ALMOST COMPLEX MANIFOLDS AND PSEUDOHOLOMORPHIC CURVES

HUNTER LIU AND KEVIN WONG

Abstract. Among the first results one learns in symplectic geometry is that symplectic transformations preserve volume. It is easily possible to transform a ball of some radius into a cylinder of another radius in a volume-preserving fashion. Then are all volume-preserving transformations symplectic? Gromov's non-squeezing theorem [4] gave one of the first cases where the answer was no: the radius of the ball is at most the radius of the cylinder for a symplectic embedding. His original proof introduces pseudoholomorphic curves on almost complex manifolds and this paper serves as an introduction to Gromov's 1985 proof and the notions he introduced therein.

Contents

1. Introduction

Symplectic geometry is strongly motivated by certain systems and constructs from physics, particularly in classical mechanics and certain branches of optics. Of note are Hamilton's equations, which are closely connected to symplectic forms on symplectic manifolds. However, with its long history came several lasting open-ended questions about symplectic manifolds, for little about them were understood.

One of these questions is related to the properties of maps between symplectic manifolds, or symplectomorphisms. These were known to preserve area, but it wasn't known whether or not they preserved other invariants as well. In fact, if all symplectomorphisms preserve area, can we say anything about the converse? In 1985, Gromov's non-squeezing theorem was one of the first and most important cases where the converse is not true. It states:

Theorem 1.1 (Gromov, 1985)**.** *If there exists a symplectic embedding of a $2n$-dimensional ball of radius r into a $2n$-dimensional cylinder of radius R, then $r \leq R$.*

In essence, it says that symplectomorphisms cannot elongate or stretch spheres into eggs or tubes, and it indicates the existence of shape-related symplectic invariants. This theorem was foundational to modern symplectic geometry and sparked a plethora of new research in its wake. In fact, one idea that was motivated by Gromov's non-squeezing theorem was the symplectic capacity. A notable example of the symplectic capacity is the Hofer-Zehnder capacity, whose existence and properties are sometimes used to give alternate proofs of Gromov's non-squeezing theorem.

Whereas the Hofer-Zehnder capacity uses tools from the study of Hamiltonian functions, Gromov's original proof utilised an important insight about almost-complex structures on symplectic manifolds. We will develop the theory of almost-complex structures and discuss how it is connected to complex analysis. It was by studying pseudoholomorphic (or sometimes J-holomorphic) curves, which are analogous to holomorphic functions on Riemann surfaces, that Gromov was able to finally give a proof of the non-squeezing theorem. We shall not be giving a full proof here due to its liberal use of non-analytic concepts such as homology theory, CW-complexes, and algebraic topology, but we will attempt to flesh out and develop the role of pseudoholomorphic curves and outline how they fit in with the rest of the proof while adding commentary and heuristic outlines as much as possible.

Finally, we will assume some background in differential geometry, specifically regarding Riemannian manifolds, tangent and cotangent spaces, and differential forms. Regardless, we develop some basics in symplectic geometry and manifold theory to establish terminology and familiar concepts.

2. Background

The smoothest introduction to symplectic geometry is a concept called *linear* symplectic geometry where we concern ourselves with well-behaved (read: linear) maps on finite-dimensional vector spaces. This section is primarily from [8].

2.1. Linear Symplectic and Complex Structure.

Definition 2.1. *Suppose V is a finite-dimensional real vector space and $\omega : V \times V \to \mathbb{R}$ is a bilinear map. We say ω is:*

- *anti-symmetric if for all $u, v \in V$, $\omega(u,v) = -\omega(v,u)$.*
- *non-degenerate if the associated map $\tilde{\omega} : V \to V^*$ defined by $\tilde{\omega}(u)(v) = \omega(u,v)$ is bijective, where V^* is the dual of V.*

Observe that we can also regard ω has a linear 2-form on V. Also, the above non-degeneracy condition is equivalent to the condition that if $\omega(u,v) = 0$ for all $v \in V$, then $u = 0$.

Definition 2.2. *A symplectic vector space is a pair (V, ω) where V is a real vector space and ω a non-degenerate, antisymmetric bilinear map. We say ω is a linear symplectic structure or form on V.*

A standard example of this is to take $V = \mathbb{R}^{2n}$ and define

$$\omega_{\text{std}}((x,u),(y,v)) = \langle x, v \rangle - \langle u, y \rangle$$

where $\langle \cdot, \cdot \rangle$ is the standard real inner product. Then (V, ω_0) is easily verified to by a symplectic vector space. However, these maps are not the only structures we can place on a space.

Definition 2.3 (Complex Structure)**.** *A complex structure on a vector space V is an automorphism $J : V \to V$ such that $J^2 = -\mathrm{Id}$. We call the pair (V, J) a complex vector space.*

Roughly speaking, this definition of J allows us to "multiply by $\sqrt{-1}$" on V. In fact, a basic example of this is $\mathbb{C}^n$ with the standard complex structure $J_0(z) := iz$. It is easy to check this defines a complex structure and moreover, if we identify $\mathbb{C}^n$ with $\mathbb{R}^{2n}$ in the usual way, we see that this is equivalently $J_0(x_i) = y_0$ and $J_0(y_i) = -x_i$ where $z = (x, y)$.

With our two structures we can place on a vector space, let us now see how these can interact with each other.

Definition 2.4. *Let (V, ω) be a symplectic vector space and J a complex structure on V. We say:*

- *J is tamed by ω if the quadratic form $\omega(v, Jv)$ is positive definite.*
- *J is compatible with ω if it is tamed by ω and J satsifies $\omega(Jv, Jw) = \omega(v, w)$.*

Let $\mathcal{J}(V,\omega)$ denote the space of ω-compatible complex structures on V. With these two guidelines, the immediate question is if $\mathcal{J}(V,\omega)$ is empty or not. In fact, we have that for any symplectic vector space, it will admit a compatible complex structure.

Theorem 2.5. *Every symplectic vector space (V,ω) admits a compatible complex structure. Moreover for all inner products $g(\cdot,\cdot)$ on V, one can canonically construct such a J.*

Proof. Consider an inner product g on V. Since g,ω are non-degenerate, there exists an endomorphism A of V such that $\omega(u,v)=g(Au,v)$. In other words, A is the transpose matrix of ω in an orthogonal basis. The map A is skew adjoint with respect to G since

$$g(A^*u,v)=g(u,Av)=g(Av,u)=\omega(v,u)=-\omega(u,v)=g(-Au,v)$$

so $A^*=-A$. Note $(AA^*)^*=A^*A=(-A)(-A^*)=AA^*$ so we have symmetric and positive-definiteness from

$$g(AA^*u,u)=g(A^*u,A^*u)>0$$

for $u\neq 0$. Symmetry gives us diagonalizability and positive-definite gives that all the eigenvalues λ_i are positive so we can write

$$AA^*=B\text{diag}(\lambda_1,\dots,\lambda_n)B^{-1}$$

so we can define $\sqrt{AA^*}=B\text{diag}(\sqrt{\lambda_1},\dots,\sqrt{\lambda_n})B^{-1}$. This is still symmetric and positive-definite. Let

$$J:=(\sqrt{AA^*})^{-1}A.$$

Since A commutes with $\sqrt{AA^*}=\sqrt{A(-A)}=\sqrt{-A^2}$, we get that J and A commute. Since $A^*=-A$, $J^*=-J$ so

$$-J^2=J^*J=\left(A^*\left(AA^*\right)^{-1}\right)\left(\left(AA^*\right)^{-1}A\right)=A^*(AA^*)^{-1}A=\text{Id}.$$

Thus, J defines a complex structure. It remains to show compatibility, but note

$$\omega(Ju,Jv)=g(AJu,Jv)=g(JAu,Jv)=g(J^*JAu,v)=g(Au,v)=\omega(u,v)$$

and

$$\omega(u,Ju)=g(Au,Ju)=g(J^*Au,u)=g(\sqrt{AA^*}u,u)>0$$

for $u\neq 0$ by positive-definiteness of $\sqrt{AA^*}$. □

In general, the positive-definite inner product $\omega(u,Jv)\neq g(u,v)$, but if J is given and $g(u,v)=\omega(u,Jv)$, then we have equality. We can also consider families of complex structures and symplectic vector spaces. If (V_t,ω_t) is a smooth family of symplectic vector spaces, then there exists a smooth family of inner products g_t to get a smooth family of complex structures J_t.

Theorem 2.6. *The space $\mathcal{J}(V,\omega)$ is contractible.*

Proof. Fix an ω-compatible complex structure J on V. Define the contraction map $f : [0,1] \times \mathcal{J}(V,\omega) \to \mathcal{J}(V,\omega)$ as follows. For all $J' \in \mathcal{J}(V,\omega)$, there exists a naturally-defined inner product g'. Let $g_t = tg + (1-t)g'$ so g_t is an inner product on V, which gives a canonically-defined continuous family of complex structures J_t. Thus $J_0 = J'$ and $J_1 = J$, so f is continuous with $f(0,J') = J'$ and $f(1,J') = J$. □

In linear theory, we observe that everything is well-behaved, but fortunately, many of the concepts and theorems discussed generalize nicely to manifolds. For example, as we will see, even if (M,ω) is a symplectic manifold, rather than a vector space, it must still be even-dimensional.

2.2. **Manifolds.** Although we assume some familiarity with manifolds, for the sake of notation, we re-introduce some concepts.

Definition 2.7 (Smooth Manifold)**.** *Let $n \in \mathbb{N}$ and M a set. A chart on M is a pair (ϕ, U) where $U \subseteq M$ and $\phi : U \to \phi(U)$ is bijection to some open subset $\phi(U)$ of $\mathbb{R}^n$. Two charts (ϕ_1, U_1) and (ϕ_2, U_2) are compatible if $\phi(U_1 \cap U_2)$ and $\phi_2(U_1 \cap U_2)$ are open and the transition map*

$$\phi_{21} = \phi_2 \circ \phi_1^{-1} : \phi_1(U_1 \cap U_2) \to \phi_2(U_1 \cap U_2)$$

is a diffeomorphism. A smooth atlas on M is a collection $\mathscr{A}$ of charts on M any two of which are compatible and such that the sets U, as (ϕ, U) range over $\mathscr{A}$, cover M. A maximal smooth atlas is an atlas containing every chart compatible with each of its members. A smooth n-manifold is a pair consisting of a set M and a maximal smooth atlas $\mathscr{A}$ on M.

Manifolds, although tedious to define rigorously, are ubiquitous in mathematics as we trust the reader may have encountered.

Definition 2.8 (Diffeomorphisms)**.** *Let $(M, \{(\phi_\alpha, U_\alpha)_{\alpha\in A}\})$ and $(N, \{(\psi_\beta, V_\beta)_{\beta\in B}\})$ be smooth manifolds. A map $f : M \to N$ is smooth if f is continuous and the map*

$$f_{\beta\alpha} = \psi_\beta \circ f \circ \phi_\alpha^{-1} : \phi_\alpha(U_\alpha \cap f^{-1}(V_\beta)) \to \phi_\beta(V_\beta)$$

is smooth for all $\alpha \in A, \beta \in B$. We say f is a diffeomorphism if it is bijective and both f, f^{-1} are smooth. The manifolds M, N are said to be diffeomorphic if there exists such an f.

In symplectic geometry, one is often concerned with doing calculus on these manifolds. Indeed, in just these first two definitions we already see familiar notions of smoothness and continuity. But these apply to *functions*, so what if we want to apply derivatives to get something like a space, a tangent space, if you will. To define tangent spaces, there are two natural motivations. The first way is to work on points on the manifold and to consider all the curves through a certain point. But some curves will have the same velocity when running through the point, so we need to impose some equivalencies. This gives rise to the first definition below. The second idea is to work in local coordinate patches where we can apply the familiar toolkit of Euclidean spaces. Recall that for $U \subseteq \mathbb{R}^n$, $V \subseteq \mathbb{R}^m$ open subsets, $f : U \to V$, and $x \in U$, the derivative of f at x is the linear map $df(x) : \mathbb{R}^n \to \mathbb{R}^m$ defined by

$$df(x) := \frac{d}{dt} f(x + th)\bigg|_{t=0} = \lim_{t \to 0} \frac{f(x+th) - f(x)}{t}$$

where $h \in \mathbb{R}^n$.

Definition 2.9 (Tangent Space)**.** *Let M be a manifold with atlas $\mathscr{A} = \{(\phi_\alpha, U_\alpha) : \alpha \in A\}$ and $p \in M$. Two smooth curves $\gamma_0, \gamma_1 : [0,1] \to M$ with $\gamma_0(0) = \gamma_1(0) = p$ are called p-equivalent if for some $\alpha \in A$ with $p \in U_\alpha$, we have*

$$\frac{d}{dt} \phi_\alpha(\gamma_0(t))\bigg|_{t=0} = \frac{d}{dt} \phi_\alpha(\gamma_1(t))\bigg|_{t=0}$$

which we denote $\gamma_0 \sim_p \gamma_1$. This defines an equivalence class so we can denote the class of a smooth curve $\gamma : [0,1] \to M$ with $\gamma(0) = p$ to be $[\gamma]_p$. Every equivalence class is a tangent vector of M at p. The tangent space of M at p is the set of equivalence classes

$$T_pM := \{[\gamma]_p : \gamma \in C^\infty([0,1], M)\}$$

with $\gamma(0) = p$.

Definition 2.10 (Tangent Space, Alternative Definition)**.** *Let M be a manifold with atlas $\mathscr{A} = \{(\phi_\alpha, U_\alpha) : \alpha \in A\}$ and $p \in M$. The $\mathscr{A}$-tangent space of M at p is the quotient space*

$$T_p^{\mathscr{A}} M := \bigcup_{\alpha : p \in U_\alpha} \{\alpha\} \times \mathbb{R}^n \Big/ \sim_p$$

where $(\alpha, v) \sim_p (\beta, w)$ if and only if $d(\phi_\beta \circ \phi_\alpha^{-1})(\phi_\alpha(p))v = w$.

In Defintion 2.10, $T_p^{\mathscr{A}} M$ is a vector space with dimension n but it is not clear immediately that T_pM is as defined in Definition 2.9. But there exsts a bijection $\varphi : T_pM \to T_p^{\mathscr{A}} M$ where $[\gamma]_p$ is mapped to the pair

$$\left(\alpha, \frac{d}{dt} \phi_\alpha(\gamma(t))\bigg|_{t=0} \right).$$

This induces a vector space structure on T_pM. Thus, these definitions are in fact interchangeable and we can drop the $\mathscr{A}$ superscript from Definition 2.10.

For each p, we can associate a tangent space T_pM, but we can also consider the manifold and all the tangent spaces to get the tangent bundle

$$TM := \bigcup_{p \in M} T_pM$$

which we will see becomes important in later sections.

3. Symplectic Manifolds

Definition 3.1 (Symplectic Manifold)**.** *A symplectic manifold is a pair (M, ω), where M is a smooth manifold and $\omega : TM \times TM \to \mathbb{R}$ is a non-degenerate closed 2-form, i.e. $d\omega = 0$. ω is called a symplectic form.*

Riemannian manifolds are equipped with Riemannian metrics, which endow the manifold with a notion of lengths and angles. Symplectic forms don't have an intuitive or natural geometric interpretation. One can think of them as an analogue to a notion of 2-dimensional area, though this should be taken with a grain of salt.

In Riemannian geometry, diffeomorphisms of Riemannian manifolds are no longer the only transformations of interest; rather, one studies isometries, which preserve the local notions of length. Similarly, maps between symplectic manifolds need to respect the symplectic forms of each manifold.

Definition 3.2. *Let (M, ω) and (M', ω') be two symplectic manifolds. A map $\varphi : M \to M'$ is a symplectomorphism if it is a diffeomorphism and $\omega = \varphi^*\omega'$, where φ^* is the pullback* [1] *by φ. $\iota : M \hookrightarrow M'$ is a symplectic embedding if it's a smooth embedding and $\omega = \iota^*\omega'$.*

A more important facet of symplectic manifolds is that not all manifolds admit a symplectic form. In contrast, one can always find a Riemannian metric on an arbitrary manifold M by embedding it in a high-enough-dimensional Euclidean space and restricting the natural dot product to M's tangent spaces.

Proposition 3.3. *Let (M, ω) be a symplectic manifold. Then, M is even-dimensional.*

Proof. Fix $p \in M$, and consider $\omega_p : T_pM \times T_pM \to \mathbb{R}$. Since ω is a non-degenerate 2-form, ω_p must be skew-symmetric and bilinear with nonzero kernel. We claim this forces T_pM, which is a finite-dimensional vector space, to be even-dimensional. This would in turn force M to be even-dimensional as well.

We will use a Gram-Schmidt-eqsue argument. Take any $v_1 \in T_pM$, and take some $w_1 \in T_pM$ such that $\omega_p(v_1, w_1) \neq 0$. Such a w_1 must exist because ω is non-degenerate, and $w_1 \notin \text{span}\{v_1\}$ because ω_p is skew-symmetric. By rescaling appropriately, we may assume $\omega_p(v_1, w_1) = 1$ without loss of generality.

Let $V_1 = \text{span}\{v_1, w_1\}$, and let $W_1 = \{v \mid \omega_p(v, w) = 0 \quad \forall w \in V_1\}$. By using bilinearity, we see that for any $a, b \in \mathbb{R}$,

$$\omega_p(av_1 + bw_1, v_1) = -b$$

whereas

$$\omega_p(av_1 + bw_1, w_1) = a.$$

Hence, if $av_1 + bw_1 \in W_1$, $a = b = 0$. In other words, $V_1 \cap W_1 = \{0\}$. Finally, we can apply a Gram-Schmidt like argument to "project" any vector $v \in T_pM$ down onto its V_1 and W_1 components. Fix any such v, and let $a = \omega_p(v, w_1)$ and $b = -\omega_p(v, v_1)$.

[1] Recall that the pullback of a symplectic form is $(\varphi^*\omega')(v, w) = \omega'(\varphi(v), \varphi(w))$.

Then, a brief computation shows that $v - av_1 - bw_1 \in W_1$, and so we have decomposed v into

$$v = (av_1 + bw_1) + (v - av_1 - bw_1).$$

Since $av_1 + bw_1 \in V_1$, we conclude that $T_pM = V_1 \oplus W_1$.

We may repeat this argument on W_1 to generate new subspaces V_2 and W_2, with $T_pM = V_1 \oplus V_2 \oplus W_2$, and then repeat this on W_2, and so on and so forth until this process terminates. Importantly, $V_1, V_2, \ldots$ are all 2-dimensional, so the dimension of the W_i's is decreasing. As T_pM is finite-dimensional, this process must eventually terminate, resulting in

$$T_pM = V_1 \oplus \cdots \oplus V_n.$$

We conclude that $\dim T_pM = \dim M = 2n$. □

This is actually a consequence of a more general result regarding skew-symmetric bilinear forms on finite-dimensional vector spaces, which gives a similar decomposition while accounting for a possible kernel. Again, this result highlights a sort of rigidity of symplectic manifolds in the sense that not every manifold can admit a symplectic form.

Importantly, we can use the above result to pick a "standard" basis for T_pM based on ω_p. In particular, we can set the basis $v_1, \ldots, v_n$, $w_1, \ldots, w_n$ such that $\omega_p(v_i, w_i) = 1$ for all $1 \leq i \leq n$ and ω_p is zero for all other pairs of basis vectors. In fact, we can produce our first example of a symplectic manifold using this idea. Simply take $\mathbb{R}^{2n}$ with the standard basis $x_1, \ldots, x_n$, $y_1, \ldots, y_n$ (labelled slightly differently for notational ease) and the symplectic form

$$\omega_{\text{std}} = \sum_{i=1}^{n} dx_i \wedge dy_i.$$

Something unexpected about symplectic manifolds is that they all locally "resemble" each other. For manifolds in general, this is far from true: no open neighbourhood on S^2 is isometric to $\mathbb{R}^2$ because their curvatures are different. In other words, Riemannian manifolds have a rich variety of *local* structures, and this is manifested in ideas like Gaussian and mean curvature. No such local invariants can exist on symplectic manifolds, and that is thanks to the following result:

Theorem 3.4 (Darboux's Theorem)**.** *Let (M, ω) be a symplectic manifold. For every point $p \in M$, there exist local coordinates $x_1, \ldots, x_n$, $y_1, \ldots, y_n$ such that*

$$\omega_p = \sum_{i=1}^{n} dx_i \wedge dy_i.$$

To reiterate, all symplectic manifolds *locally* look the same: you can always identify them with neighbourhood in $(\mathbb{R}^{2n}, \omega_{\text{std}})$. However, they do have different *global* properties, and one of them is the symplectic volume. If ω is a symplectic form on a $2n$-dimensional manifold M, then ω^n is a volume form (a differential form of maximal degree). Since ω is non-degenerate, it's clear that ω^n is nonzero.

In particular, (M, ω) and (M', ω') be symplectomorphic to each other. Then, their "symplectic volumes" are equal, i.e.

$$\int_M \omega^n = \int_{M'} (\omega')^n .$$

This follows directly from the definition of a symplectomorphism. Notably, when $M \subseteq \mathbb{R}^{2n}$ with the symplectic form ω_{std} (with the appropriate restrictions), one can check that ω_{std}^n is a scalar multiple of the standard volume form for $\mathbb{R}^n$.

There is a bit more subtlety when ω and ω' aren't *globally* induced by ω_{std}, in which case there is no obvious relationship between the above volumes and our usual Euclidean volumes. Nevertheless, the statement is still true:

Theorem 3.5. *Let M and M' be symplectic manifolds, and suppose there exists an embedding $M \hookrightarrow M'$. Then,* $\operatorname{vol} M \leq \operatorname{vol} M'$.

Here, vol refers to the usual sense of Riemannian volume, i.e. the integral of the Riemannian volume form. This is a corollary of Liouville's theorem, which is an important theorem in the study of statistical and Hamiltonian mechanics.

This result also gives us a more formal motivation for Gromov's non-squeezing theorem. This monotonicity of volumes is necessary for symplectically embeddable manifolds, but it's not obvious whether or not it is sufficient. As it so turns out, it isn't. Even if M has a smaller volume than M', there may not be any symplectic embeddings $M \hookrightarrow M'$, and this loosely boils down to factors such as the "shape" or "width" of M and M'.

4. Almost-Complex Structures

We will now take a slight detour and look at almost-complex structures. Similar to symplectic forms and manifolds, almost-complex structures have some motivation from physics. Specifically, certain Hamiltonian systems induce systems of differential equations that strongly resemble the Cauchy-Riemann equations. For more detailed discussion on the physics and mechanics buttressing symplectic geometry, we direct the reader to [5].

Definition 4.1. *Let M be a manifold, $J : TM \to TM$. J is an almost-complex structure it varies smoothly over M and if for all points $p \in M$, $J_p : T_pM \to T_pM$ is a linear automorphism such that $J_p^2 = -\mathrm{Id}$. The pair (M, J) is called an almost-complex manifold.*

Almost-complex structures are generalisations of multiplication by i, and it's easy to see that these ideas coincide for complex manifolds.

Much like symplectic forms, not every manifold can admit an almost-complex structure, and we get a similar result regarding the dimensions of such manifolds.

Proposition 4.2. *Let (M, J) be an almost-complex manifold. Then, M is even-dimensional (with respect to $\mathbb{R}$).*

Proof. Treating M as a real manifold again, fix any $p \in M$ and consider $J_p : T_pM \to T_pM$. Let $n = \dim T_pM$. As J_p is an $\mathbb{R}$-linear map, $\det J_p$ is real. However, since $J_p^2 = -\mathrm{Id}$, we have that $\det\left(J_p^2\right) = \left(\det J_p\right)^2 = (-1)^n$. It follows that n must be even. □

Notably, complex manifolds have two choices of an almost-complex structure, namely multiplication by either i or by $-i$. One interpretation of this is that $\pm i$ represent two different "orientations" of complex manifolds. Generally speaking, for any almost-complex structure J, $-J$ is also an almost-complex structure, and this demonstrates a sort of handedness present in the construct. Symplectic forms also exhibit an inherent notion of orientation because of their skew-symmetry, and so it makes sense to ask about what happens when these two senses of orientation are compatible with each other.

Definition 4.3. *Let (M, ω) be a symplectic manifold, and let J be an almost-complex structure on M. J is ω-tame or tamed by ω if $\omega(v, Jv)$ is positive-definite for all $v \in TM$. J is ω-compatible if it is ω-tame and*

$$\omega(Jv, Jw) = \omega(v, w)$$

for all $v, w \in TM$.

Notably, the condition that $\omega\,(Jv, Jw) = \omega(v, w)$ is equivalent to the condition that $(v, w) \mapsto \omega(Jv, w)$ is a Riemannian metric, i.e. that it is positive-definite and symmetric.

Returning to some results established regarding symplectic and almost-complex structures on vector spaces, we can obtain a crucially important result about $\mathcal{J}(M, \omega)$, the space of ω-compatible almost-complex structures on M, as well as about $\mathcal{J}_\tau(M, \omega)$, the space of ω-tame almost-complex structures on M.

Theorem 4.4. *$\mathcal{J}(M, \omega)$ and $\mathcal{J}_\tau(M, \omega)$ are both nonempty and contractible with respect to the C^∞ topology (cf. Theorem 2.6).*

A crucial and important corollary is that $\mathcal{J}(M, \omega)$ is connected. The proof of this theorem is much more involved than the same contractibility statement for symplectic vector spaces, and we direct the reader to [5] for a complete proof.

The nonemptiness of these spaces is the same thing as saying that every symplectic manifold admits an almost-complex structure, and this is critically important in applying the soon-to-be-developed theory of pseudoholomorphic curves to symplectic geometry. It turns out that manifolds equipped with almost-complex structures always admit skew-symmetric non-degenerate 2-forms, though they may not be closed. This is called an almost-symplectic form, and the resulting structure is called an almost-Kähler manifold. One cannot even guarantee that a complex structure on a manifold induces a symplectic form, and an example of this is the Hopf surface $S^1 \times S^3$. [2] gives a fuller discussion of examples and counterexamples in chapter 17.3.

5. Pseudoholomorphic Curves

Recall the notion of a holomorphic curve $f : U \to \mathbb{C}$ where $U \subseteq \mathbb{C}$ open can be characterized by many ways: being analytic, satisfying the Cauchy-Riemann equations, and so forth, and by the end, we find that many of these notions are equivalent. If we wanted to consider multi-valued holomorphic curves, then we can take $U \subseteq \mathbb{C}^m$ open, a smooth map $f : U \to \mathbb{C}^n$, and say that f is holomorphic if its partial derivatives $\partial_j f$ exist for all $j = 1, \dots, m$. That is, if $h \in \mathbb{C}^m$, we have the existence of

$$\lim_{h \to 0} \frac{f(z_1, \dots, z_j + h, \dots, z_m) - f(z_1, \dots, z_m)}{h}$$

for all j. The verification of the existence of this theorem is tedious and often infeasible. Fortunately, the structure and rigidity of complex numbers allow us to reduce the condition for complex differentiability to the familiar Cauchy-Riemann equations.

To see this, let us identify $\mathbb{C}^n$ with $\mathbb{R}^{2n}$ by viewing $(z_1, \dots, z_n) \in \mathbb{C}^n$ as the real vector $(x_1, y_1, \dots, x_n, y_n) \in \mathbb{R}^{2n}$ where $z_j = x_j + iy_j$. Then for all $z \in U \subseteq \mathbb{C}^m$, we get a differential $df(z) : \mathbb{C}^m \to \mathbb{C}^n$. This is exactly a real linear map from $\mathbb{R}^{2m}$ to $\mathbb{R}^{2n}$ as we saw in Section 2.2. For all $\lambda \in \mathbb{C}$, scalar multiplication $z \mapsto \lambda z$ can be treated as a real linear map from $\mathbb{R}^{2n}$ to itself. It turns out that f being holomorphic is equivalent to its differential $df(z)$ being complex linear; that is,

$$df(z)(\lambda v) = \lambda df(z)v$$

for all $v \in \mathbb{C}^m$ and $\lambda \in \mathbb{C}$. As we already have real linearity, it remains to satisfy

$$df(z) \circ (i \cdot) = i df(z)(\cdot) \tag{1}$$

as linear maps on $\mathbb{R}^{2n}$ or $\mathbb{R}^{2m}$. If $f : U \to \mathbb{C}^n$ is smooth, we can write $f(z_1, \dots, z_n) = u(x_1, \dots, x_n) + iv(y_1, \dots, y_n)$ with $z_j = x_j + iy_j$ and the Cauchy Riemann equations generalize nicely to

$$\frac{\partial u}{\partial x_j} = \frac{\partial v}{\partial y_j} \quad \text{and} \quad \frac{\partial u}{\partial y_j} = -\frac{\partial v}{\partial x_j}.$$

In fact, these are equivalent to satisfying (1) and complex linearity.

This strongly resembles the almost-complex structures that we have shown earlier. In particular, mapping $\frac{\partial}{\partial x_n} \mapsto \frac{\partial}{\partial y_n}$ and $\frac{\partial}{\partial y_n} \mapsto -\frac{\partial}{\partial x_n}$ is a skew-symmetric nondegenerate closed linear operation on TU, and it squares to -1!

Recall that if (M, J) is an almost complex manifold, we use the terminology *almost* because complex manifolds always admit a complex structure (via multiplication by i on TM), but the converse does not necessarily hold. More precisely, if M is instead a complex manifold, then any choice of holomorphic coordinates on $U \subseteq M$ will identify the tangent spaces T_pU with $\mathbb{C}^n$ and we can assign the standard complex structure to each T_pU. Moreover, this assignment is independent on the coordinates so M has a natural almost complex structure that looks like the standard one in any chart. We call such an almost complex structure integrable.

As a result from the preceeding discussion, a Riemann surface can be viewed as a complex manifold of complex dimension 1 and thus admits an integrable complex structure.

Definition 5.1 (Pseudoholomorphic)**.** *Suppose (Σ, j) is a Riemann surface, and (M, J) is an almost complex manifold. A smooth map $u : \Sigma \to M$ is pseudoholomorphic or J-holomorphic if its differential is complex linear at every point. That is,*

$$du(z) \circ j = J \circ du(z). \tag{2}$$

Observe that (2) is a non-linear first order PDE called the non-linear Cauchy Riemann equation. One can rewrite this using holomorphic local coordinate patches on Σ to turn it into a more familiar form as follows. Let $x + iy$ by a local holomorphic chart on Σ so, supposing j is integrable, $j\partial_x = \partial_y$ and $j\partial_y = -\partial_x$. Then locally, (2) is saying

$$\partial_x u + J(u)\partial_y u = 0.$$

The standard Cauchy-Riemann equations can be written $\partial_x u + i\partial_y u = 0$ so we can view the above as a perturbation of the standard equation. Indeed, we can choose coordinates near $p \in M$ so $J(p)$ can be identified with the standard complex structure and we truly have the above as a small local perturbation.

It is well-known that holomorphic functions between Riemann surfaces are very rigid: well-behaved singularities can be removed or filled in, they have unique continuations, and they have harmonic real components, to name a few properties. When f is a pseudoholorphic curve between two truly complex structures on complex manifolds, the above definition coincides with the usual notion of holomorphicity, and f of course retains these nice, rigid properties. However, even when the almost-complex manifolds do not admit actual complex structures, f continues to enjoy powerful properties of rigidity.

Theorem 5.2. *Suppose (M, J) is a smooth almost complex manifold. Then*

- *(Regularity) Every map $u \in C^1(\Sigma, M)$ solving (1) is smooth.*
- *(Local Existence) For all $p \in M$ and $v \in T_pM$, there exists a neighborhood $U \subseteq \mathbb{C}$ of the origin and a pseudoholomorphic curve $u : U \to M$ such that $u(0) = p$ and $\partial_x u = v$ in the standard coordinates $x + iy \in U$.*
- *(Intersections) Suppose $u_1 : \Sigma_1 \to M$ and $u_2 : \Sigma_2 \to M$ are two non-constant pseudoholomorphic functions with intersection $u_1(z_1) = u_2(z_2)$. Then there exist neighborhoods $U_1 \subseteq \Sigma_1$ and $U_2 \subseteq \Sigma_2$ of z_1 and z_2 respectively such that the images $u_1(U_1 \setminus \{z_1\})$ and $u_2(U_2 \setminus \{z_2\})$ are either identical or disjoint.*

For complete proofs, we direct the reader to [6] and [9]. The last property on intersections is similar to the familiar property of unique analytic continuation, and it can be rephrased as saying that the set of points where u_1 and u_2 coincide cannot accumulate anywhere unless u_1 and u_2 are identical.

Another important regularity theorem is somewhat reminiscent of how *univalent* functions can only converge to other univalent functions or to constant functions.

Proposition 5.3. *Suppose J_k is a sequence of almost-complex structures on a manifold M that converges to another almost-complex structure J on M. Let f_k be a sequence of non-constant pseudoholomorphic curves on M with respect to J_k that converges in C^∞ to f, a pseudoholomorphic curve on M with respect to J. Then, f is non-constant.*

5.1. **Energy.** Another important quantity of a pseudoholomorphic curve is its *energy*. This is closely related to the 2-dimensional notion of area of a curve. This has connections with tameness and compatibility as discussed in Section 4.3.

For holomorphic curves in $\mathbb{C}^n$, the area they trace can be found by integrating the standard symplectic structure. To generalize this, suppose (M, ω) is a symplectic manifold, $J \in \mathcal{J}_\tau(TM, \omega)$ and g_J is the bundle metric

$$g_J(v, w) = \frac{1}{2}\left[\omega(v, Jw) + \omega(w, Jv)\right].$$

If $u : (\Sigma, j) \to (M, J)$ is a pseudoholomorphic curve and we choose holomorphic coordinates (x, y) on a subset of Σ, then $\partial_y u = J\partial_x u$ implies the vectors $\partial_x u$ and $\partial_y u$ are orthogonal and have the same length under g_J. Thus, the area spanned by them is

$$\|\partial_x u\|_{g_J} \cdot \|\partial_y u\|_{g_J} = \|\partial_x u\|_{g_J}^2 = \omega(\partial_x u, J\partial_x u) = \omega(\partial_x u, \partial_y u)$$

so the area under g_J is

$$\mathrm{Area}_{g_J}(u) = \int_\Sigma u^*\omega.$$

This motivates the following definition of energy.

Definition 5.4 (Energy)**.** *Suppose (M, ω) is a symplectic manifold and $J \in \mathcal{J}_\tau(TM, \omega)$ is a tame almost-complex structure. The (harmonic) energy of a pseudoholomorphic curve $u : (\Sigma, j) \to (M, J)$ is the quantity*

$$E(u) := \int_\Sigma u^*\omega.$$

Observe that $E(u) \geq 0$ with equality holding if and only if u is constant when we restrict u to each connected component of Σ. These distinguish pseudoholomorphic curves in symplectic manifolds from those in general almost-complex ones. In the former case, the energy is a topological invariant that depends only on a characteristic of the curve (namely, its homology class) and so we can obtain a universal bound. But in the case of almost-complex manifolds, we have no *a priori* bound, but we can still establish them as we go.

For notation, let us introduce a Cauchy-Riemann operator. If we take coordinates $z = x + iy$ on Σ, which we also endow with a volume form vol_Σ, the 1-form, $\overline{\partial_J}(u)$ is given by

$$\overline{\partial_J}(u) = \frac{1}{2}\left(\partial_x u_\alpha + J(u_\alpha)\partial_y u_\alpha\right) dx + \frac{1}{2}\left(\partial_y u_\alpha - J(u_\alpha)\partial_x u_\alpha\right) dy.$$

Now we may introduce an important identity for the energy of a map.

Lemma 5.5. *Let (M, ω) be a symplectic manifold. If J is tamed by ω, then for all pseudoholomorphic $u : \Sigma \to M$,*

$$E(u) = \frac{1}{2}\int_\Sigma \|du\|_J^2 \, d\mathrm{vol}_\Sigma$$

where the norm of the linear map $du(z) : T_z\Sigma \to T_{u(z)}\Sigma$ is

$$\|du\|_J^2 = \frac{1}{|w|}\sqrt{\|du(z)w\|_J^2 + \|du(z)(jw)\|_J^2}$$

for $0 \neq w \in T_zM$. Moreover, if J is compatible with ω, then every smooth map $u : \Sigma \to M$ satisfies

$$E(u) = \int_\Sigma \left\|\overline{\partial_J}(u)\right\|_J^2 \, d\mathrm{vol}_\Sigma + \int_\Sigma u^*\omega.$$

Proof. Pick local coordinates $z = x + iy$ in a neighborhood of $p \in \Sigma$, which we can assume without loss of generality is an open subset of $\mathbb{C}$. By rescaling $z \mapsto \lambda z$ with $0 \neq \lambda \in \mathbb{C}$, we can also tale $d\mathrm{vol}_\Sigma = dx \wedge dy$. Thus,

$$\begin{aligned}\frac{1}{2}\|du\|_J^2 \, d\mathrm{vol}_\Sigma &= \frac{1}{2}\left(\|\partial_x u\|_J^2 + \|\partial_y u\|_J^2\right) dx \wedge dy \\ &= \frac{1}{2}\|\partial_x u + J\partial_y u\|_J^2 \, dx \wedge dy - \langle \partial_x u, J\partial_y u\rangle_J \, dx \wedge dy \\ &= \left\|\overline{\partial_J}(u)\right\|_J^2 \, d\mathrm{vol}_\Sigma + \frac{1}{2}\left(\omega(\partial_x u, \partial_y u) + \omega(J\partial_x u, J\partial_y u)\right) dx \wedge dy.\end{aligned}$$

If J is tamed by ω and $\partial_x u + J\partial_y u = 0$, then the first term on the right vanishes and the remaining term is equal to $u^*\omega$. If J is compatible with ω, then the last term on the right equals $u^*\omega$ and we have the claims. □

Further, pseudoholomorphic curves have a monotonicity proeprty. Let us specialize to the case when u maps into a Euclidean ball and is a nonconstant holomorphic map which is proper; i.e., u^{-1} takes compact sets to compact sets. Denote $B^{2n}(r_0)$ to be the Euclidean ball of radius r_0 in $\mathbb{R}^{2n}$.

Theorem 5.6 (Monotonicity)**.** *Let $r_0 > 0$, (Σ, j) be a Riemann surface, and $u : (\Sigma, j) \to (B^{2n}(r_0), i)$ be a nonconstant proper holomorphic map. Then for all $0 < r < r_0$,*

$$\pi r^2 \leq \int_{u^{-1}\left(\overline{B^{2n}(r_0)}\right)} u^*\omega_{std}.$$

The most important consequence of introducing the energy of these pseudoholomorphic maps is that area is closely related with our notion of energy and pseudoholomorphic curves can minimize these quantities (in their homology class). This is a special case of general results from the theory of minimal surfaces which we cannot currently develop here. One key fact in the proof of Gromov's non-squeezing theorem relies on the following observation using minimal surfaces:

Theorem 5.7. *Suppose (M, ω) is a closed symplectic manifold of dimenension $2n - 2 \geq 2$ with homotopy group $\pi_2(M) = 0$, σ an area form on S^2, and there exists a symplectic embedding*

$$\iota : \left(B^{2n}(r), \omega_{std}\right) \hookrightarrow (S^2 \times M, \sigma \oplus \omega_{std})$$

Then $\pi r^2 \leq \int_{S^2} \sigma$.

Theorem 5.7 can be proven as a corollary of Theorem 5.6 but also requires some results from homotopy and homology theory which we refer to [9] for.

We will also appeal to an analogue of Riemann's theorem on removable singularities of holomorphic functions.

Theorem 5.8. *Let $U \subset \mathbb{C}$ be open, let (M, ω) be an almost-complex manifold, J an ω-tame almost-complex structure, $a \in U$, and let $f : U \setminus \{a\} \to M$ be a pseudoholomorphic curve (with the standard complex structure on U). If f has finite energy and $\operatorname{im} f \subseteq M$ is compact, then f extends pseudoholomorphically to $\tilde{f} : U \to M$.*

Proof. For this, we will follow [6] and prove a continuous extension. Suppose without loss of generality that $U = \mathbb{D}$, the unit disc, and that $a = 0$. Let γ_r be the closed loop parameterised by $\gamma_r(t) := f\left(re^{it}\right)$ for $t \in [0, 2\pi]$, $r \in (0, 1)$. Since the energy is finite, we can write

$$E(f) = \int_0^1 \int_0^{2\pi} \frac{|\gamma_r'(t)|^2}{r^2} r dr \wedge d\theta.$$

Here, the length of γ_r' is given by the Riemannian metric induced by the ω-tame almost-complex structure. Using the Cauchy-Schwarz inequality, we can bound this by

$$E(f) \geq \int_0^1 \frac{|\gamma_r|^2}{2\pi r} dr,$$

where $|\gamma_r|$ is the arc length of γ_r. Since $E(f)$ is finite, we see that there must be a sequence of radii r_n decreasing to zero such that $|\gamma_{r_n}| \to 0$ as well.

By compactness of the image of f, there must be a subsequence of radii r_m such that the γ_{r_m} uniformly converges to a single point $p \in M$.

Now suppose towards a contradiction that for some sequence z_k, $f(z_k) \to q \neq p$. Let $0 < \delta \ll d(p, q)$ small. Then, for all $0 < r < r_0$ sufficiently small, there γ_r intersects both $B(p, \delta)$ and $B(q, \delta)$, giving us that $|\gamma_r|^2 \geq c\delta^2$ for some fixed positive constant c. Using the above integral estimate, this then gives us that

$$E(f) \geq \int_0^1 \frac{|\gamma_r|^2}{2\pi r} dr \geq \int_0^{r_0} \frac{c\delta^2}{2\pi r} dr = \infty,$$

which contradicts f having finite energy. Therefore, f extends continuously to U. □

We'll remark here that this proof subtly relies on the fact that pseudoholomorphic curves minimise areas. This is present in our argument that the length of γ_r for sufficiently small r is uniformly bounded from below; non-minimal surfaces may not have this property. The continuous extension is used to show that f extends fully to a pseudoholomorphic curve on U, though this is significantly more involved and relies on machinery such as the elliptic regularity theorem. For the full proof, we refer the reader to [1].

Finally, the ability to remove well-behaved singularities establishes a powerful compactness theorem on the space of pseudoholomorphic curves, known as Gromov compactness. For the sake of brevity, we will not fully state all of the definitions necessary to define Gromov convergence, stable maps, and the compactness theorem itself. Many standard texts do cover the theorem, as it is important to the study of pseudoholomorphic curves; for a more complete discussion of the Gromov topology on the moduli spaces of pseudoholomorphic curves, we direct the reader to [6]. However, we will still provide a brief overview of the theorem and its intuition.

Theorem 5.9. *If (M, ω) is a symplectic manifold, J_n a sequence of ω-compatible almost-complex structures converging to J in C^∞, and $f_n : S^2 \to M$ a sequence of pseudoholomorphic maps with respect to J_n with uniformly bounded energy, then f_n Gromov converges to a stable map with respect to J.*

Stable maps are, in essence, a collection of pseudoholomorphic maps defined on tangent spheres. These spheres are not allowed to form "cycles", and the points of tangency correspond to overlapping images of the corresponding pseudoholomorphic maps. More broadly, stable maps represent the ways in which convergence of pseudoholomorphic functions can fail. These failures are called "bubbles", and handling these bubbling events can be done by analysing the energy of each bubble independently. Not all of these "bubbles" need to be problematic, however; a constant bubble does not represent a failure of convergence, for instance.

Gromov convergence is stronger than typical notions of convergence. First and foremost, the parameterisation of the f_n's is not of any interest; rather it is the image of the f_n's that is being considered. Hence, one condition for convergence is that $f_n \circ \phi_n$ converges locally uniformly on compact sets of S^2 for some collection ϕ_n of Möbius automorphisms of S^2. In other words, some reparameterisation of f_n converges in the usual sense.

Gromov convergence also stipulates some relations on the sequence of f_n's on each of the bubbles that pop up. For one, energy is not allowed to dissipate from the bubbling sites; this ensures that the limit of f_n still retains much of the geometric information of the images of each f_n and does not collapse or "pinch off". The full conditions are much more technical, and again, we refer the reader to [6] for a more accurate and complete discussion of these notions.

6. Gromov's Non-Squeezing Theorem and Concluding Remarks

We will now turn our attention to Gromov's non-squeezing theorem, which is a central theorem in modern symplectic geometry.

Theorem 6.1. *Suppose there is a symplectic embedding of*

$$B^{2n}(r) := \{(x_1, \ldots, x_{2n}) : \sum_{j=1}^{2n} x_j^2 \leq r\} \subset \mathbb{R}^{2n}$$

into

$$Z^{2n}(R) := \{(x_1, \ldots, x_{2n}) : x_1^2 + x_2^2 \leq R\} \subset \mathbb{R}^{2n}.$$

Then, $r \leq R$.

Notably, the symplectic forms given to these manifolds is not determined. At least one symplectic form exists, and it is the one inherited from the ambient space $\mathbb{R}^{2n}$. However, these may not be compatible with each other.

We shall follow [7] and [9] for this proof outline. Suppose that such an embedding $B^{2n}(r) \hookrightarrow Z^{2n}(R)$ really does exist. The image of $B^{2n}(r)$ must be compact, and so by rewriting $Z^{2n}(R)$ as $B^2(R) \times \mathbb{R}^{2n-2}$, we can actually embed $B^{2n}(r)$ further into $S^2 \times T^{2n-2}$, where $T^{2n-2} := \mathbb{R}^{2n-2}/N\mathbb{Z}^{2n-2}$ for sufficiently large N. Importantly, the symplectic structure of $\mathbb{R}^{2n-2}$ factors through the quotient because ω_{std} is unaffected by the action of $N\mathbb{Z}^{2n-2}$.

The symplectic structure of S^2 is not as obvious or natural; it is instead obtained by "wrapping" the disc $B^2(R)$ onto the unit sphere S^2 equipped with an appropriately scaled symplectic form σ, which is also an area form, so that the areas are preserved. This step is actually quite subtle, and for any $\epsilon > 0$, an area form σ_ϵ is chosen such that

$$\int_{S^2} \sigma_\epsilon = \pi(R + \epsilon)^2.$$

These symplectic embeddings actually do exist and are easier to write down and work with. Then, one lets $\epsilon \to 0$ in the final inequality. We shall omit this detail for notational clarity.

Let $\iota : B^{2n}(r) \hookrightarrow S^2 \times T^{2n-2}$ be the extended embedding. Let J_0 be the almost-complex structure compatible with the above symplectic form on $S^2 \times T^{2n-2}$. This may not agree with the almost-complex structure of $\iota(B^{2n}(r))$ induced by its symplectic form, which we will call J_1. Since the space of almost-complex structures on any symplectic manifold is contractible, we may pick a continuous path J_t between J_0 and J_1, where t varies over $[0, 1]$.

The next step applies topological arguments about homologous curves from the Riemann sphere into $B^2(R) \times T^{2n-2}$, and this argument will be omitted because the fundamentals necessary were not developed here. The main idea is that by taking certain moduli spaces of pseudoholomorphic functions (with respect to J_t), one can construct a pseudoholomorphic function $f : \Sigma \to B^{2n}(r)$ passing through the origin, where Σ is a Riemann surface and $f(\Sigma)$ is contained on the boundary of $B^{2n}(r)$. At the same time, however, it can be shown that $\iota \circ f(\Sigma)$ falls into a circular "slice" of $S^2 \times T^{2n-2}$ with area no more than πR^2.

A crucial facet of this argument is the Gromov compactness theorem, which (loosely) ensures that the moduli spaces above can be compactified nicely. Pseudoholomorphic functions need not converge to other pseudoholomorphic functions. Namely, these limits may have singularities in the form of "bubbles" and "nodes". Fortunately, these cases can be excluded in this theorem by analysing certain bounds on the energies introduced by such singularities and obtaining a contradiction. A more precise formulation of this argument can be found in [9].

To conclude the proof, we restate the fact that $f(\Sigma)$ is a pseudoholomorphic curve in $B^{2n}(r)$ that contains the origin, whose boundary lies on $\partial B^{2n}(r)$, and whose area is πR^2. If $R < r$, this is impossible because pseudoholomorphic curves minimise areas. This argument is geometric and intuitive: if a loop on $\partial B^{2n}(r)$ has area less than πr^2, then it must lie entirely in one hemisphere of that sphere. However, if $f(\Sigma)$ is minimal as well, it cannot contain the origin. Hence, it follows that $r \leq R$, concluding the proof.

This concludes our discussion of how pseudoholomorphic curves play an integral role in the modern study of symplectic manifolds, specifically in the proof of Gromov's non-squeezing theorem. Some facets of the proof regrettably had to be omitted; they utilised tools from algebraic topology that would have been impossible to develop alongside our discussion of symplectic geometry and pseudoholomorphic curves. For a full proof, we refer the reader to [9] or even [4]. Symplectic geometry and the study of pseudoholomorphic curves are both also immensely rich fields of study, and we encourage readers to read [5] and [2] for more thorough coverage of each topic.

Gromov's non-squeezing theorem itself is a surprisingly intuitive yet deep theorem regarding the existence of global invariants of symplectic manifolds. One such invariant is the Gromov width of a symplectic manifold, which is informally the largest ball one can symplectically embed in a given symplectic manifold. These invariants are also called symplectic capacities, and this name may be more suiting because of desirable monotonicity and non-degeneracy properties one may want them to have. Again, symplectic capacities have been shown to exist independently of Gromov's non-squeezing theorem and can therefore be used to prove the theorem itself. This method of proof relies more heavily on the tools of Hamiltonian dynamics, however, and for a fuller discussion of this perspective, we again direct the reader to [5]. Despite the gaps in our coverage, we hope that this discussion was able to capture the richness of symplectic geometry from a purely mathematical viewpoint and hint at their importance in other applied fields of mathematics, namely physics.

References

[1] Carlos Saéz Calvo. "Gromov compactness theorem for pseudoholomorphic curves". MA thesis. University of Barcelona, 2014.

[2] Ana Cannas de Silva. *Lectures on Symplectic Geometry*. https://people.math.ethz.ch/~acannas/Papers/lsg.pdf. Jan. 2006.

[3] Simon K. Donaldson. "What is... a Pseudoholomorphic Curve". In: *Notices* 52.9 (Oct. 2005), pp. 1026–1027.

[4] M. Gromov. "Pseudo holomorphic curves in symplectic manifolds". In: *Inventiones Mathematicae* 82.2 (1985), pp. 307–347.

[5] D. McDuff and D. Salamon. *Introduction to Symplectic Topology*. 3rd ed. Oxford University Press, 2017.

[6] D. McDuff and D. Salamon. *J-holomorphic Curves and Symplectic Topology*. 2nd ed. American Mathematical Society, 2012.

[7] R. Pedrotti. *A proof of the non-squeezing theorem*. https://web.ma.utexas.edu/users/pedrotti.riccardo/Riccardo_Pedrotti-Non_squeezing_theorem.pdf. 2019.

[8] Zuoqin Wang. *Lectures on Symplectic Geometry*. http://staff.ustc.edu.cn/~wangzuoq/Courses/15S-Symp/SympGeom.html. 2015.

[9] Chris Wendl. *Lectures on Holomorphic Curves in Symplectic and Contact Geometry*. https://www.mathematik.hu-berlin.de/~wendl/pub/jhol_bookv33.pdf. May 2015.