Symplectic Topology

Lecture Notes

ZWEISTEIN

Desu-Cartes

Contents

1 Analytic and Geometric Prerequisites

The goal of this lecture is to discuss symplectic embeddings. To do this, we do quite some preparatory work.

1.1 Differential Forms

We quickly review the calculus of differential forms here. We will mainly work in $\mathbb{R}^n$ here although the formalism can be generalized to arbitrary smooth manifolds. Accordingly, we start this section with a quick review of multilinear algebra.

Definition 1.1 (Dual basis). Let $\{e_1, \ldots, e_n\}$ denote the standard basis of $\mathbb{R}^n$. The *algebraic dual space* of $\mathbb{R}^n$ is the vector space of linear forms $\alpha : \mathbb{R}^n \to \mathbb{R}$. We will usually denote it $(\mathbb{R}^n)^*$. The *dual basis* of $(\mathbb{R}^n)^*$ usually written $\{e_1^*, \ldots, e_n^*\}$ is defined to be such that $e_i^*(e_j) = \delta_{ij}$. In other words, e_i^* takes a vector and spits out its i-th coefficient.

Definition 1.2. We say that a multilinear map $\alpha : \mathbb{R}^n \times \cdots \times \mathbb{R}^n$ is *skew-symmetric* if

$$\alpha(v_1, \ldots, v_i, v_j, \ldots, v_n) = -\alpha(v_1, \ldots, v_j, v_i, \ldots, v_n).$$

In other words, exchanging elements flips the sign.

Ideally, we want to put together multiple basis elements of $(\mathbb{R}^n)^*$. We will do so by considering elements of the form $e_1^* \wedge e_2^* \wedge \cdots \wedge e_n^*$. More specifically, we set

Definition 1.3 (Wedge product). Let $\{e_1^*, \ldots, e_n^*\}$ denote the dual basis of $(\mathbb{R}^n)^*$. Let $I = (i_1, \ldots, i_k)$ be a subset of $1, \ldots, n$. The *wedge product* of two bases elements is defined as

$$e_i^* \wedge e_j^*(v_1, v_2) := \det \begin{bmatrix} e_i(v_1) & e_i(v_2) \\ e_j(v_1) & e_j(v_2) \end{bmatrix}$$

More generally, we get

$$e_{i_1}^* \wedge \cdots \wedge e_{i_k}^*(v_1, \ldots, v_k) = \det \begin{bmatrix} e_{i_1}^*(v_1) & \cdots & e_{i_1}^*(v_k) \\ \vdots & & \vdots \\ e_{i_k}^*(v_1) & \cdots & e_{i_k}^*(v_k) \end{bmatrix}$$

In the case where $k = 2$, this has the very nice geometric interpretation of sending two vectors to the signed area of the parallelogram delimited by the vectors. All of this allows us to define the vector space that will interest us.

Definition 1.4 (Exterior algebra). The *k-th exterior algebra* of $(\mathbb{R}^n)^*$ is the vector space of all skew-symmetric multilinear forms from $\mathbb{R}^n$ to $\mathbb{R}$ equipped with the wedge product. We usually denote it $\wedge^k(\mathbb{R}^n)^*$. The wedge product is associative and defined such that $\alpha \wedge \beta = (-1)^{kl}\beta \wedge \alpha$ where α takes k vectors, and β takes l vectors.

It turns out there is still one important construction we have yet to discuss. We have discussed dual spaces but as it turns out, this is more than a simple construction. Recall from linear algebra that linear maps are the morphisms in the category Vect_k of k-vector spaces.

Proposition 1.5. *The correspondence $\mathcal{D} : V \mapsto V^*$ is a contravariant functor from the category Vect_k to itself.*

There are two important aspects to this proposition. The first one is that the correspondence is a functor. Interestingly enough, this means our maps have to be sent somewhere as well. The other one is that the functor is contravariant. This means that if we have a map $\phi : V \to W$, we need to find a map $\phi^* : W^* \to V^*$. Luckily, there is a very obvious candidate: just take the transpose. Indeed, this is the construction we're looking for. Our functor thus does the following. It takes a vector space V and sends it to V^*, and takes a map $\phi : V \to W$ and sends it to $\phi^* : W^* \to V^*$. This is but one example of duality in category theory. Notice that applying our functor twice to V gives us V^{**} which is canonically isomorphic to V because our vector spaces are finite-dimensional. More concretely, we have that $(\phi^* \ell) = \ell(\phi)$ so it acts as pre-composition in a sense. We want to emulate this for our vector space $\wedge^k(\mathbb{R}^n)^*$.

Definition 1.6 (Pull-back). Let $\phi : \mathbb{R}^n \to \mathbb{R}^m$ be a linear map and $\alpha \in \wedge^k(\mathbb{R}^n)^*$. We define the *pull-back* of ϕ to be

$$\phi^*(\alpha)[v_1, \ldots, v_k] = \alpha\left(\phi(v_1), \ldots, \phi(v_k)\right).$$

We are now ready to define differential forms. We have restricted our attention to $\mathbb{R}^n$ but the general theory is not much harder. For our purposes, it will be enough to consider our manifolds as embedded in real space which explains the choice of sticking with $\mathbb{R}^n$.

Definition 1.7 (Differential k-form). A *differential k-form* is a smooth map

$$\alpha : \mathbb{R}^n \to \wedge^k(\mathbb{R}^n)^*.$$

We will usually write α_p to mean α applied at the point $p \in \mathbb{R}^n$.

The set of all differential k-forms on $\mathbb{R}^n$ is denoted by $\Omega^k(\mathbb{R}^n)$. This is a real infinite-dimensional vector space. We also set $\Omega^0(\mathbb{R}^n) := C^\infty(\mathbb{R}^n)$, the space of smooth functions on $\mathbb{R}^n$.

Remark 1.8. Unpacking the definition, we're saying that a differential k-form takes k vectors in $\mathbb{R}^n$ and spits out a number in $\mathbb{R}$. Furthermore, the form should also be multilinear, skew-symmetric and smooth. This definition is problematic however, because it's not so clear what is meant for such a map to be smooth. One way to understand it is via vector fields.

Definition 1.9 (Vector field). Let M be a smooth manifold. A vector field X on M is a smooth section of the tangent bundle of M. In other words, it is a map that takes a point p and sends it to a vector in T_pM, where T_pM denotes the tangent space of M at p.

Example 1.10. Let $M = \mathbb{R}^n$. A vector field $X = (X_1, \ldots, X_n)$ takes a point p and sends it to $(X_1(p), \ldots, X_n(p))$. Notice that here, $T_p\mathbb{R}^n \equiv \mathbb{R}^n$. Effectively, this means we're taking a point to a vector in its tangent space, but in $\mathbb{R}^n$ they are the same thing. To simplify notation, we write $X_i(p)$ as $(X_i)_p$.

As we have seen, our differential k-forms can eat vectors and spit out numbers. As vector fields take a point and spit out vectors, we can instead apply differential k-forms to vector fields. This gives us the following definition of smooth: A k-form α is smooth if the function

$$p \longmapsto \alpha_p\left((X_1)_p, \ldots, (X_n)_p\right)$$

is smooth for all smooth vector fields. This resolves the issue and insures our definition is well-posed. Lastly, we discuss bases. As we have said before, $\Omega^k(\mathbb{R}^n)$ is an infinite-dimensional vector space. When restricting to a specific tangent space $T_p\mathbb{R}^n$ however, we can define a basis for our differential k-forms.

Definition 1.11. Let $I = (i_1, \ldots, i_k)$ be a subset of $\{1, \ldots, n\}$. We define the basis elements of $\Omega^k(\mathbb{R}^n)$ at p to be

$$\left(\mathrm{d}x_{i_1} \wedge \cdots \wedge \mathrm{d}x_{i_k}\right)_p : T_p\mathbb{R}^n \times \cdots \times T_p\mathbb{R}^n \to \mathbb{R}^n.$$

They can be defined using the basis elements of $(T_p\mathbb{R}^n)^*$, namely:

$$\left(\mathrm{d}x_{i_1} \wedge \cdots \wedge \mathrm{d}x_{i_k}\right)_p (v_1, \ldots, v_k) = \det \begin{bmatrix} e^*_{i_1}(v_1) & \cdots & e^*_{i_1}(v_k) \\ \vdots & & \vdots \\ e^*_{i_k}(v_1) & \cdots & e^*_{i_k}(v_k) \end{bmatrix}$$

In a sense, we defined a global basis that restricts to our usual dual basis when zooming in on a specific tangent space. If we look closely, we will notice that $\mathrm{d}x_i$ is just the projection on the i-th coordinate which is exactly what we wanted. For instance, if $X = (X_1, X_2)$ is a vector field, we have $\mathrm{d}x_1(X) = X_1$. Another detail we need to mention is that coefficients in $\Omega^k(\mathbb{R}^n)^*$ are not numbers anymore, but smooth functions. The wedge product also naturally transfers and we have

$$(\alpha \wedge \beta)_p = \alpha_p \wedge \beta_p.$$

Remark 1.12. At this point, you might be scratching your head. Why are we doing this exactly? You might have noticed that our basis elements look awfully familiar to the usual infinitesimals from calculus. Indeed, $\mathrm{d}x$ looks just like the element used to perform integration. We do write $\int_a^b f(x)\,\mathrm{d}x$ after all. This isn't a coincidence.

If you have taken a course on differential geometry, you know that you can associate to every function $f : \mathbb{R}^n \to \mathbb{R}$ a linear map called the *differential* $\mathrm{d}f_p$ (surprise, surprise) at a point p. This differential encodes the information of all the directional derivatives at p. Indeed, we have

$$\mathrm{d}f_p(v) = \lim_{h \to 0} \left(\frac{f(p + hv) - f(p)}{h} \right).$$

This map varies smoothly with the base point and is obviously linear, which means that the map

$$\mathrm{d}f : \mathbb{R}^n \to \wedge^1(\mathbb{R}^n)^*, \quad p \mapsto \mathrm{d}f_p$$

is a differential 1-form. Accordingly, we can decompose it in the basis $\mathrm{d}x_1, \ldots, \mathrm{d}x_n$. We know from calculus that

$$\mathrm{d}f_p(v) = \langle \nabla f(p), v \rangle = \frac{\partial f}{\partial x_1}(p)v + \cdots + \frac{\partial f}{\partial x_n}v.$$

Effectively, this means that

$$\mathrm{d}f = \frac{\partial f}{\partial x_1}\mathrm{d}x_1 + \cdots + \frac{\partial f}{\partial x_n}\mathrm{d}x_n.$$

EXERCISE 1.13. Are all 1-forms differentials of functions? Set $\alpha := \alpha_1(x,y)\mathrm{d}x + \alpha_2(x,y)\mathrm{d}y$ and investigate. (Remember, the sum of differential k-forms is also a differential k-form as they form vector spaces).

In the linear case, we have defined the pullback of a function. We can extend this definition in the case of differential forms. As is common in differential geometry, we go from the smooth case to the linear case by differentiating. We denote the differential of a map $\phi : \mathbb{R}^m \to \mathbb{R}^n$ (sometimes called the pushforward) at a point p as $D\phi(p)$. In the case of real vector spaces, this is just the Jacobian.

DEFINITION 1.14. Let $\alpha \in \Omega^k(\mathbb{R}^m)$ be a differential k-form and $\phi : \mathbb{R}^n \to \mathbb{R}^m$ a smooth map. We define the *pull-back* of the form α via ϕ to be

$$(\phi^*\alpha)_p(X_1, \ldots, X_k) = \alpha_{\phi(p)}\left(D\phi(p)X_1, \ldots, D\phi(p)X_k\right).$$

This gives us a new differential k-form. Notice that the point goes from p to $\phi(p)$ because $D\phi(p) : T_p\mathbb{R}^n \to T_\phi(p)\mathbb{R}^m$.

EXERCISE 1.15. Show that $\phi^*(\alpha \wedge \beta) = \phi^*\alpha \wedge \phi^*\beta$.

EXAMPLE 1.16. At this point, an example would probably be a good idea. Let ϕ be the polar coordinate transformation, that sends

$$\begin{bmatrix} r \\ \vartheta \end{bmatrix} \longmapsto \begin{bmatrix} r\cos\vartheta = x \\ r\sin\vartheta = y \end{bmatrix}$$

. It is a classical fact from vector analysis that its differential is the matrix

$$D\phi(r,\vartheta) = \begin{bmatrix} \cos\vartheta & -r\sin\vartheta \\ \sin\vartheta & r\cos\vartheta \end{bmatrix}$$

Using the formula where we write ϕ_* instead of $D\phi$ for brevity, we get

$$\begin{aligned} \phi^*(\mathrm{d}x \wedge \mathrm{d}y) &= \det\phi_*(r,\vartheta) \cdot (\mathrm{d}r \wedge \mathrm{d}\vartheta) \\ &= r\mathrm{d}r \wedge \mathrm{d}\vartheta \end{aligned}$$

which should look familiar if you've ever used the change of variable formula from vector analysis. In a sense, we've said that $\mathrm{d}x\mathrm{d}y$ transform into $r\mathrm{d}r\mathrm{d}\vartheta$. This is actually completely true formally, but since we haven't defined integration yet, we can't really talk about it. Another less tedious way of solving this would've been to use the fact that the wedge product and the pull-back commute.

EXERCISE 1.17. Show the same result using the fact that $\phi^*(\mathrm{d}x \wedge \mathrm{d}y) = \phi^*\mathrm{d}x \wedge \phi^*\mathrm{d}y$. Hint: Remember that $\mathrm{d}x$ means the projection on the x-coordinate.

1.2 The Calculus of Differential Forms

We have now ported every tool we had in the linear case over to the smooth case. It is time to see how the smooth case changes. The most obvious difference is that we have ways of differentiating k-differential forms. The insight to be gained here is that differential calculus measures how our function changes when we change our base point. This is why it makes no sense in the linear context, as there is no dependence on a base point. We start with the exterior derivative, that we construct from scratch. We have seen above that to each smooth function f, we can associate a 1-form $\mathrm{d}f$. This construction is important and we would like to generalize it. For 0-forms (smooth functions), we define this to be the exterior derivative, that is to say,

$$\mathrm{d}f := \sum_{i=1}^{n} \frac{\partial f}{\partial x_i}\mathrm{d}x_i$$

Since we know how to differentiate smooth functions, and those are the coefficients of forms of higher degree (the degree being k for a k-differential form), we have a very simple way of defining d for k-differential forms.

DEFINITION 1.18 (Exterior derivative). Let $\alpha = \sum_I f_I \mathrm{d}x_I$ be a k-differential form where $|I| = k$ and $\mathrm{d}x_I = \mathrm{d}x_{i_1} \wedge \cdots \wedge \mathrm{d}x_{i_k}$. We define the *exterior derivative* to be the map $\mathrm{d} : \Omega^k(\mathbb{R}^n) \to \Omega^{k+1}(\mathbb{R}^n)$ that sends

$$\sum_I f_I \mathrm{d}x_I \longmapsto \sum_I \mathrm{d}f_I \wedge \mathrm{d}x_I.$$

EXAMPLE 1.19. Let $\alpha = f(x,y,z)\mathrm{d}x \wedge \mathrm{d}y \wedge \mathrm{d}z$. Then $\mathrm{d}\alpha = \mathrm{d}f \wedge \mathrm{d}x \wedge \mathrm{d}y \wedge \mathrm{d}z$. For instance, if $f(x)$ is the identity function $f(x) = x$, then $\mathrm{d}f(x) = \mathrm{d}x$, our usual projection.

This derivative enjoys a few nice properties.

Proposition 1.20. *The exterior derivative satisfies the following properties:*

1. *We have that $d^2 = d \circ d = 0$. This means that applying the exterior derivative to a form twice gives us zero.*

2. *The exterior derivative is compatible with the wedge product: $d(\alpha \wedge \beta) = d\alpha \wedge \beta + (-1)^k \alpha \wedge d\beta$, for a k-form α and an ℓ-form β.*

3. *The exterior derivative commutes with the pull-back: If $\phi : \mathbb{R}^n \to \mathbb{R}^m$ is a smooth map and $\alpha \in \Omega^k(\mathbb{R}^m)$, then $d\phi^*\alpha = \phi^* d\alpha$.*

Definition 1.21. We say that a differential k-form α is *closed* if $\mathrm{d}\alpha = 0$. If furthermore $\alpha = \mathrm{d}f$ for some smooth function f which we call a *primitive*, we say that α is *exact*.

Notice that since $\mathrm{d}^2 = 0$, we have that exact implies closed. Understanding when closed implies exact is the impetus behind de Rham cohomology.

Exercise 1.22. Show that the 2-form $\eta_{(x,y)} = \frac{x\mathrm{d}y - y\mathrm{d}x}{x^2+y^2}$ is closed, but not exact. Hint: Try pulling it back using polar coordinates like we did above. Can you tell what this form represents?

An important point is that our way of differentiating differential k-forms yields $(k+1)$-forms. Is there a way to differentiate forms in a way that preserves the degree? The answer is yes, but it is of a completely different nature. We will generalize the idea behind the directional derivative of vector calculus. First, we need to discuss flow.

Definition 1.23. Let X be a set. We define a *flow* on X to be an action of the group $(\mathbb{R}, +)$ on X. This is a mapping $\varphi_t(x) : \mathbb{R} \times X \to X$ such that $\phi_0(x) = x$ and $\phi_s \circ \phi_t = \phi_{s+t}$.

As it turns out, we can always associate a flow locally to a vector field X.

Proposition 1.24. *Let $X : T\mathbb{R}^n \to \mathbb{R}^n$ be a vector field. We can always associate locally a flow to X, that is to say, there exists a map*

$$\varphi : (-\varepsilon, \varepsilon) \times \mathbb{R}^n \to \mathbb{R}^n, \quad (t, x) \mapsto \varphi_t(x)$$

such that

$$\left.\frac{d}{dt}\right|_{t=0} \varphi_t(x) = X(x), \quad \varphi_0(x) = x.$$

It is with this construction that we can define our other derivative.

Definition 1.25 (Lie derivative). For X any vector field, we define $\mathcal{L}_X : \Omega^k(\mathbb{R}^n) \to \Omega^k(\mathbb{R}^n)$, the *Lie derivative in the direction of* X by setting

$$\mathcal{L}_X \alpha = \left.\frac{\mathrm{d}}{\mathrm{d}t}\right|_{t=0} \varphi_t^* \alpha.$$

While this formula is perfectly okay that way, there is a simpler way to compute the Lie derivative.

Definition 1.26 (Interior product). Let X be a vector field on $\mathbb{R}^n$ and α a k-differential form. The *interior product* is the map ι_X that sends α to the $k-1$-form defined by the property that

$$(\iota_X \alpha)(X_1, \ldots, X_n) = \alpha(X, X_1, \ldots, X_n).$$

What this does is contract the differential form α and the vector field X. This effectively fills the first spot in the k-differential form α with X. Using this product, we arrive at a really surprising and interesting result.

Theorem 1.27 (Cartan's magic formula)**.** *We have*

$$\mathcal{L}_X = \iota_X \circ d + d \circ \iota_X.$$

Proof. First, it is easy to see that this holds for functions and exact one-forms. Indeed, $\mathcal{L}_X(f) = \iota_X \mathrm{d}f$ and $\mathrm{d}(\mathrm{d}f) = 0$. Now notice that k-differential forms are nothing but sums of wedges of exact 1-forms multiplied by smooth functions. If we can show that the formula is compatible with the wedge product, we're done. To this end, define $D_X := \iota_X \circ \mathrm{d} + \mathrm{d} \circ \iota_X$. First, it is clear that $\mathcal{L}_X(\alpha \wedge \beta) = \mathcal{L}_X(\alpha) \wedge \beta + \alpha \wedge \mathcal{L}_X(\beta)$ from the definition of the Lie derivative and the fact that the wedge product is bilinear. Now notice that our operator D_X satisfies the same property. Suppose that α is a k-differential form. This holds true because $\mathrm{d}(\alpha \wedge \beta) = (\mathrm{d}\alpha) \wedge \beta + (-1)^k \alpha \wedge (\mathrm{d}\beta)$ and we also have $\iota_X(\alpha \wedge \beta) = (\iota_X \alpha) \wedge \beta + (-1)^k \alpha \wedge (\iota_X \beta)$. This in turn shows that $(\mathcal{L}_X - D_X)\alpha = 0$ by decomposing in elementary summands of the form $f \mathrm{d}x_{i_1} \wedge \cdots \wedge \mathrm{d}x_{i_k}$. This completes the proof. □

Having gone over ways to differentiate k-differential forms, we now look for ways to integrate them. We already know how to integrate smooth functions, but in a sense, we even know a little more. Let γ be a smooth curve in $\mathbb{R}^n$. We have seen in calculus how to integrate over such curves. There is a very natural way to generalize this. As a small disclaimer, we'll assume in this chapter that all our forms are have compact support to avoid any technicalities.

Definition 1.28 (Integration of 1-forms). Let $\gamma : [a, b] \to \mathbb{R}^n$ be a smooth curve in $\mathbb{R}^n$ and $\alpha \in \Omega^1(\mathbb{R}^n)$. We define *the integral of* α *along* γ to be

$$\int_\gamma \alpha := \int_a^b \alpha_{\gamma(t)} \left(\dot{\gamma(t)} \right) dt.$$

The fundamental theorem of calculus says that the integral of an exact 1-form $\mathrm{d}f$ is equal to the difference of fs evaluated at its end points.

Exercise 1.29. What does the integral of the one form $y\,\mathrm{d}x$ represent? Does it look familiar? What's different?

Suppose now that α is a k-dimensional form. It is important to realize that we will only be able to integrate α over k-dimensional submanifolds. The simplest case is that of an open set.

Definition 1.30 (Integration of k-dimensional forms). Let $U \subseteq \mathbb{R}^k$ be an open set and let $\alpha = f(x_1, \ldots, x_k)\mathrm{d}x_1 \wedge \cdots \wedge \mathrm{d}x_k$ be a k-differential form. The integral of α is the oriented usual integral in $\mathbb{R}^k$, that is to say:

$$\int_U \alpha := \int_U f(x_1, \ldots, x_k)\, \mathrm{d}x_1, \ldots, \mathrm{d}x_k.$$

The only real difference is that by skew-symmetry, exchanging two of the $\mathrm{d}x_i$ flips the sign of the integral. For instance,

$$\int_U g(x,y)\mathrm{d}x \wedge \mathrm{d}y = -\int_U g(x,y)\mathrm{d}y \wedge \mathrm{d}x.$$

Extending the integral to submanifolds is not hard now that we know how to do this.

Definition 1.31 (Integration over submanifolds). Let $\varphi_U : U \to V \subseteq M$ be a chart for a k-dimensional submanifold M. Let $\alpha \in \Omega^k(U)$. We define the integral over V to be

$$\int_V \alpha = \int_U \varphi_V^* \alpha.$$

Notice how we have transferred the problem of integrating in a manifold to a simple integral in $\mathbb{R}^k$ via the pull-back of the form. Chaining those integrals together via a partition of unity lets us integrate over the entire manifold M. Lastly, we wish to discuss the cornerstone of vector calculus, the celebrated theorem of Stokes. Most so-called fundamental theorems in calculus such as the FTC, Green-Riemann and so on are just applications of this theorem. Here it is in its full glory.

Theorem 1.32. *Let $(M, \partial M)$ be an oriented $(k-1)$-dimensional manifold with boundary and let $\alpha \in \Omega^{(k-1)}(M)$ (with compact support). Then,*

$$\int_M \alpha = \int_{\partial M} d\alpha.$$

Exercise 1.33. Derive the fundamental theorem of calculus from this formula. Hint: Take $\alpha = f(x)$ and $M = [a, b]$. What about the divergence theorem and Green-Riemann?

1.3 Cohomology

In a first course on homology, we usually discuss multiple flavors of homology functors. Usually, we discuss (semi-)simplicial homology first and extend our simplices to get singular homology. Another particular kind of homology that's very nice for computations is cellular homology. We have a few ways of computing actual examples:

1. $H_*^{\mathrm{cell}}(X)$, the cellular homology groups,
2. The long exact sequence,
3. The Mayer-Vietoris sequence,
4. Künneth's Formula.

Perhaps the reader is not familiar with the last one which states that $H_k(X \times Y, K) = \sum_{i+j=k} H_i(X,K) \otimes H_j(Y,K)$, where K is a field. The usual homology construction uses chains. Cohomology is the dual notion: what happens if we consider cochains instead? It is often the case that cohomology is a more natural tool that homology. Indeed, cohomology has a few advantages to homology that we will discuss now. Suppose we have a chain:

$$\cdots \longleftarrow C^{k-1} \longleftarrow C^k \longleftarrow C^{k+1} \longleftarrow \cdots$$

How do we make the arrows go the other direction? Well one idea is to apply a Hom functor. Indeed, we have

$$\cdots \longrightarrow C^{k-1} \longrightarrow C^k \longrightarrow C^{k+1} \longrightarrow \cdots$$

where $C^k := \mathrm{Hom}(C_k, R)$ where R is a ring.

EXERCISE 1.34. Show that if $R = K$ a field, then $H_i(X,K) \equiv H^i(X,K)$.

We write d_i for the induced homomorphism between Hom-sets. One important aspect of (singular) cohomology is that there is a striking link to analysis. As we have seen earlier, the exterior derivative squares to zero. This means that we can define a cohomology on the spaces of differential forms.

DEFINITION 1.35 (De Rham cohomology). Let $\mathrm{d}_k : \Omega^k(M) \to \Omega^{k+1}(M)$ denote the exterior derivative. The cochain complex

$$\cdots \longrightarrow \Omega^{k-1} \xrightarrow{\mathrm{d}_{k-1}} \Omega^k \xrightarrow{\mathrm{d}_k} \Omega^{k+1} \longrightarrow \cdots$$

induces cohomology groups $H^k_{\mathrm{dR}}(M) := \ker(\mathrm{d}_i)/\mathrm{im}(\mathrm{d}_{i-1})$.

We usually denote the cycle corresponding to ω by $[\omega]$. So what do those equivalence classes look like? We say that two forms α and β in $\Omega^k(M)$ are *cohomologous* if they differ by an exact form, that is to say, $\alpha - \beta$ is exact. This is all good and well, but this is just another cohomological theory, what does it have to do with singular cohomology? The key theorem is that of Georges de Rham.

Theorem 1.36 (De Rham's Theorem)**.** *Let M be a smooth manifold. Then the map*

$$\Psi^* : H^k_{dR}(M) \to H^k(M; \mathbb{R})$$
$$[\omega] \longmapsto \left([c] \mapsto \int_c \omega\right)$$

is an isomorphic of vector spaces.

Note here that the cycle $[c]$ is in H_k, not in H^k_{dR}. This theorem provides us with analytic tools to understand the topology of manifolds. Interestingly enough, this also shows that the differential forms depend on the manifold M, but *not* on its smooth structure. This is one clear advantage of working with cohomology.

2 Symplectic Structures

2.1 Symplectic Geometry

We finally start our study of symplectic geometry in earnest. Now that we are familiar with differential forms and how to work with them, we can finally define what symplectic geometry is about. Following the astonishing work of Felix Klein, modern geometry is commonly understood as studying a space and a group of transformations associated with it. In Riemannian Geometry for instance, we study smooth manifolds equipped with a Riemannian metric. It is not so much the space that is interesting but the structure on that space, as is often the case in mathematics as a whole. The question then becomes, what is the inherent structure on a symplectic manifold? As it turns out, it's a special kind of 2-form.

DEFINITION 2.1 (Symplectic form on M). A *symplectic form* on M is a closed non-degenerate 2-form ω.

What do we mean by non-denegerate exactly? From the 2-form, we can induce a bilinear pairing on the tangent spaces T_pM. The form being non-degenerate means that if $\omega(v, w) = 0$ for all $w \in T_pM$, then $v = 0$. Notice that we need our manifold to be even-dimensional here for this to make sense. Let's look at the simplest example there is, $\mathbb{R}^{2n}$.

EXAMPLE 2.2. Let $M = \mathbb{R}^{2n}$. This is a symplectic manifold (usually called a symplectic vector space). The canonical symplectic form is $\omega = \sum_{i=1} \mathrm{d}x_i \wedge \mathrm{d}y_i$ where $\mathbb{R}^{2n} = (x_1, \dots, x_n, y_1, \dots, y_n)$. Remember that $\mathrm{d}x_i = e_i^*$, so applying our standard form is akin to taking multiple determinants.

Having defined what symplectic spaces look like, it is only natural to ask what the morphisms look like in this category. The structure is the symplectic form, so a morphism should preserve this.

DEFINITION 2.3. Suppose (M, ω) and (N, ω') are two symplectic manifolds. A *symplectomorphism* is a smooth map $\varphi : M \to N$ such that

$$\varphi^*\omega = \omega'$$

If M and N are nothing but $\mathbb{R}^{2n}$, we usually say the symplectomorphisms are linear. Those linear symplectomorphisms form a group that we denote $\mathrm{Sp}(2n)$. In dimension 1, this is nothing but the volume-preserving maps, but in higher dimensions, symplectomorphisms have more structure. Indeed, any linear symplectomorphism is volume-preserving.

EXERCISE 2.4. Show that if $\varphi \in \mathrm{Sp}(2n)$, then $\det \varphi = 1$. Conclude that any linear symplectomorphism preserves volume.

Historically, these notions came out of classical mechanics. More specifically, the key equation in classical mechanics is Newton's Law.

Principle 2.5 (Newton's Law)**.**

$$F = m\ddot{x}.$$

It is not immediate what this equation has to do with our symplectic manifolds. And indeed, it doesn't really look like geometry at this point. It turns out that Lagrange (and especially Hamilton) managed to give another equation (equivalent to Newton's) that also governs classical mechanics. Instead of considering the object's position function, we instead focus our attention on the energy of the system. More formally, let us define a new set of coordinates that we will call canonical coordinates. We let $r = (p, q)$ where q is the vector of so-called generalized coordinates, and p is the vector of conjugate momenta. What matters is that, as you can see, this is even-dimensional. This is because if $p = (p_1, p_2, p_3)$ (say), then $q = (q_1, q_2, q_3)$. The Hamiltonian function is then defined as $H = T + V$ where T is potential energy and V is kinetic energy. It turns out that the phase space of our coordinates is a symplectic manifold. Locally, the form is given as $\omega = \sum_{i=1}^{n} \mathrm{d}p_i \wedge \mathrm{d}q_i$. We can now reformulate Newton's equations in a set of two equations.

Proposition 2.6. *Newton's Law is equivalent to the following:*

$$\begin{cases} \dot{x}(t) = \frac{\partial H_t}{\partial y}(x(t), y(t)) \\ \dot{y}(t) = -\frac{\partial H_t}{\partial x}(x(t), y(t)) \end{cases}$$

where $H_t(x(t), y(t)) = \frac{1}{2}\left\|y\right\|^2 - V_t(x)$.

The first term is kinetic energy, the second is potential energy. And now we can kind of see why symplectic geometry comes in. Notice how the second term has a negative sign. This is key. We can now define a vector field on $\mathbb{R}^{2n}$ in the following way.

Proposition 2.7. *Let* ω_0 *be the standard symplectic form on* $\mathbb{R}^{2n}$*. We define the vector field* X_{H_t} *to be the one that satisfies the following equation:*

$$\omega_0(X_{H_t}(z), \cdot) = dH_t(z)(\cdot).$$

More concretely, we have

$$X_{H_t}(z) = \left(\frac{\partial H}{\partial y}(z), -\frac{\partial H}{\partial x}(z)\right).$$

We have $\dot{z}(t) = X_{H_t}(z(t))$ which gives us a flow φ_H^t. The important fact is now to notice that this flow is symplectic.

Proposition 2.8. *The flow* φ_H^t *is symplectic.*

Proof. We want to show that $(\varphi_H^t)^*\omega_0 = \omega_0$. Since φ_H^0 is the identity, it is enough to show that $\frac{\mathrm{d}}{\mathrm{d}t}(\varphi_H^t)^*\omega_0 = 0$. Using Cartan's formula, we get

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t}(\varphi_H^t)^*\omega_0 &= (\varphi_H^t)^*\mathcal{L}_{X_H}\omega_0\\ &= (\varphi_H^t)^*(\mathrm{d}\iota_{X_H} + \iota_{X_H}\mathrm{d})\omega_0\\ &= (\varphi_H^t)^*(\mathrm{d}(\omega(X_H,\cdot)) + \iota_{X_H}0)\\ &= (\varphi_H^t)^*(\mathrm{d}(\mathrm{d}H(\cdot)))\\ &= 0.\end{aligned}$$

where $\mathrm{d}\omega_0 = 0$ since the form is closed. ☐

Exercise 2.9. Show Liouville's Theorem: $\mathrm{Vol}(\varphi_H^t(U)) = \mathrm{Vol}(U)$ with U in phase space.

2.2 Symplectic Embeddings

So far, we have only discussed the geometry of symplectic manifolds. It turns out topological properties of symplectic manifolds are also quite interesting. The problem of symplectic embeddings give us a way to understand how rigid symplectic structures actually are. Interestingly, the first fundamental result will require us to step into the world of complex geometry.

Definition 2.10 (Almost complex structures). Let M be a smooth manifold. An *almost complex structure* on M is a smooth vector bundle endomorphism $J : TM \to TM$ such that $J^2 = -\mathrm{Id}$. Effectively, this means a family of linear maps $J_p : T_pM \to T_pM$ such that $J_p \circ J_p = -\mathrm{Id}_{T_pM}$.

The idea will then be to equip our symplectic manifold (M, ω) with an almost complex structure J to show interesting things about M. We first need to settle the question of existence. We assume our symplectic manifold also possesses a Riemannian metric g.

Proposition 2.11. *Let (M, ω, g) be a symplectic manifold. We set $\mathcal{J}(M) := \{J$ almost complex structure on $M\}$. Then, $\mathcal{J}(M)$ is non-empty and contractible.*

Proof. Let $g(M)$ denote all the metrics we can put on M. By convexity, $g(M)$ is contractible. We can always construct our almost-complex structure from our symplectic form and our metric. Indeed, we have $J(u) = (\iota_u(g))^{-1}(\iota_u(\omega(u))$. Since $\omega \to \omega$ is the identity, and $g \mapsto g'$ is a homeomorphism, we indeed have that $\mathcal{J}(M)$ is contractible. ☐

Definition 2.12 (ω-tame almost complex structures). Let (M, ω, J) be a symplectic manifold with an almost complex structure J. We say that J is *ω-tame* if $\omega(u, Ju) > 0$ for all $u \neq 0$. The family of all ω-tame almost complex structures is denoted by $\mathcal{J}_\omega(M)$.

This lets us define the key notion that will help us prove our principal theorem.

DEFINITION 2.13 (J-holomorphic curves). Let (M, ω, J) be a symplectic manifold with almost complex structure J. A curve u is said to be *J-holomorphic* if we have

$$J \circ \mathrm{d}_u = \mathrm{d}_u \circ i$$

where i is the canonical complex structure on $\mathbb{C}$.

We are now ready to outline the most fundamental theorem in the theory.

Theorem 2.14 (Gromov's Nonsqueezing Theorem)**.** *If $B^{2n}(a)$ embeds symplectically in $Z^{2n}(A)$, then $a \leq A$. Here $Z^{2n}(A)$ means the $2n$-dimensional cylinder whose disk section has area A.*

The proof involves moduli spaces and cohomology. We sketch it here in case the reader is sufficiently proficient with said material to understand it and fill in the gaps.

Proof. Suppose first that our symplectic embedding φ also preserves the standard complex structure on $\mathbb{C}^n$, that is $J_0 = i \oplus i \oplus i \oplus \cdots \oplus i$. First, notice that $Z^{2n} = D(A) \times \mathbb{C}^{n-1}$. Let $D_{z_0}(A) = D(A) \times \{z_0\}$ be the disc that contains $\varphi(0)$. Set $S := \varphi^{-1}\left(\varphi(B^{2n}(a)) \cap D_{z_0}(A)\right)$ is a 2-dimensional complex submanifold passing through the origin. This means that the area of S is greater or equal to a by Lelong's inequality. Since J is complex, we have that

$$\begin{aligned} a &\leq \mathrm{area}_{J_0}(S) \\ &= \mathrm{area}_{\omega_0}(S) \\ &= \int_S \omega_0 \\ &= \int_S \varphi^* \omega_0 \\ &= \int_{\varphi(S)} \omega_0 \\ &\leq \int_{D_{z_0}(A)} \omega_0 \\ &= A. \end{aligned}$$

The third equality holds because $\varphi^*\omega_0 = \omega_0$ (the embedding is symplectic), the fourth one is just the change of variable formula, and the fifth one comes from the monotonicity of the integral. This proof is slightly absurd. If φ preserves both the symplectic form and the complex structure, it preserves the metric. Here, the metric is Euclidean which means the embedding is a translation followed by a rotation, which makes the entire proof trivial. Our goal will now be to mimick this idea using the contractibility of $\mathcal{J}(M)$.
Suppose then that φ is only symplectic. Let J_φ be an almost complex structure

on $Z^{2n}(A)$ such that $J_\varphi = \varphi_* J_0$. Here, $\varphi_* J_0$ means $\mathrm{d}\varphi \circ J_0 \circ \mathrm{d}\varphi^{-1}$. We can find such a J_φ because both J_0 and $\varphi_* J_0$ are ω-tame. Indeed, a simple calculation shows that $\omega_0(\varphi_* J_0 u, u) = \omega_0\left(J_0 \circ \mathrm{d}\varphi^{-1}(\cdot)u, \mathrm{d}\varphi^{-1}(\cdot)u\right) > 0$. We're now looking for a J_φ-holomorphic disk in $Z^{2n}(A)$ passing through $\varphi(0)$ with boundary on the boundary of $Z^{2n}(A)$ with ω_0 non-negative everywhere along the disc. If such a disk exists, we can repeat the argument above. To simplify matters, we start by compactifying the disk. Take $A' > A$ and compactify $D(A)$ into $S^2(A')$. We can endow $M := S^2(A') \times \mathbb{C}^{n-1}$ with the product symplectic form $\omega = \omega_{S^2} \oplus \omega_0$ because there always exists a symplectic form on S^2. We thus need to show the following lemma:

Lemma 2.15. *The moduli space*

$$\begin{aligned}\mathcal{M}(J') := \{u : \left(S^2, i\right) \to \mathbb{C}^{n-1} \times S^2(A') \mid \overline{\partial}_{J'} u = 0,\ u(N) = \varphi(0) \text{ and} \\ \left[u(S^2)\right] = \left[\{*\} \times S^2(A')\right] / PSL(2, \mathbb{C})\}\end{aligned}$$

for $J' \in \mathcal{J}_\omega(M)$ is not empty. Here N is the North pole on $S^2(A')$.

Proof. First, notice that $\mathcal{M}(J_0) = \{u(x) = (*, x)\}$. There's only one solution because for all u J_0-holomorphic, by the maximum principle we can we can always squeeze restrict u to be constant. Choose a path from J_0 to J, with J_t, $t \in [0,1]$, $J_t \in \mathcal{J}_\omega(\mathbb{C}^{n-1} \times S^2(A')$. Let $\mathcal{M} := \coprod_{0 \le t \le 1} \{t\} \times \mathcal{M}(J_t) \subset [0,1] \times C^\infty(S^2, \mathbb{C}^{n-1} \times S^2)/$reparametrization. If we can show that $\mathcal{M}$ is compact, we are done because then $\mathcal{M}(J)$ is non-empty. Suppose instead that $\mathcal{M}^t$ becomes empty at t^*. We have a sequence $\{t_k^*\} \to t^* < 1$. For all k, we have $u_k(S^2) := S_k \subset U(J_{t_k})$. Gromov's compactness theorem tells us that, after passing to a subsequence, the spheres S_k converge in a suitable sense to a finite union of J^{t^*}-holomorphic spheres $\mathcal{S}^1, \ldots, \mathcal{S}^m$ whose homology classes $[C_i]$ add up to $[C] = \left[S^2(A') \times \{*\}\right]$. But now we're done, because each sphere is J^{t^*}-holomorphic and thus $0 < \int_{\mathcal{S}_i} \omega = [\omega]\, C_i$. This means that $m = 1$ and $n_1 = 1$, because $C_i = n_i C_i$ in $H_2(M; \mathbb{Z}) = \mathbb{Z}$ with $n_1 \ge 1$, and $\sum_i = 1^m n_i = 1$. This in turn shows that $\mathcal{M}^{t^*}$ is not empty. □

The rest of the proof is easy. Let u be such a J_φ-holomorphic disk in $\mathcal{M}(J_\varphi)$. We then have that u is J_φ-holomorphic and thus setting $S := \varphi^{-1}\left(\varphi(B^{2n}(a) \cap u(S^2)\right)$, we get $a \le A$ using the Lelong inequality. □