simplifies things. For physics reasons, we would like to consider more general groups G, e.g. $SU(2)$ for the weak force, $SU(3)$ for the strong force, and so on. This leads us to Yang–Mills theory.

Let $P \to X$ be a principal G-bundle. A connection A on P is called *Yang–Mills* if

$$\nabla_A^* F_A = 0.$$

Note that the Bianchi identity implies that $\nabla_A F_A = 0$. When $G = SO(2)$ and P is trivial, we recover Maxwell's equations.

Here, ∇^* is the adjoint of the covariant derivative, defined with respect to the L^2-inner products. An equivalent statement is that

$$\nabla_A * F_A = 0.$$

We can now start doing some numerology. Suppose X is a 4-manifold. Then the Hodge star defines a map

$$* : \Omega^2(X) \to \Omega^2(X),$$

with $*^2 = 1$. Thus, we get an eigendecomposition

$$\Omega^2(X) = \Omega^2_+(X) \oplus \Omega^2_-(X),$$

where $\Omega^2_+(X)$ consists of the *self-dual* forms and $\Omega^2_-(X)$ the *anti-self-dual* forms.

In particular, since curvature is a 2-form, we say that A satisfies the *self-dual Yang–Mills equation* if

$$*F_A = F_A$$

and A satisfies the *anti-self-dual Yang–Mills equation* if

$$*F_A = -F_A.$$

6.3 Moduli space of solutions and 4-manifold invariants

Before, we saw that the space of solutions to Maxwell's equations is a vector space, whose dimension is b^1. A natural question is: what about the space of solutions to the Yang–Mills equations?

As it turns out, the space of solutions can have very complicated geometry. Donaldson studied the moduli space of anti-self-dual connections, and related the geometry of the moduli space to topological properties of the underlying 4-manifold. This is the first *gauge theoretic invariant*.

The general recipe of gauge theory is similar. We write down a PDE, construct a moduli space of solutions, and then try to extract topological information from the geometry of the moduli space. One thing we can do is integrate a suitable cohomology class over the moduli space to get a number, and this is closely related to *enumerative geometry*, which is a reading group which Noah is organising.

Remark 6.1. The main topic in Noah's reading group seems to be Donaldson–Thomas theory, which is an enumerative invariant defined on Calabi-Yau 3-folds. One can think of this as a holomorphic version of Chern–Simons theory on 3-manifolds. Another topic which will show up in Noah's reading group is Gromov–Witten theory, which Taubes related to Seiberg–Witten theory on symplectic 4-manifolds.

There is a very interesting programme proposed by Donaldson, Thomas and Segal for defining enumerative invariants in dimensions 6, 7 and 8 using gauge theory, but that is well beyond the scope of this reading group.

7 Seiberg–Witten theory

Donaldson's invariants can be used to prove many interesting results, but they are very difficult to compute. Witten introduced the Seiberg–Witten equations, which are simpler to work with, and can be used to recover many of Donaldson's results.

7.1 Spin geometry

Recall that we have a double cover $SU(2) \to SO(3)$ of Lie groups. In general, we can define $\mathrm{Spin}(n)$ to be the universal cover of $SO(n)$. Roughly speaking, spinors correspond to things which get sent to its negative after a rotation by 360 degrees, and looks something like a Möbius band.

In the second talk, we'll see what it means to have a spin, or spin^c structure on a manifold or a vector bundle, and what obstructios exist.

When a vector bundle has more structure, it makes sense to ask for a connection to preserve that structure. For example, we can talk about *orthogonal connections* on a vector bundle, which are ones which preserve an inner product. In the third talk, we will understand what it means for a connection to preserve a spin (or spin^c) structure. Using this, we can define what is known as a *Dirac operator* $\mathcal{D}$.

Morally, a Dirac operator is a square root of the Laplacian. More precisely, we have the *Weitzenböck formula*

$$\mathcal{D}^2 = \nabla^*\nabla + \text{lower order terms.}$$

7.2 Seiberg–Witten equations

With all of this set-up, we can now define the Seiberg-Witten equations for a 4-manifold, which are equations for a connection A and a *monopole* Φ:

$$\mathcal{D}_A\Phi = 0 \qquad \text{(Dirac equation)}$$
$$F_A^+ + \eta = \sigma^+((\Phi\Phi^*)_0). \qquad \text{(Monopole equation)}$$

The precise form of the equations doesn't matter for now. Our goal is to study the geometry of the moduli space

$$\mathcal{M} = \mathcal{A}/\mathcal{G},$$

where $\mathcal{A}$ is the space of solutions to the Seiberg-Witten equations, and $\mathcal{G}$ is the gauge group. The main result is that when $b^+(X) > 0$, $\mathcal{M}$ is a finite-dimensional compact oriented manifold, for a generic choice of perturbation η.

- finite-dimensional - This follows from $\mathcal{D}$ being a *Fredholm* operator, and the dimension of the moduli space can be calculated using the *Atiyah–Singer index theorem*.

- compactness - A priori there should be no reason why the moduli spaces are compact. For example, in the case of Yang-Mills theory, one needs to add in *ideal connections* (or for those who are algebro-geometrically minded, coherent sheaves) to make the moduli space compact.

  There is a general compactness result proven by Karen Uhlenbeck, and the main result when proving compactness is proving an a priori estimate, so that we can show that any Cauchy sequence has a convergent subsequence.

- manifold - It's not obvious why the space $\mathcal{A}$ is even a manifold. We can think of the Seiberg-Witten equations as being given by a map
$$S : V \to W,$$
  where V, W are infinite-dimensional Banach spaces, and $\mathcal{A} = S^{-1}(0)$. In general, 0 may not be a regular value of S, and so $\mathcal{A}$ won't be well behaved.

  In this case, we can use an infinite-dimensional version of Sard's theorem (set of critical values has measure zero/is meagre in the sense of Baire), to show that for a generic choice of η, $\mathcal{M}$ is a manifold.

At this point, one may rightly object: how do we know that the moduli space does not depend on the choice of perturbation? In some cases, we can show that for different choices, we have a cobordism between the moduli spaces. However, in other cases, we can't do this, and what we get is *wall-crossing* phenomenon.

At the end of the day, what we want to extract are numbers.

- If $\mathcal{M}$ is 0-dimensional, we can count the number of points in $\mathcal{M}$, with signs.
- In general, by integrating a suitable cohomology class over $\mathcal{M}$, we can define the *Seiberg–Witten invariant*.

7.3 Properties

How do we construct 4-manifolds? One reasonable approach is to start off with a class of "nice" 4-manifolds, and then do things to them, such as connect sum, cut-and-paste, blow-up and so on. Thus, to understand how the Seiberg–Witten invariant behaves, its important to understand

(i) what the Seiberg–Witten invariants of our "nice" manifolds are,

(ii) how the Seiberg–Witten invariants behave under various operations.

For (i), a reasonable choice of "nice" manifolds are Kähler surfaces. In particular, many of these come from *algebraic geometry*, and we can use techniques coming from algebraic geometry to understand them as well. Whenever we have a property about Kähler manifolds, a natural question is then whether they hold for symplectic manifolds as well.

In fact, the Seiberg–Witten invariants can be used to construct obstructions to manifolds admitting a symplectic structure.

For (ii), there are formulae for the Seiberg–Witten invariants of connect sums and blow-ups, which means that in some concrete examples, we can compute how the Seiberg-Witten invariants change.

7.4 Embedded surfaces

Taubes proved that the Seiberg–Witten invariants agreed with the Gromov–Witten invariants, which are given by counting pseudoholomorphic curves (i.e. Riemann surfaces) within a symplectic manifold. Thus, it is natural to conjecture that the Seiberg–Witten invariants should be able to give us information about embedded surfaces in our 4-manifold. In fact, we have a lower bound on the genus of such a surface.

8 Things which I did not talk about

8.1 Physics

Many of the equations which we have seen, and the ones which are studied, often come from physics. For example, Yang–Mills theory comes from quantum field theory, and the Seiberg–Witten equations come from supersymmetric field theory. More generally, there are many connections between the (mathematical) physics and the applications to topology.

8.2 Analysis

The equations are often non-linear (elliptic) partial differential equations. To make sense of various constructions, we would need *completeness*. In particular, we would need to work with Sobolev spaces, and use various tools from functional analysis and partial differential equations. This is a very technical topic, and I have not talked about it today.

8.3 Floer theory

Floer theory is an infinite-dimensional analogue of Morse theory. Much as we can recover the cohomology *ring* using Morse theory, Floer theory allows us to build stronger invariants, which are not just numbers, but groups, rings, A_∞-categories and so on. This is also closely related to the notion of *categorification*.

8.4 Homotopical aspects

Given a continuous map $f : X \to Y$, we get an induced map $f^* : H^*(Y) \to H^*(X)$, which only depends on the homotopy class of f. So far, everything we have mentioned is closer to working with f^* than with f. A lot of recent work has been dedicated to try and understand *homotopical* aspects of gauge theory. For example, the *Seiberg–Witten Floer stable homotopy type* was used by Manolescu to show that non-triangulable manifolds exist in dimensions ≥ 5.