Lecture II: Basics

Jonathan Evans

30th September 2010

In this lecture we will cover some very basic material in symplectic geometry.

- First, we will discuss classical Hamiltonian mechanics, reinterpret it in our symplectic setting and show that it was all worthwhile by defining cotangent bundles.
- Second we will talk about the linear algebra of alternating forms, the special subspaces we get in a symplectic vector spaces and the topology of various homogeneous spaces (Lagrangian Grassmannian, Siegel upper half space) which arise in this context.

Hamiltonian dynamics

Another motivation for studying symplectic geometry is that it provides a good setting for talking about classical dynamics.

Let $\{q_j\}_{j=1}^n$ be coordinates in $\mathbb{R}^n$ and let v_j be the component of velocity in the q_j-direction.

A classical dynamical system can be described by a Lagrangian function

$$\mathcal{L}(q, v, t)$$

as follows. The Lagrangian defines an *action functional*

$$L : \Omega(x, y) \to \mathbb{R} \qquad L(\gamma) = \int_0^1 \mathcal{L}(\gamma, \dot{\gamma}, t)dt$$

on the space of paths $\Omega(x, y)$ from x to y. The classical motions of the system are the critical points of this functional, i.e. solutions to the Euler-Lagrange (E-L) equation

$$\frac{d}{dt}\left(\frac{\partial \mathcal{L}}{\partial v_j}\right) = \frac{\partial \mathcal{L}}{\partial q_j}$$

This is a system of second-order PDE for the q_j. If we write

$$p_j = \frac{\partial \mathcal{L}}{\partial v_j}$$

the E-L equation becomes

$$\dot{p}_j = \frac{\partial \mathcal{L}}{\partial q_j}$$

We want to eliminate v from our equations and write everything in terms of q and p. We can do this (via the implicit function theorem) if the Jacobian matrix

$$\frac{\partial p_i}{\partial v_j} = \frac{\partial^2 \mathcal{L}}{\partial v_i \partial v_j}$$

has non-zero determinant.

Assume that $\mathcal{L}$ satisfies this condition, then we can write v_j in terms of the q and p. We now consider (instead of $\mathcal{L}$) the Hamiltonian function

$$H = \sum_j v_j p_j - \mathcal{L}$$

Differentiating H with respect to v_j gives

$$\frac{\partial H}{\partial v_j} = p_j - \frac{\partial \mathcal{L}}{\partial v_j} = 0$$

by definition of p so we have eliminated v-dependence[1].

[1] Of course this isn't really what's going on: you might like to think about how to phrase this in terms of the vertical cotangent bundle of the tangent bundle (i.e. the bundle over TX whose fibre at (x, v) is the space of linear functionals on $T_x X$) and its canonical function evaluating vertical covectors on tangent vectors ("$p \cdot v$").

Differentiating with respect to q_j gives

$$\frac{\partial H}{\partial q_j} = -\frac{\partial \mathcal{L}}{\partial q_j} = \dot{p}_j$$

by the E-L equation. Thus we have a pair of first order PDE called *Hamilton's equations of motion*:

$$\dot{q}_j = \frac{\partial H}{\partial p_j}$$
$$\dot{p}_j = -\frac{\partial H}{\partial q_j}$$

These describe the time evolution of a system with coordinates q and "conjugate momenta" p.

In the preceeding discussion we took a function $H : \mathbb{R}^{2n} \to \mathbb{R}$ and produced a vector field $(\dot{q}, \dot{p})$ on $\mathbb{R}^{2n}$. Certainly the vector field depends only on the first derivatives of H, so somehow it is a rephrasing of the information contained in the 1-form dH.

At a point (p, q) a 1-form is an element of the dual space $T^*_{(p,q)}\mathbb{R}^{2n}$ so we're looking for a map $T^*\mathbb{R}^{2n} \to T\mathbb{R}^{2n}$, something like the musical $\sharp$ isomorphism in Riemannian geometry. Equivalently, we're looking for a nondegenerate bilinear form on $T\mathbb{R}^{2n}$. It is clear from Hamilton's equations that the bilinear form in question is the alternating 2-form

$$\omega = \sum_{i=1}^{n} dq_i \wedge dp_i$$

By this I mean that

Exercise

$\diamondsuit$: *Hamilton's equations can be rewritten as*

$$\iota_{(\dot{q},\dot{p})}\omega = dH \tag{1}$$

Definition

A vector field v is called Hamiltonian if

$$\iota_v \omega = dH$$

for some function H. The flow of a Hamiltonian vector field is called a Hamiltonian flow.

Lemma

The 2-form ω is preserved by any Hamiltonian flow.

Proof.

By Cartan's formula for the Lie derivative of a form

$$\begin{aligned} \mathfrak{L}_v\omega &= \iota_v d\omega + d\iota_v\omega \\ &= ddH \\ &= 0. \end{aligned}$$

□

This proof relies on Hamilton's equation $\iota_v\omega = dH$ (in fact all we need is that $\iota_v\omega$ is closed, which is the same as exact in $\mathbb{R}^{2n}$) and on the fact that $d\omega = 0$ which is clear for this particular ω.

We can therefore generalise Hamiltonian dynamics to manifolds other than this linear phase space providing those manifolds admit a *symplectic form*, that is a nondegenerate, closed 2-form. This 2-form provides a dictionary to translate between closed 1-forms and vector fields whose flow preserves ω.

Definition

A symplectic manifold is a manifold X equipped with a symplectic form ω.

In this more general context one easily sees classical properties of Hamiltonian dynamics still hold:

Exercise

$\diamondsuit$: *Show that on a symplectic manifold the Hamiltonian flow of a function H preserves the Hamiltonian function H (of course if we formulate the theory with a time-dependent Hamiltonian this obviously won't be true any more. However, the symplectic structure is still preserved.)*

The cotangent bundle

The cotangent bundle is the simplest extension of dynamics on $\mathbb{R}^n$. The phase space $\mathbb{R}^{2n}$ is just the space of coordinates and conjugate momenta for dynamics on $\mathbb{R}^n = \{p \equiv 0\}$. What happens if we're interested in dynamics on a sphere (e.g. the surface of the Earth)? Or on the complement of the trefoil knot (because we like pretty pictures)? Clearly the phase space should locally look like $(\mathbb{R}^{2n}, \sum_{i=1}^n dq_i \wedge dp_i)$ because a manifold looks locally like $\mathbb{R}^n$. But how are we supposed to patch together these charts into a global phase space?

In the Lagrangian picture, changing q coordinates by a local diffeomorphism $q \mapsto q'$ gives rise to a velocities $v \mapsto v'$ via $v_i' = \frac{d\phi_i}{dt} = \frac{\partial q_i'}{\partial q_j} v_j$ by the chain rule. The quantity $p_i = \frac{\partial \mathfrak{L}}{\partial v_i}$ changes to

$$p_i' = \frac{\partial \mathfrak{L}}{\partial v_i'} = \frac{\partial \mathfrak{L}}{\partial v_j}\frac{\partial v_j}{\partial v_i'} = p_j \frac{\partial q_j}{\partial q_i'}$$

Tensorially, then, v transforms as a vector and p as a covector. Therefore, unsurprisingly, the correct global phase space for Hamiltonian dynamics on a manifold M is the cotangent bundle T^*M.

If we pick local coordinates q on a patch in M and let p be the vertical coordinates in T^*M defined as follows: if η is a 1-form then

$$\eta = \sum_i p_i(\eta) dq_i$$

The 2-form ω will be defined so that it agrees with the form in $\mathbb{R}^{2n}$ over any patch, i.e.

$$\omega = \sum_i dq_i \wedge dp_i$$

Remarkably, this does not depend on the choice of q coordinates[2]:

[2] Observe that we don't have to pick p coordinates: they come for free once we have picked q coordinates. This naturality is what underlies the fact that ω is canonical.

Lemma

If q' is another choice of q-coordinates and ω' the corresponding symplectic form then $\omega = \omega'$.

Proof.

When we change coordinates,

$$dq'_i = \frac{\partial q'_i}{\partial q_j} dq_j = \sum_i M_{ij} dq_j$$

The conjugate momenta p'_j are still linear coordinates in a fibre (by construction) and since

$$\eta = \sum_i p_i(\eta) dq_i = \sum_j p'_j(\eta) dq'_j = \sum_j p'_j(\eta) M_{ji} dq_i$$

we see that $p' = M^{-1}p$. In the combination $\sum_i dq'_i \wedge dp'_i$, the M and the M^{-1} cancel and leave $\sum_i dq_i \wedge dp_i$. $\square$

Theorem

The cotangent bundle T^*M of a manifold M is equipped with a natural symplectic 2-form ω. Naturality is meant in the following precise sense: If $\phi : M \to M$ is a diffeomorphism then define $\Phi = (\phi^{-1})^* : T^*M \to T^*M$ (which is a bundle map living over ϕ, i.e. it is linear on fibres). Then

$$\Phi^*\omega = \omega.$$

This follows from what we have said, but let's give a fresh proof to shed new light on what's really going on. There is a tautological 1-form λ on T^*M which at a point (q, p) is just the form p. By this I mean: to evaluate λ on a vector $V \in T_{(q,p)}T^*M$, push V forward along the cotangent bundle projection to a vector v at q and then set $\lambda(V) = p(v)$. In coordinates this 1-form is

$$\lambda = \sum_i p_i dq_i$$

and $d\lambda = \omega$ is the 2-form we have defined above. Not only is ω closed, but it is also *exact* in this case.

Naturality of ω follows from naturality of the tautological form

$$\Phi^*\lambda = \lambda$$

To see this latter, let $x = (q, p) \in T^*M$, $V \in T_x T^*M$, $\pi : T^*M \to M$, $v = \pi_* V$ and $\Phi = (\phi^{-1})^* : T^*M \to T^*M$.

$$\begin{aligned}
(\Phi^*\lambda)_x(V) &= \lambda_{\Phi(x)}(\Phi_* V) \\
&= (\Phi(x))(\phi_* v) \\
&= ((\phi^{-1})^* p)(\phi_* v) \\
&= p(v) \\
&= \lambda(V).
\end{aligned}$$

Geodesic flow as a Hamiltonian system

Let (M, g) be a Riemannian manifold. There is a very natural Hamiltonian function we can write down on T^*M

$$H(q,p) = \frac{1}{2}|p|^2$$

We need g to make sense of the norm $|\cdot|^2$. This is the Legendre transform of the Laplacian $\mathfrak{L} = \frac{1}{2}|v|^2$ on TM, whose Euler-Lagrange equations give rise to the geodesic flow. It should come as no surprise that:

Exercise

$\heartsuit$: *The Hamiltonian flow of H gives rise to the cogeodesic flow which is $\flat$-dual to the geodesic flow.*

Since the Hamiltonian is a conserved quantity of its own flow, we see that for every $\ell \geq 0$ the geodesic flow preserves the sphere bundle of radius ℓ in TM (of course, we already knew this).

Exercise

$\diamondsuit$: *Let (M, g) be the round 2-sphere. Describe the radius-ℓ circle bundle and the flowlines of the geodesic flow. Do the same for the flat 2-torus.*

Definition

A symplectic vector space is a pair (V, ω) where V is a real vector space and ω is a nondegenerate alternating bilinear form.

For example, $V = \mathbb{R}^{2n}$, $\omega = dx_1 \wedge dy_1 + \cdots + dx_n \wedge dy_n$. This example is called the standard symplectic $\mathbb{R}^{2n}$.
The form ω provides an isomorphism

$$\begin{aligned} \iota.\omega : V &\to \check{V} \\ X &\mapsto \omega(X, -) \end{aligned}$$

An important concept is the symplectic orthogonal complement of a subspace $W \leq V$:

$$W^{\omega} = \{v \in V | \omega(w, v) = 0 \forall w \in W\}$$

equivalently $W^{\omega} = (\iota.\omega)^{-1}\mathrm{ann}(W)$ where ann is the annihilator. Since W and its annihilator have complementary dimension we see that

$$\dim W + \dim W^{\omega} = \dim V$$

The symplectic orthogonal complement allows us to distinguish important subspaces of V:

- Isotropic subspaces, for which $W \subset W^\omega$, i.e. $\omega|_W = 0$,
- Symplectic subspaces, for which $W \cap W^\omega = \{0\}$, i.e. $\omega|_W$ is symplectic,
- Coisotropic subspaces, for which $W^\omega \subset W$.

Exercise

$\diamondsuit$: *Give an example of a subspace of each type in the standard symplectic* $\mathbb{R}^{2n}$. *Give an example of a subspace which falls into none of these types.*

Observe:

Lemma

If (V, ω) is a symplectic vector space then $\dim V = 2n$ *is even and there is a linear isomorphism* $V \to \mathbb{R}^{2n}$ *intertwining* ω *and* ω_0*, the standard symplectic form*

$$\omega_0 = dx_1 \wedge dy_1 + \cdots + dx_n \wedge dy_n$$

on $\mathbb{R}^{2n}$*.*

Proof.

Assume true for all symplectic vector spaces of dimension less than V (it is certainly true for $\dim V = 0$, so assume $\dim V > 0$). Pick a vector $v_1 \in V$ and let v_2 be such that $\omega(v_1, v_2) = 1$ (nondegeneracy). The span $\langle v_1, v_2 \rangle$ is a symplectic subspace and its symplectic orthogonal complement is a symplectic subspace W of dimension $\dim V - 2$. By induction we can find a basis $v_3, \ldots, v_{2n}$ of W such that the symplectic form gives the only nonvanishing products as $\omega(v_{2i-1}, v_{2i}) = 1 = -\omega(v_{2i}, v_{2i-1})$. The union of this basis with $\{v_1, v_2\}$ does the trick. $\square$

So there is precisely one symplectic vector space up to isomorphism in every even dimension. Symplectic manifolds are therefore even dimensional. They are also naturally oriented, because the n-th power of the symplectic form ω_0 is the nondegenerate volume form

$$n! dx_1 \wedge dy_1 \wedge \cdots \wedge dx_n \wedge dy_n$$

That is

Lemma

A symplectic vector space (V, ω) has a natural nondegenerate volume form ω^n.

- A submanifold of a symplectic manifold is called isotropic, symplectic or coisotropic respectively if its tangent spaces are of the corresponding type.
- Notice that a random submanifold will not be of any specific type: it may have tangencies of varying type. Therefore these classes of submanifold are special.
- One particularly special class is the class of Lagrangian submanifolds:

Definition

An isotropic subspace has dimension at most $\dim V/2$ *and a Lagrangian subspace is an isotropic subspace of this dimension.*

Examples of global questions in symplectic geometry ask about the existence of Lagrangian or symplectic submanifolds and the topology of the (infinite-dimensional) space of all Lagrangian or symplectic submanifolds.

- Today we're only interested in linear theory, so we might as well as the linear question: what is the space of Lagrangian subspaces of a symplectic vector space (the Lagrangian Grassmannian)?
- This will turn out to be a homogeneous space just like the usual Grassmannian so we need to find a group acting transitively on the space of Lagrangian subspaces.

So let's look at some groups.

Symplectic linear group

The symplectic linear group $Sp(2n)$ is the group of all linear automorphisms of $\mathbb{R}^{2n}$ which preserve the standard symplectic structure (equivalently one could look at $Sp(V)$ but since $V \cong \mathbb{R}^{2n}$ we might as well just look at $Sp(2n)$). To preserve the standard symplectic structure means

$$\omega_0(\psi v, \psi w) = \omega_0(v, w)$$

or explicitly as matrices relative to a symplectic basis

$$\psi^T \omega_0 \psi = J_0$$

where ω_0 is the matrix

$$\begin{pmatrix} 0 & \mathrm{id} \\ -\mathrm{id} & 0 \end{pmatrix}$$

in coordinates $x_1, \ldots, x_n, y_1, \ldots, y_n$.

- Note that $-\omega_0 = \omega_0^T =: J_0$ looks like the diagonal complex matrix with is on the diagonal when we write $GL(n, \mathbb{C}) \subset GL(2n, \mathbb{R})$ by $z_k = x_k + iy_k$.
- We can actually use J_0 to recover the standard Euclidean metric g_0 on $\mathbb{R}^{2n}$

$$g_0 = \omega_0(-, J_0-)$$

- The interaction of these three geometric structures: a symplectic form, a metric and a complex structure, is central to symplectic topology.

Lemma

$$Sp(2n) \cap O(2n) = Sp(2n) \cap GL(n, \mathbb{C}) = GL(n, \mathbb{C}) \cap O(2n) = U(n)$$

Proof.

This is not hard once we write out the criteria for a matrix to live in one of these groups

- $\psi \in Sp(2n)$ if and only if $\psi^T \omega_0 \psi = \omega_0$,
- $\psi \in O(2n)$ if and only if $\psi^T \psi = 1$,
- $\psi \in GL(n, \mathbb{C})$ if and only if $\psi J_0 = J_0 \psi$.

To see that the common intersection is $U(n)$, consider $GL(n, \mathbb{C}) \cap O(2n)$. A unitary matrix is a complex matrix U for which $U^\dagger U = 1$. But in our representation the conjugate-transpose operation $\dagger$ on complex n-dimensional matrices is just the transpose on their $2n$-dimensional real representatives. □

In fact, the subgroup $U(n) \subset Sp(2n)$ captures all of the topology of $Sp(2n)$:

Lemma

The inclusion $U(n) \to Sp(2n)$ is a homotopy equivalence.

- Let ψ be a symplectic matrix. It has a polar decomposition:

$$\psi = PQ = (\psi\psi^T)^{1/2}((\psi\psi^T)^{-1/2}\psi)$$

  where P is positive definite symmetric and Q is orthogonal.
- Here the square root of a positive symmetric matrix is defined by conjugating it to a diagonal matrix, taking the square roots of the diagonal entries and then conjugating back.
- The deformation retract of $Sp(2n)$ onto $U(n)$ will be given by

$$\psi_t = P^{-t/2}\psi$$

  and it remains to show that $\psi_t \in Sp(2n)$ for all t (equivalently that $P^{-t/2} \in Sp(2n)$).

- For starters, $\psi\psi^T \in Sp(2n)$ because it's a product of symplectic matrices. Therefore the lemma will follow from the assertion that $R^t \in Sp(2n)$ for any positive symmetric symplectic matrix R.
- To show this, write $\mathbb{R}^{2n}$ as a sum of eigenspaces V_λ of R (eigenvalue λ). The subspace V_λ is also the λ^t-eigenspace for R^t.
- Now let $A \in V_\lambda$ and $B \in V_\mu$:

$$\omega_0(R^t A, R^t B) = (\lambda\mu)^t \omega_0(A, B)$$

- When $t = 1$ we have $\omega_0(RA, RB) = \omega_0(A, B)$ so

$$(\lambda\mu)\omega_0(A, B) = \omega_0(A, B) \tag{2}$$

- Without loss of generality assume that $\omega_0(A, B) \geq 0$. Then taking the positive t-th root of Equation (2) and multiplying by $\omega_0(A, B)^{1-t}$ gives

$$\omega_0(R^t A, R^t B) = \omega_0(A, B)$$

  as required.

Complex structures

The useful matrix J_0 will now be generalised.

Definition

A complex structure on a vector space is a matrix J such that $J^2 = -1$. A complex structure J on a symplectic vector space (V, ω) is said to be compatible with ω if

$$g_J(-,-) = \omega(-, J-)$$

is a positive-definite metric and if ω is J-invariant, i.e.

$$\omega(J-, J-) = \omega(-,-)$$

In the absence of this latter condition we say that ω tames J. Let us write $\mathcal{J}(V, \omega)$ for the space of compatible complex structures on (V, ω).

Lemma

The space $\mathcal{J}(V,\omega)$ is equal to $Sp(2n)/U(n)$ and is therefore contractible.

Proof.

WLOG $\omega = \omega_0$. Since $U(n)$ is the stabiliser of J_0 under the action of $Sp(2n)$ on $\mathcal{J}(\mathbb{R}^{2n},\omega_0)$ it suffices to show that $Sp(2n)$ acts transitively on $\mathcal{J}(\mathbb{R}^{2n},\omega_0)$. Equivalently we need to find a symplectic basis which is g_J-orthonormal. Pick v_1 of unit length and let $v_2 = Jv_1$ (so that $g_J(v_1,v_2) = -\omega(v_1,v_1) = 0$).
Now look at the symplectic orthogonal complement of $\langle v_1, v_2\rangle$. This is preserved by J since ω_0 is preserved by J. Now we are done by induction. $\square$

Recall that an almost complex structure on a manifold is a smoothly varying choice of complex structure on each tangent space. We say that an almost complex structure on a symplectic manifold is compatible with (respectively tamed by) the symplectic form if this is true on each tangent space.

Lemma

Let (X, ω) be a symplectic manifold. The space $\mathcal{J}(X, \omega)$ of compatible almost complex structures on X is contractible.

Proof.

$\mathcal{J}(X, \omega)$ is the space of sections of a bundle with contractible fibres and hence contractible. $\square$

Exercise

$\diamondsuit$: *Find a more explicit proof using the deformation retract we constructed earlier.*

All this means that the tangent bundle of our symplectic manifold (which is a $Sp(2n)$-bundle) has a canonical lift (up to homotopy) to a $U(n)$-bundle. Therefore it has Chern classes so we can talk about the Chern classes of a symplectic manifold. Indeed the same is true of any $Sp(2n)$-bundle.

Lagrangian Grassmannian

Let $L(n)$ denote the space of all Lagrangian subspaces of a $2n$-dimensional symplectic vector space.

Lemma

$U(n)$ acts transitively on the space of Lagrangian subspaces of the standard symplectic $\mathbb{R}^{2n}$ with stabiliser $O(n)$. That is, $L(n) \cong U(n)/O(n)$.

Proof.

Suppose that L is some Lagrangian subspace in the standard symplectic $\mathbb{R}^{2n}$. If J is the standard complex structure then JL is orthogonal to L. Pick an orthonormal basis $X_1, \ldots, X_n$ of L. Then $X_i, Y_i = JX_i$ is an orthonormal symplectic basis of $\mathbb{R}^{2n}$. Since $U(n)$ acts transitively on orthonormal symplectic frames ($U(n) = Sp(2n) \cap O(n)$) this proves the lemma. $\square$

Example

Here are some examples of Lagrangian subspaces.

- The subspace $\{(x_1, 0, x_2, 0)\}$ in $\mathbb{R}^4$ when using the standard symplectic form $dx_1 \wedge dy_1 + dx_2 \wedge dy_2$.
- The graph of an antisymplectic map $\phi : (V, \omega) \to (V, \omega)$ i.e. $\phi^*\omega = -\omega$.

Here are some examples of Lagrangian submanifolds.

- The antidiagonal in $S^2 \times S^2$, the graph of the antipodal map on the sphere, where ω is the usual area form.
- $S^1 \times S^1 \subset \mathbb{C} \times \mathbb{C}$.
- The zero-section of a cotangent bundle with its standard symplectic structure.

We're now going to make some remarks about the topology of the Lagrangian Grassmannian. First think about the group $U(n)$. It has $\pi_1(U(n)) = \mathbb{Z}$ because it fibres over $U(1)$ with simply-connected fibres $SU(n)$ (via the determinant map). Now consider the fibration

$$\begin{array}{ccc} O(n) & \longrightarrow & U(n) \\ & & \downarrow \\ & & U(n)/O(n) = L(n) \end{array}$$

The homotopy long exact sequence of this fibration gives

$$\pi_1(O(n)) \to \pi_1(U(n)) \to \pi_1(L(n)) \to \pi_0(O(n)) \to \pi_0(U(n))$$

When $n = 1$ we have $L(1) = S^1/2 = \mathbb{RP}^1$. To see this, note that $L(1)$ is just the space of unoriented lines: the group $U(1)$ acts by rotating and when it gets to $\pm 1 \in O(1)$ we return to the original position.

- More generally, observe that the composition $O(n) \to U(1) \stackrel{\det}{\to} U(1)$ has image ± 1 so the inclusion $O(n) \to U(n)$ kills $\pi_1(O(n))$.
- Therefore $\pi_1(L(n))$ is an extension of $\mathbb{Z}/2 = \pi_0(O(n))$ by $\mathbb{Z} = \pi_1(U(n))$ by the exact sequence above.
- One can see explicitly that it is the extension $\mathbb{Z} \to \mathbb{Z} \to \mathbb{Z}/2$ by looking at the π_1-injective subgroup $U(1) \subset U(n)$: half-way around we get to $-1 \in O(n)$.

- So in all cases, $\pi_1(L(n)) = \mathbb{Z}$.
- This isomorphism is actually canonical: to see this pick an isomorphism of the symplectic vector space with the standard one (the space of these is connected) and take as a generator the loop $\phi_t \mathbb{R} \times \mathbb{R}^{n-1} \subset \mathbb{R}^{2n}$ where ϕ_t is an anticlockwise rotation by $t\pi$ radians in one of the symplectic factors of $\mathbb{R}^{2n}$ (which makes sense because the symplectic subspace is canonically oriented).
- Therefore we can actually assign an integer to a loop $\gamma : S^1 \to L(n)$ of Lagrangian subspaces. We call it the Maslov index of γ.

Suppose that $u : D^2 \to X$ is an immersed disc in a symplectic manifold whose boundary is contained in a Lagrangian submanifold L. Then we can trivialise $u^* TX$ symplectically (i.e find an isomorphism with the standard product bundle $D^2 \times \mathbb{R}^{2n}$) and the loop of tangent spaces to L pulls back to a loop of Lagrangian subspaces in $\mathbb{R}^{2n}$. We can thereby define the Maslov index of a disc with Lagrangian boundary conditions.

Exercise

$\diamondsuit$: *Compute the Maslov index of the unit disc in* $\mathbb{C}$.

Exercise

$\diamondsuit$: *We have seen that $U(n)$ acts transitively on Lagrangian subspaces. Show that $Sp(2n)$ acts 2-transitively on transverse Lagrangian subspaces (i.e. any two transverse Lagrangian subspaces can be moved simultaneously by a symplectic linear map to the two standard Lagrangians $\mathbb{R}^n \times \{0\}$ and $\{0\} \times \mathbb{R}^n$ in $\mathbb{R}^{2n}$.*

Exercise

$\heartsuit$: *Let's work in $\mathbb{R}^4$ with the standard symplectic structure. We say two symplectic planes S_1 and S_2 are simultaneously holomorphisable if there exists an ω-compatible complex structure J for which $JS_1 = S_1$ and $JS_2 = S_2$. First show that this implies $S_1 \pitchfork S_2$. Now assume S_1 is the standard (x_1, y_1)-plane and write S_2 as the graph of a 2-by-2 matrix M (why are we justified in doing this?). What condition on M implies that S_2 is symplectic? What condition on M implies S_1 and S_2 are simultaneously holomorphisable? What does this say about the action of $Sp(2n)$ on pairs of transverse symplectic planes?*