CHAPTER 5: SYMPLECTIC AND CONTACT GEOMETRY

DANNY CALEGARI

Abstract. These are notes on Symplectic and Contact geometry, mostly in 4 and 3 dimensions, which are being transformed into Chapter 5 of a book on 3-Manifolds.

Contents

1. Symplectic geometry

Symplectic geometry is a vast subject, and we barely scratch the surface here. Our main interest in symplectic geometry is twofold. Firstly, we are interested in symplectic 4-manifolds insofar as the might bound or contain 3-manifolds in geometrically interesting ways. Secondly, we are interested in the symplectic geometry of certain parameter spaces associated to 3-manifolds, and the ways in which the Floer–Gromov theory of pseudoholomorphic curves can be used to obtain invariants of 3-manifolds.

An excellent introduction to Symplectic Geometry in [17]. An introduction more focused on the relationship of symplectic 4-manifolds with 3-manifolds is [6]. The gold standard of exposition of the analytic foundations of pseudoholomorphic curves is [18]. Gromov's paper [13] is inspirational and laconic, and should be read in concert with more sober references.

1.1. Linear algebra.

Definition 1.1. Let V be a real vector space. A *symplectic structure* on V is a bilinear 2-form $\omega : V \times V \to \mathbb{R}$ that is

(1) antisymmetric: $\omega(u, v) = -\omega(v, u)$ for all $u, v \in V$; and
(2) nondegenerate: for any nonzero $u \in V$ there is $v \in V$ with $\omega(u, v) \neq 0$.

With respect to a basis for V, any bilinear 2-form ω is represented by a matrix J so that $\omega(u, v) = v^T Ju$. The form is antisymmetric if and only if J is antisymmetric, and it is nondegenerate if and only if $\det(J) \neq 0$. The *symplectic group* $\mathrm{Sp}(V, \omega)$ is the group

Date: November 28, 2023.

of automorphisms of V that preserve ω. With respect to a basis as above, these are the matrices A for which $A^T JA = J$.

An odd dimensional antisymmetric matrix has kernel; thus, only even dimensional vector spaces admit symplectic structures.

Example 1.2 (Standard symplectic $\mathbb{R}^{2n}$). On $\mathbb{R}^{2n}$ with coordinates $x_1, \cdots, x_n, y_1, \cdots, y_n$ define a 2-form ω by

$$\omega(x_i, x_j) = \omega(y_i, y_j) = 0; \quad \omega(x_i, y_j) = \begin{cases} 0 \text{ if } i \neq j \\ 1 \text{ if } i = j \end{cases}$$

In this basis ω is represented as above by the block matrix $J = \left(\begin{smallmatrix} 0 & -I_n \\ I_n & 0 \end{smallmatrix}\right)$ where I_n is the $n \times n$ identity matrix.

Definition 1.3. If V, ω is symplectic and $U \subset V$ is a linear subspace, define the *orthogonal complement* $U^\perp$ to be the subspace

$$U^\perp := \{v \in V \text{ such that } \omega(u, v) = 0 \text{ for all } u \in U\}$$

The subspace U is said to be

(1) symplectic if $U \cap U^\perp = 0$;
(2) isotropic if $U \subset U^\perp$;
(3) coisotropic if $U^\perp \subset U$; and
(4) Lagrangian if $U = U^\perp$.

There is a map $\sigma : V \to V^*$ that takes $v \in V$ to $\alpha_v \in V^*$ defined by $\alpha_v(u) = \omega(u, v)$. Nondegeneracy of ω implies that σ is an isomorphism. An inclusion $U \to V$ is dual to a surjection $V^* \to U^*$ whose kernel is $\sigma(U^\perp)$. Thus $\dim(U) + \dim(U^\perp) = \dim(V)$, and Lagrangian subspaces all have dimension $\dim(V)/2$.

In the standard symplectic $\mathbb{R}^{2n}$ the subspaces X spanned by the x_i and the subspace Y spanned by the y_i are Lagrangian.

Lemma 1.4. *(1) Any isotropic subspace $U \subset V, \omega$ is contained in a Lagrangian subspace.*
(2) Any Lagrangian subspace $L \subset V, \omega$ has a Lagrangian complement; i.e. a Lagrangian subspace L' with $L \cap L' = 0$ and $L \oplus L' = V$.
(3) If $L \subset V, \omega$ is Lagrangian, and L' is any Lagrangian complement, any isomorphism $L \to X \subset \mathbb{R}^{2n}$ extends to a unique symplectic isomorphism $V \to \mathbb{R}^{2n}$ sending $L' \to Y$.
(4) If $U \subset V, \omega$ is symplectic, any symplectic isomorphism $U \to \mathbb{R}^{2m}$ extends to a symplectic isomorphism $V \to \mathbb{R}^{2n}$.

Proof. If $U \subset V$ is isotottropic but not Lagrangian then U is properly contained in $U^\perp$. If $v \in U^\perp - U$ is arbitrary, $\langle U, v\rangle$ is isotropic. This proves (1).

Suppose $L \subset V$ is Lagrangian. Let U be isotropic of dimension less than $\dim(V)/2$ with $L \cap U = 0$. The map $\sigma : L \to V^* \to U^*$ is surjective, for otherwise there are distinct $u_1, u_2 \in U$ that pair the same way with every element of L, which implies that $u_1 - u_2 \in L$ contrary to the definition of U. Thus, if w is any nonzero vector with $\dim(\langle U, w\rangle) = \dim(U) + 1$ and $L \cap \langle U, w\rangle = 0$ then there is $v \in L$ for which v and w have the same image in U^*. It follows that $\langle U, w-v\rangle$ is isotropic, contains U, and has dimension equal to $\dim(U)+1$. This proves (2).

The map $\sigma : \mathbb{R}^{2n} \to (\mathbb{R}^{2n})^*$ takes Y isomorphically to X^*. Likewise, if we choose any splitting $V = L \oplus L'$ the map $\sigma : V \to V^*$ takes L' isomorphically to L^*. Thus there is a unique isomorphism $L' \to Y$ respecting these maps, and therefore a unique symplectic isomorphism $L \oplus L' \to X \oplus Y$. This proves (3).

(3) implies that any symplectic vector space is isomorphic to the standard symplectic vector space of the same dimension. If $U \subset V$ is symplectic and we have an isomorphism $U \to \mathbb{R}^{2m}$ then $U^\perp$ is also symplectic so by (2) there is an isomorphism $U^\perp \to \mathbb{R}^{2n-2m}$, and therefore $U \oplus U^\perp \to \mathbb{R}^{2m} \oplus \mathbb{R}^{2n-2m}$, proving (4). □

Example 1.5 (Standard Hermitian $\mathbb{C}^n$). Let $W = \mathbb{C}^n$ with coordinates $z_1, \cdots, z_n$ and the standard Hermitian form $h(u,v) = v^*u$ where v^* means conjugate transpose. This is complex linear in the first argument and complex antilinear in the second argument. Let $V = W_\mathbb{R}$, the underlying real vector space of W. If we write $z_j = x_j + iy_j$ then x_j, y_j define coordinates on V making it isomorphic to $\mathbb{R}^{2n}$.

With respect to this isomorphism, the real part of h is the usual Euclidean inner product on $\mathbb{R}^{2n}$ and the imaginary part of h is the standard symplectic structure on $\mathbb{R}^{2n}$.

A real subspace of W is said to be *totally real* if the restriction of h to W takes real values. The maximal totally real subspaces of W are precisely the Lagrangian subspaces of $\mathbb{R}^{2n}$.

Let's consider $\mathbb{C}^n$ with the standard Hermitian form h with real part g and imaginary part ω. The complex linear automorphisms of $\mathbb{C}^n$ are $\mathrm{GL}(n,\mathbb{C})$, the real linear automorphisms preserving g are $\mathrm{O}(2n,\mathbb{R})$ and the real linear automorphisms preserving ω are $\mathrm{Sp}(2n,\mathbb{R})$. Finally, the complex linear automorphisms preserving h are $\mathrm{U}(n)$. There are isomorphisms

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

Example 1.6 (Lagrangian as a coset space). Let $\mathcal{L}_n$ denote the space of Lagrangian subspaces of standard symplectic $\mathbb{R}^{2n}$, equivalently the space of maximal totally real subspaces of standard Hermitian $\mathbb{C}^n$. The group $\mathrm{U}(n)$ acts transitively on $\mathcal{L}_n$ with stabilizers conugate to $\mathrm{O}(n;\mathbb{R})$; thus $\mathcal{L}_n$ is isomorphic to the coset space $\mathrm{U}(n)/\mathrm{O}(n)$. The map $\det^2 : \mathrm{U}(n)/\mathrm{O}(n) \to S^1$ exhibits $\mathcal{L}_n$ as a fiber bundle over S^1.

$\mathcal{L}_1$ is a circle. It may be identified with $\mathbb{RP}^1$, the space of lines in $\mathbb{R}^2$.

$\mathcal{L}_2$ is an S^2 bundle over S^1. In fact it is the twisted (nonorientable) S^2 bundle over S^1, with monodromy the antipodal map on S^2.

Definition 1.7. Let V be a real vector space of dimension $2n$. A *complex structure* on V is an endomorphism J with $J^2 = -\mathrm{Id}$. A complex structure and a symplectic structure ω are *compatible* if $\omega(u,v) = \omega(Ju,Jv)$ for all $u,v \in V$ and $\omega(v,Jv) > 0$ for all nonzero $v \in V$.

A compatible complex and symplectic structure ω, J together determine a J-invariant positive definite inner product $g(u,v) = \omega(u,Jv)$, and thereby a J-invariant Hermitian form with real part g and imaginary part ω. This witnesses $\mathrm{GL}(n,\mathbb{C}) \cap \mathrm{Sp}(2n,\mathbb{R}) = \mathrm{U}(n)$.

Lemma 1.8. *The space of compatible complex structures on the standard symplectic $\mathbb{R}^{2n}, \omega$ is the coset space $\mathrm{Sp}(2n,\mathbb{R})/\mathrm{U}(n)$. This space is contractible.*

Proof. We have already seen that a complex structure compatible with a symplectic structure determines a Hermitian form, whose stabilizer is a copy of $\mathrm{U}(n)$ in $\mathrm{Sp}(2n,\mathbb{R})$ conjugate to the standard one. This proves the first claim. To see that the coset space is contractible, observe that a symplectic matrix $A \in \mathrm{Sp}(2n,\mathbb{R})$ has a polar decomposition $A = PU$ where P is symmetric positive definite and U is orthogonal. A calculation shows that P and U are both separately symplectic, and therefore the one parameter family P^tU as t goes from 1 to 0 defines a deformation retraction of $\mathrm{Sp}(2n,\mathbb{R})$ onto $\mathrm{U}(n)$. □

1.2. **Symplectic manifolds.** Let W be an oriented $2n$-manifold. A *symplectic structure* on W is a 2-form ω which is

(1) closed: $d\omega = 0$; and
(2) nondegenerate: ω^n is nowhere zero.

Since ω^n is nowhere zero, it determines an orientation on W. Since ω is closed, it determines a cohomology class $[\omega] \in H^2(W;\mathbb{R})$ and if W is closed, $[\omega]^n \neq 0$ in $H^{2n}(W;\mathbb{R})$.

Remark 1.9. Any nondegenerate 2-form ω induces a symplectic structure on the fibers of TW (i.e. it gives TW the structure of a symplectic vector bundle).

Example 1.10. Oriented surfaces with area forms are symplectic. Products of symplectic manifolds are symplectic. Open submanifolds of symplectic manifolds are symplectic.

Example 1.11. Smooth complex projective varieties are symplectic. To see this, first realize $\mathbb{CP}^n$ as S^{2n+1}/S^1 for S^{2n+1} the unit sphere in $\mathbb{C}^{n+1}$. This determines a Hermitian form on $T\mathbb{CP}^n$ by identifying tangent vectors to $\mathbb{CP}^n$ with equivalence classes of vectors in S^{2n+1} perpendicular to the S^1 orbits, and restricting the standard Hermitian form on $\mathbb{C}^{n+1}$. The imaginary part of the induced Hermitian form on $\mathbb{CP}^n$ is closed, and non-degenerate on complex subspaces of $T\mathbb{CP}^n$.

Example 1.12. If M^n is a smooth manifold, there is a 'tautological' 1-form α on T^*M defined as follows. Let $\pi : T^*M \to M$ denote projection, inducing $d\pi : TT^*M \to TM$. If $p \in M$ and $q \in T^*_pM$ and $v \in T_{(p,q)}T^*M$ then $\alpha(v) := q(d\pi(v))$. If p_i are local coordinates on M, then dp_i is a basis for T^*M locally and we may choose coordinates q_i on the fibers in this basis, so that p_i, q_i become local coordinates on T^*M, and then $\alpha = \sum q_i dp_i$. Thus $\omega := -d\alpha = \sum dp_i \wedge dq_i$ is nondegenerate and therefore symplectic.

1.3. **Theorems of Moser and Darboux.**

Theorem 1.13 (Moser)**.** *Let $\omega(t)$ be a 1-parameter family of symplectic forms on a closed manifold W, and suppose the cohomology classes $[\omega(t)]$ are constant. Then the $\omega(t)$ are all pulled back from $\omega(0)$ by a 1-parameter isotopy.*

Proof. By hypothesis $\dot\omega(t)$ are all exact, so we can find a smooth 1-parameter family of 1-forms $\alpha(t)$ with $d\alpha(t) = \dot\omega(t)$. Since each $\omega(t)$ is nondegenerate there is a smooth family of vector fields $X(t)$ so that $\alpha(t)(Y) = \omega(t)(X,Y)$ for all t; i.e. $\iota_{X(t)}\omega(t) = \alpha(t)$. By Cartan's magic formula

$$\mathcal{L}_{X(t)}\omega(t) = \iota_{X(t)}d\omega(t) + d\iota_{X(t)}\omega(t) = \dot\omega(t)$$

Thus the $\omega(t)$ are all pulled back from $\omega(0)$ by the flow generated by $X(t)$. □

Theorem 1.14 (Darboux). *Let W^{2n}, ω be symplectic. Then every point has a neighborhood symplectomorphic to an open subset of the standard symplectic $\mathbb{R}^{2n}$.*

Proof. For arbitrary $p \in W$ choose a smooth embedding $\phi : U \to \mathbb{R}^{2n}$ for some open neighborhood U of p, and let ω' be the pullback of the standard symplectic form to U. After composing ϕ with a linear automorphism of $\mathbb{R}^{2n}$ if necessary we may assume that $\omega' = \omega$ at p; thus there is a smaller open neighborhood V of p for which $\omega(t) := t\omega + (1-t)\omega'$ is symplectic for $t \in [0,1]$.

Using ϕ we may identify V with its image in $\mathbb{R}^{2n}$ and think of $\omega(t)$ as a 1-parameter family of symplectic forms on a neighborhood of 0. Since $\mathbb{R}^{2n}$ is contractible we may write $\dot{omega}(t) = \alpha(t)$ for some 1-parameter family of 1-forms, and by subtracting a constant 1-form from each $\alpha(t)$ we may assume $\alpha(t)$ vanishes at the origin. Then as above $\alpha(t) = \iota_{X(t)}\omega(t)$ for some family of vector fields $X(t)$ definined in a neighborhood of 0. Since each $X(t)$ vanishes at 0, the vector fields generate a flow on an open neighborhood of 0 defined for all $[0,1]$. Thus, as in the proof of Moser, ω and ω' are symplectically isomorphic on some neighborhoods of 0 (and the same is true at p). □

1.4. **Almost complex structures.**

Definition 1.15 (Almost complex). If W is a (necessarily even dimensional oriented) manifold, an *almost complex* structure on W is a choice of (smoothly varying) complex structure on the fibers of TW.

An almost complex structure on W is equivalent to the data of a section J of the bundle of automorphisms of TW with $J^2 = -\mathrm{Id}$ (the automorphism J acts on each fiber as multiplication by i). The existence of an almost complex structure on a manifold is a purely homotopy theoretic condition, that the classifying map of the tangent bundle lifts to complex Grassmannian.

Definition 1.16 (Tame). Let J be an almost complex structure on W. A closed 2-form ω on W is said to *tame* J if ω is strictly positive on the complex lines of TW.

A closed 2-form that tames any J is necessarily symplectic. Furthermore, for a fixed almost-complex structure, the set of closed 2-forms that tame it form a convex subspace of $\Omega^2(W)$; this fact allows partition of unity arguments in the construction of symplectic structures, as we shall see in the proof of Theorem 1.18.

Definition 1.17 (Compatible). A closed 2-form ω is *compatible with* J if the structures are compatible on each fiber of TW in the sense of Definition 1.7.

A 2-form ω compatible with J necessarily tames it. If W, ω is symplectic, there is always a J tamed by ω (resp. compatible with ω); in fact the set of J tamed by ω (resp. compatible with ω) is contractible. This follows from Lemma 1.8. Thus the tangent bundle of any symplectic manifold is a complex vector bundle in a canonical way (up to isomorphism) and we may refer to the Chern classes $c_j \in H^{2j}(W;\mathbb{Z})$ of this complex structure.

1.5. **Lefschetz fibrations.**

A (4-dimensional) *topological Lefschetz fibration* is the data of a closed oriented 4-manifold W and a map $\pi : W \to S^2$ such that

(1) π is a submersion away from finitely many *singular points* where there are local (oriented) complex co-ordinates z_1, z_2 on W and z on S^2 so that in these co-ordinates π has the form $z = z_1^2 + z_2^2$;
(2) there is at most one singular point on each fiber; and
(3) there is some $c \in H^2(W)$ with $c[F] = 1$ for every (nonsingular) fiber F.

The adjective 'topological' will be sometimes be omitted in the sequel out of laziness. We make some remarks on this definition.

(1) W as above admits an almost complex structure J for which the tangent space to the fibers are complex lines. This is obvious near the nonsingular fibers (take J to preserve an orthogonal splitting $TW = TF \oplus TF^\perp$ for some Riemannian metric) and one sees that this structure pieces together with the complex structure given by hypothesis near the singular points.
(2) Since all nonsingular fibers are homologous, (3) is equivalent to the condition that the nonsingular fibers F are nontrivial in homology (with real coefficients). Since $[F] \cap [F] = 0$ it follows that the first Chern class c_1 associated with the almost complex structure satisfies $c_1[F] = \chi(F)$. Thus (3) is automatically satisfied when F has genus $\neq 0$.
(3) Since π is a submersion away from the singular points, a neighborhood of a nonsingular fiber is a product $F \times D$ where F is a closed oriented surface, and the fibers are $F \times$ point. If $x \in W$ is a singular point, and $p = \pi(x)$, and $p \in D \subset S^2$ is a disk for which x is the only singular point in $\pi^{-1}(D)$, then $\pi^{-1}(D - p)$ is an F bundle with monodromy (in the positive direction) equal to a right-handed Dehn twist around some essential simple closed curve $\gamma \subset F$ called a *vanishing cycle*. As $q \to p$ the curves γ_q in the fibers F_q degenerate to the point x; i.e. F_p is homeomorphic to the quotient of F with γ pinched to a point.
(4) One may generalize the definition of a topological Lefschetz fibration by allowing the base to be any closed oriented surface Σ; by abuse of notation we refer to these more general structures also as topological Lefschetz fibrations.

Theorem 1.18 (Lefschetz fibration is symplectic)**.** *Let $W^4 \to \Sigma^2$ be a (generalized) topological Lefschetz fibration. Then W is symplectic.*

Proof. By (3) there is $c \in H^2(W)$ with $c[F] = 1$ for every fiber F. We first construct a closed 2-form η on W^4 which is strictly positive on TF for every fiber, and such that $[\eta] = c$. Fix a representative 2-form ζ of the de Rham class c.

A neighborhood N of a nonsingular fiber F is diffeomorphic to a product $F \times D^2$ and we may therefore project it to F and pull back an area form (of total area 1) to obtain a form strictly positive on $TF|N$ and of the form $\zeta + d\alpha$ where α is a 1-form on N.

If $x \in F$ is a singular point, then by (2) near x there are (oriented) local complex coordinates $(z_1, z_2) \in \mathbb{C}^2$ for which $\pi(z_1, z_2) = z_1^2 + z_2^2$. Choose a closed 2-form in a neighborhood of x modeled on the standard symplectic form on a ball in $\mathbb{C}^2$, and extend it to a closed 2-form on a neighborhood N of F which is an area form on each fiber. Since $H^2(N, \mathbb{R}) = \mathbb{R}$ this will be of the form $\zeta + d\alpha$ where α is a 1-form on N.

Take a partition of unity ρ_y on S^2 whose supports are contained in small disks D_y for which each $N_y := \pi^{-1}(D_y)$ is a neighborhood as above of an ordinary or singular fiber with associated 2-form $\zeta + d\alpha_y$. Define

$$\eta := \zeta + d\sum_y (\rho_y \circ \pi)\alpha_y = \sum_y (\rho_y \circ \pi)(\zeta + d\alpha_y) + \sum_y (d\rho_i \circ d\pi) \wedge \alpha_y$$

Each term of $\sum_y (d\rho_i \circ d\pi) \wedge \alpha_y$ vanishes on TF, and the remainder is strictly positive on each TF by construction. Moreover, η as above is closed.

Now let β be an area form on S^2 and define $\omega := \eta + t\pi^*\beta$. This form is closed; we claim that for sufficiently large t it is nondegenerate (and therefore symplectic). To see this, put a complex structure J on TW preserving an orthogonal splitting $TW = TF \oplus TF^\perp$. The form $\pi^*\beta$ is non-negative on all complex lines, and strictly positive on all except the vertical lines (those tangent to TF). The form η is strictly positive on TF and (by continuity) on some open neighborhood of the vertical complex lines. Furthermore, η tames J in a neighborhood of the singular points. Thus some $\eta + t\pi^*\beta$ tames J and is therefore symplectic. $\square$

1.6. **Character varieties.**

Definition 1.19. Let G be a (real or complex) matrix Lie group with Lie algebra $\mathfrak{g}$ and discrete center Z, and let M be a manifold with finitely generated fundamental group.

The *representation variety* is the set $R(M, G) = \mathrm{Hom}(\pi_1(M), G)$, thought of as an algebraic variety with coordinates the matrix entries of the $\rho(g_j)$ where the g_j are generators of $\pi_1(M)$, and $\rho : \pi_1(M) \to G$ is a homomorphism, which is cut out by equations of the form $\rho(w) = \mathrm{Id}$ for all relations w in the generators and their inverses.

Let $R^*(M, G)$ denote the set of irreducible representations. The group G acts on R and R^* by conjugation. The conjugation action of G on $R^*(M, G)$ is locally free (it factors through a free action of G/Z), and we denote the quotient $X^*(M, G)$.

Now let's specialize to the case that M is a closed oriented surface Σ of genus g, and let Σ^* denote Σ minus a point p. The fundamental group of Σ^* is free on $2g$ generators $a_1, b_1, \cdots, a_g, b_g$, and the fundamental group of Σ is the quotient of this by the relation $\prod_i [a_i, b_i] = 1$. For convenience let's denote $c := \prod_i [a_i, b_i]$.

Lemma 1.20. *Let Σ have genus g. Then $X^*(\Sigma, G)$ is a smooth manifold of dimension $(2g - 2)\dim(G)$.*

Proof. Evidently $R(\Sigma^*, G) = G^{2g}$. The inclusion $\Sigma^* \to \Sigma$ induces a surjection on π_1 and therefore an inclusion $R(\Sigma, G) \to R(\Sigma^*, G)$. Define $\sigma : R(\Sigma^*, G) \to G$ by $\sigma(\rho) = \rho(c)$. If $\rho \in R(\Sigma, G)$ then $\sigma(\rho) = \mathrm{Id}$ and $d\sigma : T_\rho R(\Sigma, G) \to \mathfrak{g}$. We claim that when ρ is irreducible, $d\sigma$ is surjective. By the implicit function theorem this will imply that $R^*(\Sigma, G)$ is a smooth manifold of dimension $(2g - 1)\dim(G)$, and therefore (since G acts locally freely on the irreducible representations) $X^*(\Sigma, G)$ is a smooth manifold of dimension $(2g - 2)\dim(G)$ as claimed.

So, let ρ be irreducible with $\sigma(c) = \mathrm{Id}$. We may apply an automorphism of Σ^* (if necessary) to assume that $\alpha := \rho(a_g)$ and $\beta := \rho(b_g)$ are nonzero, and together their image in G is irreducible. Since the image is irreducible, the only elements of G that simultaneously commute with α and β are in the center; consequently $(\mathrm{Ad}_{\alpha^{-1}} - 1)\mathfrak{g}$ and

$(\mathrm{Ad}_{\beta^{-1}} - 1)\mathfrak{g}$ together span $\mathfrak{g}$. Then if $u, v \in \mathfrak{g}$ are arbitrary,

$$\begin{aligned}\frac{d}{dt}\Big|_{t=0}[\alpha(1+tu), \beta(1+tv)] &= \frac{d}{dt}\Big|_{t=0}\alpha(1+tu)\beta(1+tv)\alpha^{-1}(1-t\mathrm{Ad}_\alpha(u))\beta^{-1}(1-t\mathrm{Ad}_\beta(v)) \\ &= \alpha u\beta\alpha^{-1}\beta^{-1} + \alpha\beta v\alpha^{-1}\beta^{-1} - \alpha\beta u\alpha^{-1}\beta^{-1} - \alpha\beta\alpha^{-1}v\beta^{-1} \\ &= [\alpha,\beta]\mathrm{Ad}_{\beta\alpha}\left(\mathrm{Ad}_{\beta^{-1}}(u) + v - u - \mathrm{Ad}_{\alpha^{-1}}(v)\right)\end{aligned}$$

and therefore $d\sigma$ is surjective at ρ. □

Let $[\rho] \in X^*(\Sigma, G)$ and let ρ be a homomorphism in the equivalence class of $[\rho]$. If ρ_t is a 1-parameter family of representations deforming $\rho_0 = \rho$, for any $g, h \in \pi_1(\Sigma)$ we may differentiate the identity $\rho_t(g)\rho_t(h) = \rho_t(gh)$ at $t = 0$ to obtain $\rho'(g)\rho(h) + \rho(g)\rho'(h) = \rho'(gh)$. Setting $c(g) := \rho(g)^{-1}\rho'(g) \in \mathfrak{g}$ this becomes the identity

$$\mathrm{Ad}_{\rho(h)^{-1}}c(g) + c(h) = c(gh)$$

Functions $c : \pi_1(\Sigma) \to \mathfrak{g}$ satisfying this equation are (by definition) group 1-cocycles with values in the $\pi_1(\Sigma)$-module $\mathfrak{g}$ (with the module structure coming from Ad_ρ).

If the class of $[\rho_t]$ is constant in $X^*(\Sigma, G)$ to first order, then there is some smooth $\gamma : [0, 1] \to G$ with $\gamma(0) = \mathrm{Id}$ so that up to first order, $\rho_t(g) := \gamma(t)\rho(g)\gamma(t)^{-1}$. For such a family, with $\gamma'(0) = v$ we have $c(g) = \mathrm{Ad}_{\rho(g)^{-1}}(v) - v$. Cocycles of this kind are (by definition) group 1-coboundaries. The quotient of the space of 1-cocycles by the space of 1-coboundaries is the cohomology group $H^1_\rho(\Sigma, \mathfrak{g})$. Thus we obtain an injective map $T_\rho X^*(\Sigma) \to H^1_\rho(\Sigma, \mathfrak{g})$.

Surjectivity is a tricky question for general Σ. If one tries to write down a an honest deformation ρ_t (say, as power series in t), 1-cocycles give values for the linear terms for which $\rho_t(g)\rho_t(h) = \rho_t(gh)$ to first order. In general there is an obstruction in $H^2_\rho(\Sigma, \mathfrak{g})$ to finding compatible quadratic terms for which this formula holds to second order; however, this obstruction vanishes for surface groups; see Goldman [12] Thm. ??. To see this, observe firstly that since $\pi_1(\Sigma^*)$ is free, the obstruction vanishes in this case and therefore

$$\dim T_\rho X^*(\Sigma^*) = \dim H^1_\rho(\Sigma^*, \mathfrak{g}) = (2g-1)\dim(\mathfrak{g})$$

Group 1-cocycles on $\pi_1(\Sigma^*)$ may take an arbitrary value on c (this is essentially equivalent to Lemma 1.20) so $\dim H^1_\rho(\Sigma, \mathfrak{g}) \le (2g-2)\dim(\mathfrak{g}) = \dim T_\rho X^*(\Sigma)$ and we conclude $T_\rho X^*(\Sigma) \to H^1_\rho(\Sigma, \mathfrak{g})$ is an isomorphism.

Cup product defines an antisymmetric bilinear form on $H^1_\rho(\Sigma, \mathfrak{g})$ with values in $H^2_\rho(\Sigma, \mathfrak{g}\otimes \mathfrak{g})$. If G is reductive, the *Killing form* on $\mathfrak{g}$, defined by $\langle A, B\rangle := \mathrm{tr}(AB)$, is a nondegenerate symmetric form. One obtains in this way an alternating 2-form ω on $X^*(\Sigma)$ by

$$H^1_\rho(\Sigma, \mathfrak{g}) \otimes H^1_\rho(\Sigma, \mathfrak{g}) \to H^2_\rho(\Sigma, \mathfrak{g}\otimes\mathfrak{g}) \to H^2_\rho(\Sigma, \mathbb{R}) \to \mathbb{R}$$

Goldman Thm. ?? shows that this form is symplectic. This may be proved by direct calculation though we do not do it here. A more conceptual (though technically more complicated) proof identifies X^* with the Marsden–Weinstein quotient for the action of the gauge group $\mathrm{Map}(\Sigma, G)$ on the (naturally symplectic) space of irreducible G-connections on a trivial G bundle over Σ.

1.7. **Symmetric products.** Let Σ be a closed oriented surface. A symplectic form on S is the same thing as an (oriented) area form, and any such form arises from some Riemannian metric. A Riemannian metric determines a conformal structure with respect to which S is isomorphic to a smooth complex projective curve.

The nth symmetric product $S^n\Sigma$ is the quotient of $\Sigma^n = \Sigma \times \cdots \times \Sigma$ by the symmetric group acting by permutation of the coordinates. Rather surprisingly, this turns out to have the natural structure of a smooth complex projective variety (of complex dimension n).

Example 1.21. Although it is not closed, we may consider the case $\Sigma = \mathbb{C}$. Then it turns out $S^n\mathbb{C} = \mathbb{C}^n$. To see this, we may think of $S^n\mathbb{C}$ as the space of unordered n-tuples of complex numbers, and $\mathbb{C}^n$ as the space of ordered n-tuples. The isomorphism is then given by the (biholomorphic!) map that sends a monic degree n polynomial to its set of roots.

2. Contact structures

Let M be an oriented 3-manifold. A *contact structure* is a smooth 2-plane field ξ that is nowhere integrable. If α is a local 1-form with $\xi = \ker(\alpha)$, nowhere integrability is equivalent to $\alpha \wedge d\alpha \neq 0$ everywhere. The contact structure is *positive* if some (hence any) α satisfies $\alpha \wedge d\alpha > 0$ everywhere, and *negative* if $\alpha \wedge d\alpha < 0$. This is well-defined independent of the choice of α, since multiplying α by a nowhere zero function f multiplies $\alpha \wedge d\alpha$ by f^2. Changing the orientation on M changes positive contact structures to negative ones and vice versa.

A contact structure is *co-oriented* if ξ is co-oriented (equivalently, oriented since M is oriented). This is equivalent to the existence of a global 1-form α as above. In the sequel we assume our contact structures are all co-oriented unless we explicitly say otherwise.

Example 2.1 (Standard contact structure on $\mathbb{R}^3$). On $\mathbb{R}^3$ the *standard* contact structure is $\xi := \ker(\alpha)$ where $\alpha = dz + xdy$; see Figure 1. Since $\alpha \wedge d\alpha = dx \wedge dy \wedge dz$ this is a positive contact structure. Positivity means that ξ rotates anticlockwise as one moves tangent to ξ in any direction;

Example 2.2 (Connection 1-form). On a surface S with a Riemannian metric the Levi–Civita connection defines a distribution ξ on the unit tangent bundle $M := UTS$ for which curves tangent to ξ are parallel. If α is the 1-form with $\ker(\alpha) = \xi$ and the restriction of α to each circle fiber is the signed angular measure (so that the integral of α around each fiber is 2π) then $-d\alpha$ is the (pullback from S of the) curvature 2-form of the metric on S. Thus if S is a surface of negative (resp. positive) curvature, M inherits a canonical positive (resp. negative) contact structure.

If ξ is a contact structure and α an associated 1-form, the contact condition implies that $d\alpha$ is non-degenerate on ξ. Thus $\ker(d\alpha)$ is transverse to ξ, so that there is a unique vector field X transverse to ξ with $X \in \ker(d\alpha)$ and $\alpha(X) = 1$. The vector field X with these properties is the *Reeb vector field* associated to ξ. By Cartan's formula $\mathcal{L}_X(\alpha) = 0$; i.e. the 1-parameter family of diffeomorphisms generated by X preserve α.

2.1. **Fillings and thickenings.** The contact geometry of 3-manifolds is related to the symplectic geometry of 4-manifolds by *filling* and *thickening*.

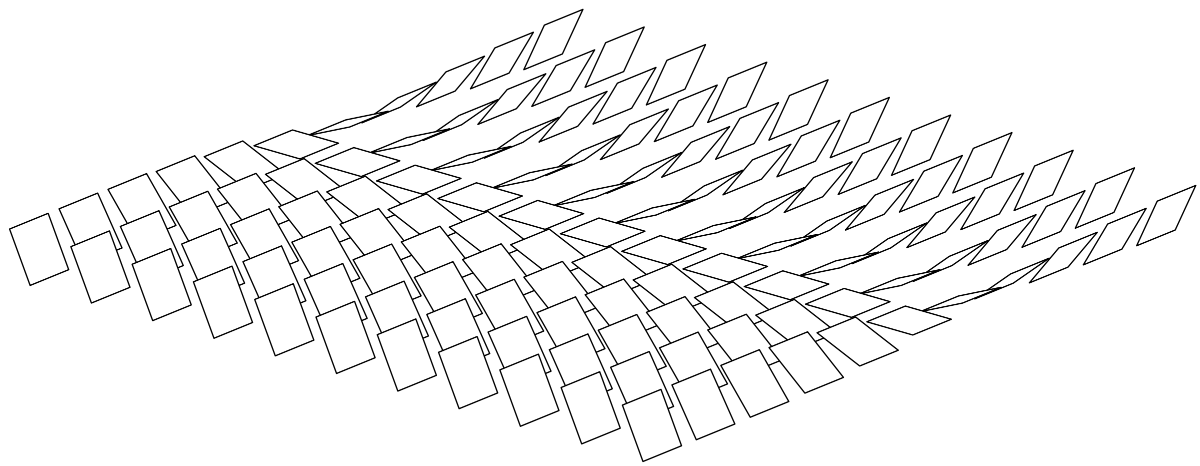

FIGURE 1. The standard contact structure on $\mathbb{R}^3$

Definition 2.3. A compact symplectic manifold W, ω is *strongly convex* if there is a vector field X defined in a neighborhood of ∂W satisfying $\mathcal{L}_X \omega = \omega$.

If W^4 is strongly convex, the form $\alpha := \iota_X \omega$ satisfies $d\alpha = \mathcal{L}_X \omega = \omega$ so that

$$\alpha \wedge d\alpha = \alpha \wedge \omega = \iota_X \omega \wedge \omega$$

which is a positively oriented volume form on $M := \partial W$ with respect to the orientation induced as the boundary of a symplectic manifold.

Definition 2.4 (Strongly fillable)**.** A (positive) contact structure M, ξ is *strongly fillable* (resp. *strongly semi-fillable*) if there is a strongly convex symplectic manifold W, ω with $M = \partial W$ (resp. M is one component of ∂W) inducing ξ on M.

Example 2.5 (Symplectic thickening). Let M, ξ be positive contact with 1-form α. The *symplectization* of M is $M \times \mathbb{R}$ with the symplectic form $\omega := d(e^t \alpha)$ (where t is the coordinate on $\mathbb{R}$). The vector field $X := \partial/\partial t$ satisfies $\mathcal{L}_X \omega = \omega$, and therefore $\iota_X \omega$ induces M, ξ on the slice $M \times 0$. However, $M \times (-\infty, 0]$ does not define a strong filling of M, ξ because it is non-compact.

Definition 2.6 (Weakly fillable)**.** A (positive) contact structure M, ξ is *weakly fillable* (resp. *weakly semi-fillable*) if there is a compact symplectic manifold W, ω with $\partial W = M$ so that $\omega|\xi$ is everywhere positive (resp. M, ξ is one component of a possibly disconnected M', ξ' that is weakly fillable).

Since $d\alpha$ is positive on ξ whenever α is a positive contact form, it follows that any strongly (semi-)fillable contact structure is weakly (semi-)fillable.

Example 2.7 (Standard contact structure on S^3). Let W be the closed unit ball in $\mathbb{C}^2$ with the standard complex/symplectic structure ω. The induced contact structure on S^3 is strongly filled by W, ω. If we remove any point from S^3, this induces the standard contact structure on $\mathbb{R}^3$.

Example 2.8 (Overtwisted contact structure on $\mathbb{R}^3$). Give $\mathbb{R}^3$ cylindrical coordinates r, θ, z and let ξ be the contact structure with contact form $\alpha = r\sin(r)d\theta + \cos(r)dz$. The disk D in the plane $z = 0$ with radius π is tangent to ξ at the center and along ∂D; see Figure 2.

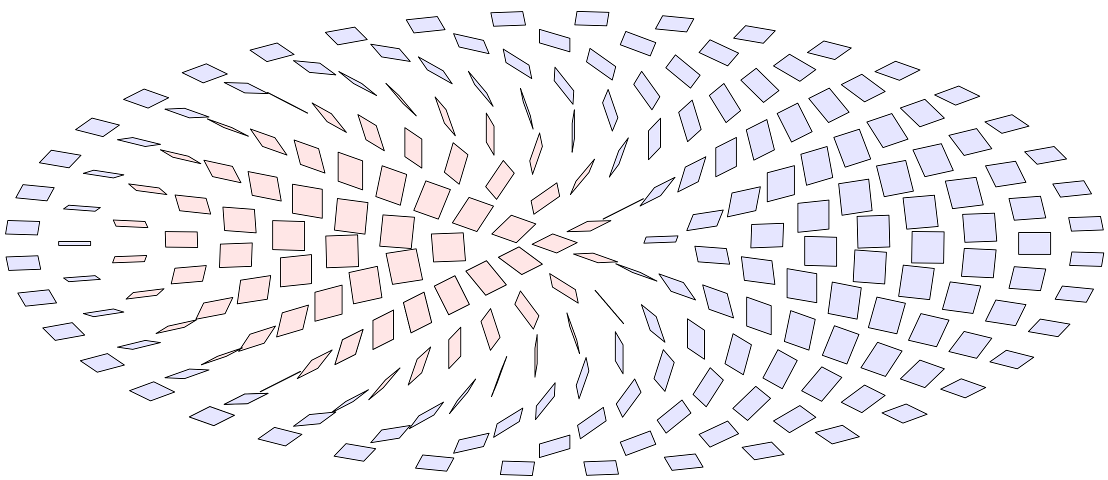

FIGURE 2. An overtwisted disk

Definition 2.9 (Overtwisted disk)**.** An *overtwisted disk* for a contact structure M, ξ is an embedded closed disk $D \subset M$ tangent to ξ along ∂D. A contact structure is *overtwisted* if it contains an overtwisted disk, and is *tight* otherwise.

Some (many) authors define an overtwisted disk to be one whose boundary is *transverse* to ξ (keeping ∂D tangent to ξ). A disk overtwisted in either sense may be perturbed to be overtwisted in the other sense, so which definition is used is a matter of convenience. What is really important is whether the contact structure admits such a disk (of either kind).

2.1.1. *Characteristic foliations.* Here is how overtwisted disks arise in practice. A compact embedded surface S may be isotoped to be in general position with respect to ξ; this means that it has isolated tangencies with ξ, and the boundary (if any) may be made transverse to ξ. For a surface in general position, the intersection with ξ is a line field with isolated singularities, that may be integrated to a singular foliation, known as the *characteristic foliation* and denoted S_ξ. An embedded subdisk whose boundary is a (nonsingular) closed curve of the characteristic foliation is an overtwisted disk (in the second sense above). The singularities (i.e. tangencies of S with ξ) are *elliptic* or *hyperbolic* depending on whether the local index of the characteristic foliation near the singularity is 1 or -1, and singularities are *positive* or *negative* depending on whether the (co)-orientations of S and ξ agree or disagree.

Note that if S and ξ are oriented, the characteristic foliation may also be oriented in such a way that $(\nu(\xi), \nu(S), S'_\xi)$ is an oriented basis (where $\nu(S)$ and $\nu(\xi)$ are the positive normals to S and ξ respectively). Thus a positive elliptic tangency is a source, whereas a negative elliptic tangency is a sink.

Theorem 2.10. *Suppose (M,ξ) is weakly semi-filled by W,ω. Then ξ is tight.*

Proof. We give a proof due to Gromov ([13], $2.4.D_2'$). This proof uses the machinery of pseudoholomorphic curves that will be developed in § 3, and until the reader has absorbed that section, it may read like science fiction.

Let's suppose (M,ξ) admits an embedded disk D overtwisted for ξ. Near the center tangency p the inclusion $p \in D \subset M \subset W$ is locally (symplectically) modeled on

$$(0,1) \in \{\mathrm{Im}(z_2)=0\} \cap S^3 \subset S^3 \subset B^4 \subset \mathbb{C}^2$$

Near p there is a family of disks E_t, holomorphic for the standard complex structure on $\mathbb{C}^2$ parameterized by $0<t$ real. In these local coordinates, for each fixed t

$$E_t := \{(z_1,t) \in \mathbb{C}\times\mathbb{R} \text{ such that } |z_1|^2 + (1-t)^2 \le 1\}$$

The Maslov index (and hence the real dimension) of this family is 1.

Choose a generic J-holomorphic structure compatible with ω (extending the model one in a neighborhood of the center) regular for the E_t, and continue the family for some maximal open interval $t \in (0,t_0)$. Since the boundary of W is pseudo-convex (in particular, it is strictly mean convex in holomorphic directions) the E_t may never make an interior tangency with M; thus in particular ∂E_t is never tangent to ξ (and therefore is strictly contained in the interior of D). Thus by Gromov's Compactness Theorem we can extract a limit E_{t_0} which is necessarily singular (or we could continue the family further).

But such a singular curve must be of the form $E_{t_0} = E' \vee S^2$ where S^2 is holomorphic, and therefore the Maslov index of E' is -1 which can't occur for generic J. This contradiction shows that no overtwisted D exists, so that ξ is tight as claimed. ∎

2.2. **Legendrian knots.**

Definition 2.11. Let (M,ξ) be a contact structure. A knot K in M is *Legendrian* if it is tangent to ξ.

A Legendrian knot K has a canonical framing, given by the intersection of the contact structure ξ with the normal bundle. If M is a rational homology sphere, this framing determines a self-linking number, known for historical reasons as the *Thurston–Bennequin number*, and usually denoted $\mathrm{tb}(K)$.

If K is an oriented Legendrian knot there is another invariant called the *rotation number* $r(K)$ defined with respect to a trivialization of $\xi|K$ as an $\mathbb{R}^2$ bundle, that measures how many times K' twists relative to the trivialization. Reversing the orientation of K changes the sign of $r(K)$.

Theorem 2.12 (Thurston–Bennequin inequality)**.** *Let K be a Legendrian knot in S^3 with the standard contact structure, with self-linking (i.e. Thurston–Bennequin) number $\mathrm{tb}(K)$. Let S be any Seifert surface for K, and let $r(K)$ be the rotation number of K with respect to a trivialization of $\xi|S$. Then $\mathrm{tb}(K) \pm r(K) \le -\chi(S)$.*

Proof. Orient K, and let v be a section of $\xi|K$ perpendicular to K' so that the orientation given by v, K' agrees with the orientation coming from the trivialization of $\xi|S$. Push K slightly in the direction of v (so that it is positively transverse to ξ in the given orientation), and let K^+ be obtained by pushing K off itself slightly further in the direction v, and K^S by pushing K in a direction which is parallel in the given trivialization. By definition

$$\mathrm{tb}(K) = \mathrm{link}(K, K^+) = \mathrm{link}(K, K^S) + r(K)$$

We shall show $\mathrm{link}(K, K^S) \le -\chi(S)$; the inequality follows from this.

Let $e \in H^2(S, \partial S)$ be the Euler class of $\xi|S$, relative to the trivialization of $TS|\partial S$ coming from K'. Then $\mathrm{link}(K, K^S) = -e[S]$. On the other hand, if we put S in general position with respect to ξ we may compute $e[S]$ from the characteristic foliation. Let e^+ (resp. e^-) denote the number of positive (resp. negative) elliptic tangencies, and likewise define h^+ (resp. h^-) to be the number of positive (resp. negative) hyperbolic tangencies. Then $\chi(S) = e^+ + e^- - h^+ - h^-$ whereas $e[S] = e^+ - e^- - h^+ + h^-$ and therefore

$$\mathrm{link}(K, K^S) = -\chi(S) + 2(e^- - h^-)$$

Thus the theorem will be proved if we can arrange for $e^- = 0$.

If an elliptic and hyperbolic pair of the same sign is joined by a trajectory of X we may eliminate them by a local move. Furthermore, in general position there are no trajectories between pairs of hyperbolic singular points. A negative elliptic tangency p is a sink of X; let E_p be the attracting basin of this sink. Then E_p is a disk, whose boundary is a (possibly singular) Legendrian curve. If ∂E_p contained no tangencies of ξ with S, it would be an overtwisted disk, contradicting the fact that S^3 is tight. Thus there must be some tangencies of ξ with S in ∂E_p and by what we have said so far these must alternate between positive elliptic tangencies and negative hyperbolic ones. Thus p is joined by a trajectory of X to a negative hyperbolic tangency, and therefore *a posteriori* does not exist. So $e^- = 0$ and the inequality is proved. □

2.3. **Confoliations.** The purpose of this section is to prove the following theorem:

Theorem 2.13 (Eliashberg–Thurston [5]; foliation to contact structure)**.** *Let $\mathcal{F}$ be any C^2 foliation other than the product foliation of $S^2 \times S^1$ by spheres. Then $T\mathcal{F}$ may be approximated by positive and negative contact structures.*

The proof is carried out in two steps: first perturb the foliation to an intermediate structure called a *confoliation*, and then perturb the confoliation further to a contact structure.

Here is the definition of a confoliation:

Definition 2.14. A 2-plane field ξ for which there is locally some 1-form α with $\ker(\alpha) = \xi$ and $\alpha \wedge d\alpha \ge 0$ (resp. ≤ 0) is a *positive (resp. negative) confoliation.*

A foliation may not be perturbed to a contact structure locally; if we think of α locally as a connection 1-form on a line bundle and $-d\alpha$ as its curvature, Stokes' Theorem prevents us creating some local positive contact structure without creating an 'equal amount' of negative contact structure. However: such a perturbation exists in a neighborhood of a loop with suitable holonomy.

Definition 2.15 (Contracting and weakly contracting)**.** Let $\gamma \subset \lambda$ be a loop in a leaf λ of $\mathcal{F}$. Say that γ has *contracting holonomy* if there is a transversal τ parameterized as $[-1, 1]$ (with $0 = \tau \cap \lambda$) so that holonomy transport around γ restricted to τ is conjugate to $f : [-1,1] \to [-1,1]$ with $x < f(x) < 0 < f(y) < y$ for all $x < 0 < y$, and it has *weakly contracting holonomy* if $x_i < f(x_i) < 0 < f(y_i) < y_i$ for some sequences x_i, y_i converging to 0 from either side.

Holonomy which exhibits the behaviour in Definition 2.15 is sometimes said to be *two-sided (weakly) contracting*, to distinguish it from holonomy which exhibits the given behaviour only on one side of λ.

Lemma 2.16 (Holonomy perturbs)**.** *Let $\mathcal{F}$ be a foliation of M, and let γ be an essential simple loop in a leaf λ of $\mathcal{F}$ with two-sided weakly contracting holonomy. Let N be a regular open solid torus neighborhood of γ. Then $\mathcal{F}$ may be (C^0) perturbed in any neighborhood of γ to a (positive or negative) confoliation which is contact near γ, and agrees with $\mathcal{F}$ near the boundary.*

Proof. First consider the case that γ has contracting holonomy $f : [-1,1] \to [-1,1]$. Note that any such f is topologically conjugate to any other. Let E be a transverse rectangle that cuts N into a parallelepiped P. The intersection with $\mathcal{F}$ induces a characteristic foliation $\mathcal{F}|\partial P$ on ∂P. Perturb $\mathcal{F}|\partial P$ to a new foliation $\mathcal{T}$ on $E^\pm$ by tilting it to the left on E^+ (and therefore on the right on E^-) so that it agrees in N (see Figure 3).

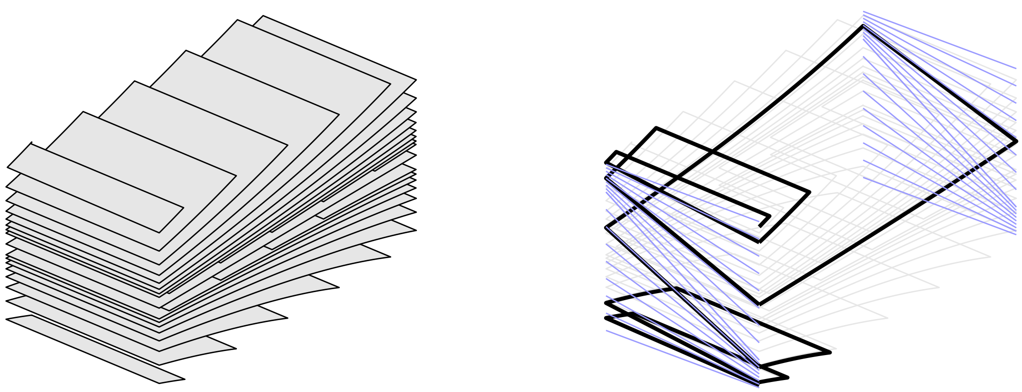

FIGURE 3. A foliated parallelepiped P. If we tilt the characteristic foliation on $E^\pm$ the integral curves spiral around ∂P.

The result gives a characteristic foliation on ∂P that spirals (positively) from a unique maximum to a unique minimum, and this characteristic foliation may be filled in over P with a positive contact structure that is tangent to $\mathcal{F}$ on $\partial P - E^\pm$ and to $\mathcal{T}$ on $E^\pm$ and therefore descends to a positive contact structure on N that may be extended by $\mathcal{F}$ on $M - N$ to give a positive confoliation on M.

Now suppose γ has weakly contracting holonomy conjugate to $h : [-1,1] \to [-1,1]$ on some transversal τ. Without loss of generality we may assume $x < h(x)$ for all $x \in [-1, x_0]$ and $y > h(y)$ for all $y \in [y_0, 1]$ for some specific $-1 < x_0 \le 0 \le y_0 < 1$. Let $g : [-1,1] \to [-1,1]$ be any homeomorphism that takes $[x_0, y_0]$ to $[-\epsilon, \epsilon]$ for some very small ϵ, and let $f' = ghg^{-1}$. Evidently we may choose f' to be ϵ-close (in the C^0 topology) to any fixed f as above, so after isotopy we may assume that the foliation $\mathcal{F}'$ in a neighborhood of γ is C^0 close to some $\mathcal{F}$ as above and equal to it outside the ϵ-neighborhood of λ. We perturb the characteristic foliation to the $\mathcal{T}$ on $E^\pm$ as for $\mathcal{F}$, and observe that the integral curves of this new foliation are ϵ-close to the integral curves as above and agree with them outside the ϵ-neighborhood of λ. Thus they still spiral from a unique maximum to a unique minimum so that they may be filled in with a positive contact structure over P. □

Example 2.17 (Linear holonomy). If $\mathcal{F}$ is co-oriented C^1, the derivative of holonomy determines a representation from $\pi_1(\lambda)$ (for any leaf λ) to the (multiplicative) group $\mathbb{R}^+$. A loop γ has nontrivial linear holonomy if the image under this representation is not equal to 1. Evidently, nontrivial linear holonomy is two-sided expanding or two-sided contracting depending on whether the image is greater or less than 1.

For sufficiently smooth foliations, nontrivial linear holonomy may be found along most minimal sets. Recall that in Chapter 4 we proved Sacksteder's Theorem that for a C^2 foliation, every exceptional minimal set has nontrivial linear holonomy; and if $\mathcal{F}$ itself is minimal, either there is nontrivial linear holonomy or $\mathcal{F}$ has no holonomy at all. Recall also that we showed that a C^2 minimal foliation with trivial holonomy is transversely measured, and may be perturbed to a surface bundle.

If ξ is a confoliation, there is an open submanifold of M where ξ is a positive contact structure, and a closed complement where ξ is tangent to a foliation. A positive confoliation is *saturated* if every point may be joined to a point in the contact region by a path tangent to ξ.

The following theorem is due to Altschuler [1]

Theorem 2.18 (Altschuler; Saturated splits). *Let ξ be a saturated confoliation. Then ξ may be perturbed to a C^∞ close contact structure.*

Proof. Altschuler considers a leafwise heat flow □

Using these ingredients we may now give the proof of Theorem 2.13.

Proof. Let $\mathcal{F}$ be a foliation of M, and let Λ be a minimal set. By hypothesis and the Reeb stability theorem there are no spherical leaves, so by perturbing maximal foliated I-bundles we may assume that there are finitely many minimal sets, and that they are all one of the following three kinds:

(1) an isolated closed leaf;
(2) all of M (i.e. $\mathcal{F}$ is already minimal); or
(3) an *exceptional minimal set* — one intersecting a transversal in a Cantor set.

Let S be an isolated closed leaf. Since S is isolated, there are simple essential closed loops $\gamma^\pm$ with weakly contracting holonomy on the positive resp. negative side. We may blow up S to $S \times I$, and insert a foliated bundle which is contracting near 1 along γ^+ and contracting near 0 along γ^- to produce a new C^0 close foliation in which S has been replaced by two isolated closed leaves, each of which has a loop with two-sided weakly contracting holonomy.

If $\mathcal{F}$ is minimal, then by Proposition **??** either some leaf of $\mathcal{F}$ has two-sided contracting holonomy, or else M is T^3 and $\mathcal{F}$ is (isotopic to) a linear foliation; in the latter case we may perturb $\mathcal{F}$ to a foliation by closed tori and then perturb again to reduce to the first case.

If Λ is an exceptional minimal set, then by Proposition **??** some curve has two-sided contracting holonomy.

In each case by Lemma 2.16 we may perturb $\mathcal{F}$ to a positive contact structure in an arbitrarily small neighborhood of such a curve, so that the perturbations may be made disjointly and simultaneously. The result is a saturated (positive) confoliation, which splits open to a positive contact structure by Theorem 2.18. □

Corollary 2.19. *Let M be a closed, oriented 3-manifold and let $\mathcal{F}$ be a co-oriented taut foliation. Then $M \times (-1,1)$ may be given the structure of a symplectic 4-manifold with pseudoconvex boundary.*

Proof. Let ω be a closed 2-form positive on $T\mathcal{F}$ and let $\alpha^\pm$ be positive and negative contact structures close to α with $\ker(\alpha) = T\mathcal{F}$. For small ϵ the form $\omega + \epsilon d(t\alpha)$ is symplectic on $M \times [-1,1]$, and is strictly positive on $\alpha^\pm$ which are both positive contact structures on the boundary components with the induced orientation. □

2.4. **Open book decompositions.**

Definition 2.20. An *open book decomposition* of an oriented 3-manifold M is an oriented 1-manifold X (the *binding*) and a fibration $\pi : M - X \to S^1$ whose fibers (the *leaves*) are open oriented surfaces that may be compactified to surfaces with boundary X.

Thus, M admits an open book decomposition with binding X whenever X is a fibered link (the fibration however may not be unique). We denote the data of an open book decomposition as (M, X, π).

The data of the open book may be recovered (up to isomorphism) from the mapping class $\phi \in \mathrm{Mod}(\Sigma, \partial\Sigma)$ where Σ is the fiber of π. We let $M(\Sigma, \phi)$ denote the isomorphism class of open book decomposition associated to this data

Definition 2.21. A positive co-oriented contact structure ξ on an oriented 3-manifold M is *supported by* an open book decomposition (M, X, π) if there is a (positive) contact form α for ξ with $\alpha|TX > 0$ for the oriented binding X and $d\alpha|T\Sigma > 0$ for the oriented fibers Σ of π.

Theorem 2.22 (Thurston–Winkelnkemper; open book to contact structure)**.** *Any open book decomposition (M, X, π) supports some positive co-oriented contact structure (M, ξ).*

Proof. A neighborhood of each component of X is a solid torus $D^2 \times S^1$. We put coordinates r, θ on D^2 and ϕ on S^1. Let $f : [0,1] \to [0,1]$ be a smooth increasing diffeomorphism, tangent to 0 to first order at the endpoints (i.e. a 'sigmoid function') and define $\alpha := (1 - f(r))d\phi + f(r)d\theta$. Then

$$\alpha \wedge d\alpha = f'(r) dr \wedge d\theta \wedge d\phi$$

Since $f'(r) > 0$ in $(0,1)$ and $f'(r)$ goes to zero to first order at 0 and 1, the form $\alpha \wedge d\alpha$ is strictly positive on the interior of the solid torus, and tapers off to 0 at the boundary where $\ker(\alpha)$ becomes tangent to the fibers of $\theta : (D^2 - 0) \times S^1 \to S^1$. Thus $\ker(\alpha)$ defines a positive contact structure on a neighborhood of the binding which extends to a positive confoliation on M by making it tangent to the fibers of π outside this neighborhood. This positive confoliation is saturated, and may be perturbed to a positive contact structure as in § 2.3. □

The converse of Theorem 2.22 is a theorem of Giroux:

Theorem 2.23 (Giroux; contact structure to open book)**.** *Let (M,ξ) be a positive co-oriented contact structure. Then there is an open book decomposition (M,X,π) supporting (M,ξ).*

First we give a definition:

Definition 2.24. Let (M,ξ) be a positive contact structure. A *contact cell decomposition* for (M,ξ) is a finite CW complex structure on M such that

(1) the 1-skeleton is Legendrian (i.e. tangent to ξ);
(2) every 2-cell D the push off of ∂D given by the framing ξ has intersection number -1 with D; and
(3) every 3-cell embeds in the standard contact structure on $\mathbb{R}^3$.

A contact cell decomposition for (M,ξ) always exists: if we choose a sufficiently small cellulation, we may embed the 3-cells in the standard contact $\mathbb{R}^3$, and then C^0 perturb the 1-skeleton to be Legendrian. Each 2-cell D may be thought of as living in the standard contact $\mathbb{R}^3$, and therefore the ξ-framed push off of ∂D intersects D with intersection number $\mathrm{tb}(\partial D) \le 1$. If $\mathrm{tb}(\partial D) < -1$ we may decompose D by Legendrian arcs into subcells D_i each with $\mathrm{tb}(\partial D_i) = -1$.

Now, let R be a ribbon surface with core the 1-skeleton, contained in a sufficiently small neighborhood so that ∂R is positively transverse to ξ. We claim that ∂R is the binding of an open book decomposition with R a fiber, and supporting (M,ξ). To see this, cut open M along R to obtain a sutured handlebody H whose boundary contains two copies $R^{\pm}$ of R meeting along ∂R. We claim this sutured handlebody is a product. To see this, observe that each 2-disk D gives a compressing disk for H whose boundary intersects the sutures twice. This proves the claim.

2.5. **Stabilization.**

Definition 2.25. Let Σ be a compact oriented surface with nonempty boundary, and let $\phi \in \mathrm{Mod}(\Sigma,\partial\Sigma)$ be a mapping class. Let Σ' be obtained from Σ by attaching a 1-handle, and let γ be an essential simple closed curve in Σ' intersecting the core of the 1-handle transversely once. We may extend ϕ by the identity to obtain a mapping class $\phi' \in \mathrm{Mod}(\Sigma',\partial\Sigma')$ and then define the positive (resp. negative) *stabilization* ϕ_γ along γ to be the composition $\tau_\gamma\phi'$ where τ_γ denotes a right handed (resp. left handed) Dehn twist.

If $M(\Sigma,\phi) = (M,X,\pi)$ and Σ,ϕ_γ is obtained from Σ,ϕ by positive stabilization, then $M(\Sigma',\phi_\gamma) = (M,Y,\pi')$ where Y is obtained from X by plumbing with a right handed Hopf band with core γ. Positive stabilization is compatible with positive contact structures

Lemma 2.26. *Suppose $M(\Sigma,\phi) = (M,X,\pi)$ and $M(\Sigma',\phi_\gamma) = (M,Y,\pi')$ where Σ,ϕ_γ is obtained from Σ,ϕ by positive stabilization. If a positive contact structure (M,ξ) is supported by (M,X,π) then it is also supported by (M,Y,π').*

Furthermore, one has the following theorem of Giroux:

Theorem 2.27 (Giroux; stabilization)**.** *If (M,ξ) is a positive contact structure, any two open book decompositions (M,X,π), (M,Y,π') supporting (M,ξ) become equivalent after sufficiently many positive stabilizations.*

2.6. **Right veering diffeomorphisms.** The results in this section are due to Honda–Kazez–Matić [15].

Definition 2.28. A nontrivial mapping class $\phi \in \mathrm{Mod}(\Sigma, \partial\Sigma)$ is *right veering* if for any $p \in \partial\Sigma$ and any essential embedded arc α with an endpoint at p, the geodesic representative of $\phi(\alpha)$ is to the right of (or is equal to) the geodesic representative of α near p, with respect to any hyperbolic metric on Σ. If C is a component of $\partial\Sigma$ then ϕ is right veering with respect to C if it is right veering as above for some (equivalently, any) $p \in C$.

Here is an equivalent definition of right veering. Choose a hyperbolic structure on Σ, and let $\tilde{\Sigma} \subset \mathbb{H}^2$ denote the universal cover, and let $\tilde{p} \in \partial\tilde{\Sigma}$ be any lift of p. The union of the limit set of $\tilde{\Sigma}$ with $\partial\tilde{\Sigma}$ is a topological circle S^1 compactifying $\tilde{\Sigma}$. Any representative of ϕ lifts to a unique equivariant diffeomorphism of $\tilde{\Sigma}$ fixing p, that extends continuously to a homeomorphism of S^1. Furthermore, this extension is independent of the choice of representative, and defines a faithful representation $\mathrm{Mod}(\Sigma, \partial\Sigma) \to \mathrm{Homeo}^+(S^1)$ fixing $\tilde{p}$. A mapping class ϕ is right veering with respect to the component C containing p if and only if for all $q \in S^1 - \tilde{p}$ either $\phi(q) = q$ or $p, \phi(q), q$ is positively ordered (in the cyclic order on S^1).

Evidently for each C the mapping classes right veering for C form a monoid $\mathrm{Veer}_C(\Sigma)$; the intersection of these monoids over all components C is therefore also a monoid $\mathrm{Veer}(\Sigma)$.

Example 2.29. A right handed Dehn twist (and therefore any product of right handed Dehn twists) is in Veer. It follows that no nontrivial product of right handed Dehn twists is ever trivial in $\mathrm{Mod}(\Sigma, \partial\Sigma)$; this is in contrast to the case of the mapping class group of a closed surface, in which every element is a product of right handed Dehn twists.

Example 2.30. Let Σ', ϕ' be obtained from Σ, ϕ by attaching a 1-handle and extending ϕ over the 1-handle as the identity. If $\phi \in \mathrm{Veer}(\Sigma)$ then $\phi' \in \mathrm{Veer}(\Sigma')$.

Here is the proof. Let $p \in \partial\Sigma'$ lie also in $\partial\Sigma$. The universal cover $\tilde{\Sigma}'$ is a tree of copies of $\tilde{\Sigma}$ plumbed together along arcs covering the 1-handle, and the action of ϕ' on $S^1(\Sigma')$ is obtained from the action of ϕ on $S^1(\Sigma)$ by recursively blowing up segments and inserting actions of ϕ on the complements of fixed components. Since all the inserted actions of ϕ are right veering, the same is recursively true of ϕ'.

It follows that if the monodromy of an open book is right veering, then so is the monodromy of any positive stabilization.

Theorem 2.31 (Honda–Kazez–Matić; tight implies right veering)**.** *Let $M(\Sigma, \phi)$ be an open book supporting a tight (positive) contact structure (M, ξ). Then the monodromy ϕ is right veering.*

2.7. **Closing symplectic manifolds.**

Theorem 2.32 (Eliashberg, Etnyre; pseudoconvex to closed)**.** *Let W, ω be a pseudoconvex symplectic 4-manifold. Then W is an open submanifold of a closed symplectic 4-manifold.*

3. J-holomorphic curves

3.1. **Holomorphic curves.** Let W^{2n}, ω be a symplectic manifold, and let J be a compatible almost complex structure with associated metric $g(X,Y) := \omega(X, JY)$.

Definition 3.1. Let S be a Riemann surface. A smooth map $u : S \to W$ is *holomorphic* if $du(iX) = Jdu(X)$ for all $X \in TS$.

The terms J-holomorphic and pseudo-holomorphic are also standard; we sometimes use the term J-holomorphic when we want to emphasize the dependence on the choice of J.

Lemma 3.2 (Calibration). *Let ξ be an oriented 2-plane in T_pW. Then $|\omega(\xi)| \le \text{area}_g(\xi)$ and $\omega(\xi) = \text{area}_g(\xi)$ if and only if ξ is J-invariant and oriented compatibly with the induced complex structure.*

Proof. Let $X, Y \in T_pW$ be orthonormal with respect to g. Then $\omega(X, Y) = -g(X, JY)$ so $|\omega(X, Y)| \le 1 = \|X \wedge Y\|$ and $\omega(X, Y) = 1$ if and only if $Y = JX$. □

Corollary 3.3. *Holomorphic curves are globally least area minimal surfaces in their homology class (among compactly supported variations).*

Proof. Let S be a holomorphic curve, and let S' be homologous to S. Then

$$\text{area}(S) = \int_S \omega = \int_{S'} \omega \le \text{area}(S')$$

□

3.2. **Cauchy–Riemann equations.** Under suitable circumstances, one may show that the space of J-holomorphic curves (with fixed domain, and in a suitable homology class) is a smooth finite dimensional manifold. The proof of this is a routine but technically involved exercise in the theory of elliptic PDE; for details we refer the reader to McDuff–Salamon [18], and content ourselves here with a summary.

For a compact Riemann surface S and a smooth map $u : S \to W$, write $\bar\partial_J u := (1/2)(du + J\,du\,i)$. The map $u : S \to W$ is holomorphic if and only if $\bar\partial_J u = 0$. We may choose Darboux local coordinates on W around $u(p)$ for some $p \in S$, symplectically isomorphic to an open subset of standard symplectic $\mathbb{R}^{2n}$, and after composing with a linear automorphism we may assume that J acts on $T_0\mathbb{R}^{2n}$ as J_0, sending $\partial/\partial x_j$ to $\partial/\partial y_j$ for all j. At p the equation $\bar\partial_J u = 0$ reduces to the ordinary Cauchy–Riemann equations; thus, near p, the equation $\bar\partial_J u = 0$ is uniformly elliptic.

Let $\mathcal{B}$ denote the space of smooth maps $u : S \to W$ in some fixed homology class $A \in H_2(W; \mathbb{Z})$, and let $\mathcal{M}(S, A, J)$ (or just $\mathcal{M}$) denote the subspace of $\mathcal{B}$ consisting of J-holomorphic maps. If $u : S \to W$ is any smooth map, and u_t is a smooth 1-parameter variation of u, then $\xi := u'(0)$ is a section of u^*TW. The pullback bundles u_t^*TW are not isomorphic; however, using the Levi-Civita connection ∇ on W associated to g we may trivialize this family and think of $\bar\partial_J u_t$ as a family of 1-forms on S with values in u^*TW. These forms are complex anti-linear, in the sense that they entwines multiplication by i and by $-J$; thus this family is a section of the complex bundle $\Lambda^{0,1}T^*S \otimes_J u^*TW$. Let $\mathcal{E}$ therefore denote the bundle over $\mathcal{B}$ whose fiber over u is the space of smooth sections of $\Lambda^{0,1}T^*S \otimes_J u^*TW$. The map $\bar\partial_J$ defines a section $\mathcal{B} \to \mathcal{E}$, and we would like to understand when $\bar\partial_J$ is transverse to the zero section, whose intersection is precisely $\mathcal{M}$. Via the Levi–Civita connection we obtain a splitting (locally) of the tangent space to $\mathcal{E}$ at $\bar\partial_J(u)$ as $T_u\mathcal{B} \oplus \mathcal{E}_u$.

The projection of $d\bar\partial_J$ to $\mathcal{E}_u$ with respect to this splitting is a linear operator D_u from sections of u^*TW to sections of $\Lambda^{0,1}T^*S \otimes_J u^*TW$; as a formula, one has for any vector field X on S,

$$(D_u\xi)(X) = (1/2)\left(\nabla_{du(X)}\xi + J\nabla_{du(X)}\xi i + (\nabla_\xi J)du(iX)\right)$$

Since $\nabla_\xi J$ is tensorial (i.e. 0th order) in ξ, it follows that $D_u\xi = \bar\partial_J\xi+$ terms 0th order in ξ. Thus D_u is elliptic, and one may show that it induces a Fredholm map between suitable Banach space completions of $T_u\mathcal{B}$ and $\mathcal{E}_u$.

Now, let $\mathcal{J}$ denote the space of ω-compatible almost complex structures. Each $J \in \mathcal{J}$ determines $\mathcal{M}(S, A, J)$ as above, and the union is a bundle $\pi : \mathcal{M}^\mathcal{J} \to \mathcal{J}$. One would like to argue that suitable completions of $\mathcal{M}^\mathcal{J}$ and $\mathcal{J}$ are smooth Banach manifolds, and therefore that a residual subset of the completion of $\mathcal{J}$ consists of regular values of π. This can be shown providing one has enough freedom to deform J independently at the image of different points in S; the key condition is that u should be injective on an open dense set. This motivates the following definition:

Definition 3.4. A holomorphic curve $u : S \to W$ is *simple* if it is nonconstant, and does not factor through a (holomorphic) branched cover of Riemann surfaces.

Remark 3.5. Simplicity is automatic under certain homological conditions.

(1) Since $\text{area}(u(S)) = \int_S u^*\omega$ for a holomorphic curve, $u \in \mathcal{M}(S, A, J)$ is constant if and only if $A = 0$.
(2) If $u : S \to W$ factors through $S \to S' \to W$ where $S \to S'$ is a branched cover of degree n, the homology class A satisfies $A = nA'$ where A' is the homology class of $S' \to W$.

It follows that if A is primitive and nonzero, every $u \in \mathcal{M}$ is simple.

Let $\mathcal{M}_s \subset \mathcal{M}$ denote the subspace of simple holomorphic curves. It is evidently open. Furthermore, any simple u is evidently injective on an open dense subset of S. We may define $\mathcal{M}_s^\mathcal{J}$ to be the union of the $\mathcal{M}_s$ over all $\mathcal{J}$. One then shows that $\mathcal{M}_s^\mathcal{J}$ and $\mathcal{J}$ admit completions with respect to which they are both smooth Banach manifolds.

Definition 3.6. An almost-complex structure J is *regular* if D_u is surjective for every simple $u \in \mathcal{M}_s$.

Proposition 3.7. *If J is regular, then $\mathcal{M}_s$ is a smooth almost-complex manifold of dimension equal to the index of D_u.*

Proof. An infinite dimensional implicit function (which depends on certain estimates that may be found e.g. in [18] § 3.3) imply that if J is regular, $\mathcal{M}_s$ is a smooth manifold locally modeled on the kernel of D_u. This kernel is not quite J-invariant (because of the first-order terms) but we may homotop D_u (through surjective Fredholm maps) to a nearby operator whose kernel is J-invariant, and therefore complex. This gives an almost-complex structure to $\mathcal{M}$. □

One may show that J is regular in this sense precisely when J is a regular value of $\pi : \mathcal{M}_s^\mathcal{J} \to \mathcal{J}$. By Sard–Smale the set of regular values of a map between separable Banach manifolds is residual. A priori these regular J are not smooth, but only lie in a suitable completion of the space $\mathcal{J}$. However, Taubes (see [18], § 3.4) shows one may find a residual subset of regular values in $\mathcal{J}$. Precisely, one has:

Theorem 3.8 (Taubes)**.** *For any W, ω, for any closed Riemann surface S and any homology class $A \in H_2(W;\mathbb{Z})$, there is a residual subset of the space of smooth ω-compatible J that are regular. For such a J the moduli space $\mathcal{M}_s(S, A, J)$ of simple J-holomorphic curves is a smooth oriented manifold of dimension equal to the index of D_u at any u.*

Furthermore, for any pair J_0, J_1 of J as above, there is a smooth path J_t interpolating between J_0 and J_1 for which the union of the $\mathcal{M}_s(S, A, J_t)$ is a smooth oriented cobordism between $\mathcal{M}_s(S, A, J_0)$ and $\mathcal{M}_s(S, A, J_1)$.

3.3. **Local structure.** Because the operator $\bar{\partial}_J$ agrees with the Cauchy–Riemann operator to leading order, we expect that the local structure of a J-holomorphic curve should resemble the local structure of an 'honest' holomorphic curve in an algebraic variety. This intuition is realized by the following theorem of Micallef and White [19], Thms. 6.1 and 6.2:

Theorem 3.9 (Micallef–White)**.** *Let $u : \Sigma \to M^{2n}$ be a non-constant J-holomorphic curve for some Riemann surface Σ (not necessarily connected). Then for any point $x \in u(\Sigma)$ and for all $p_j \in u^{-1}(x)$ there are neighborhoods $p_j \in U_j \subset \Sigma$ and $x \in V \subset M$ and coordinate charts $\psi : V \to \mathbb{C}^n$, $\phi_j : U_j \to \mathbb{C}$ with $\psi(x) = 0$ and $\phi_j(p_j) = 0$, and so that*

$$\psi u \phi_j^{-1}(z) = (z^{Q_j}, f_j(z))$$

where $f_j(z) \in \mathbb{C}^{n-1}$ vanishes to order $\geq Q_j \geq 1$ at $z = 0$.

Furthermore, if J is C^2, then ψ is C^1 and each ϕ_j is $C^{2,\alpha}$ for some positive α.

This implies (for instance) that multiple points x in the image of u are isolated unless u factors through a (branched) cover. It also lets us control the geometry of a singularity.

Theorem 3.9 is actually a corollary of a more general theorem about the local structure of singularities of minimal surfaces in Riemannian manifolds, and has nothing fundamentally to do with symplectic or almost-complex geometry per se. We refer the reader to [19] for the proof.

3.4. **Computation of the index.** If S is a Riemann surface and L is a holomorphic line bundle over S, the kernel of $\bar{\partial}$ is the space of holomorphic sections of L, and the cokernel is the space of holomorphic sections of $K \otimes L^*$ where K is the canonical line bundle (i.e. T^*S). The index (as a Fredholm map between *complex* Banach spaces) may therefore be computed from the Riemann–Roch formula as $\deg(L) + 1 - g$; the real index is twice this, i.e. $2\deg(L) + 2 - 2g$. The degree of a line bundle over a Riemann surface is $c_1(L)[S]$.

If E is a holomorphic vector bundle, the splitting principle says that the index of $\bar{\partial}$ may be computed as though E were a sum of line bundles. Thus the *real* index in this case is $\dim_{\mathbb{C}}(E)(2 - 2g) + 2c_1(E)[S]$.

If M, ω is a symplectic manifold of dimension $2n$ and $u : S \to M$ is a holomorphic curve (for some J), the pullback u^*TM is a complex vector bundle of complex dimension n, and the operator D_u is homotopic to $\bar{\partial}$ and therefore has the same index as above. Thus:

Lemma 3.10 (Index formula)**.** *Let M^{2n}, ω be symplectic, and let $u : S \to M$ be a holomorphic curve of genus g for some compatible J. The index of D_u (and therefore the dimension of $\mathcal{M}$ near u if u is simple and J is regular) is $n\chi(S) + 2c_1(M)u_*[S]$.*

3.5. **Curves with boundary.** Let W, ω be symplectic and let $L \subset W$ be a Lagrangian submanifold. Let S be a compact oriented surface with nonempty boundary. A smooth map $u : (S, \partial S) \to (W, L)$ is holomorphic if $\bar{\partial}_J u = 0$. If we fix a relative homology class $A \in H_2(W, L; \mathbb{Z})$ we may define $\mathcal{M}(S, A, J)$ exactly as before. The Fredholm theory goes through as in § 3.2 with essentially no modification. The index ...

3.6. **Gromov's compactness theorem.** For any Riemann surface S, for any homology class $A \in H_2(W; \mathbb{Z})$ and any J, the group $\mathrm{Aut}(S)$ acts on $\mathcal{M}(S, A, J)$ by precomposition. The quotient is denoted $\overline{\mathcal{M}}(S, A, J)$ or just $\overline{\mathcal{M}}$ for short. The action is not typically free, but is free when restricted to $\mathcal{M}_s$. The group $\mathrm{Aut}(S)$ has complex dimension 3 when S is a sphere, 1 when S is a torus, and it is discrete (and typically trivial of $\mathbb{Z}/2\mathbb{Z}$) when S has higher genus.

For $S = S^2$ the space $\mathcal{M}_s(S, A, J)$ is never compact unless it is empty, since the orbit of the automorphism group $\mathrm{Aut}(S^2)$ on any simple curve is proper and noncompact. But even after we quotient out by this action the spaces $\overline{\mathcal{M}}$ and $\overline{\mathcal{M}}_s$ can fail to be compact. Let's understand this in some simple examples.

Example 3.11 (Rational maps). A degree d rational map $f : \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ has a graph Γ_f which is a holomorphic curve in the Kähler manifold $\hat{\mathbb{C}} \times \hat{\mathbb{C}}$ representing the homology class $A := [\hat{\mathbb{C}} \times \text{point}] + d[\text{point} \times \hat{\mathbb{C}}]$.

For simplicity let's suppose $f(\infty) = 1$ so that f is determined by its divisor $\mathrm{div} f$ (supported in $\mathbb{C}$) which can be written in terms of zeros and poles as $Z - P$ where each of Z and P is an unordered list of d complex numbers. Repetitions are allowed, but Z and P are disjoint. Some of the noncompactness of the space of rational maps arises in families where elements of Z and P collide. Let $a \in Z$ and $b \in P$ so that f has a factor of the form $(z-a)/(z-b)$. When $|a - b| = \epsilon$ is very small and $|z - a| > \sqrt{\epsilon}$, this factor differs from 1 by $\sim \sqrt{\epsilon}$. If we deform f in a family $f(t)$ by adjusting a so that $a(t) \to b$, the graphs $\Gamma_{f(t)}$ converge in the Hausdorff topology to a nodal curve $\Gamma_g \cup b \times \hat{\mathbb{C}}$, the union of the holomorphic graph Γ_g (where g of degree $(d-1)$ is obtained from f by canceling the $(z-a)/(z-b)$ factor) and the 'vertical' curve $b \times \hat{\mathbb{C}}$. One calls this phenomenon *bubbling off*, where $b \times \hat{\mathbb{C}}$ is the 'bubble'.

Here is another way to look at the local picture of the degeneration $\Gamma_{f(t)} \to \Gamma_g \cup b \times \hat{\mathbb{C}}$ near the singular point $(b, g(b))$. For $\epsilon \in \mathbb{C}$ let S_ϵ be the curve $z_1 z_2 = \epsilon$ in $\mathbb{C}^2$. When $\epsilon \neq 0$ the curve S_ϵ is an annulus, with a holomorphic isomorphism to $\mathbb{C} - 0$ given by projection to either the z_1 or the z_2 axis. As $\epsilon \to 0$ the S_ϵ degenerate to S_0 which (ignoring the embedding) is homeomorphic to the quotient space obtained from an annulus by pinching a meridian curve to a point.

It will turn out that the space of (unparameterized) holomorphic curves in a fixed homology class can be compactified by adding such nodal curves (more commonly called *cusp curves* in the symplectic topology literature). Here we may allow curves of any (fixed) genus, so throughout this section we do not assume that the genus is zero.

If S is a smooth surface and γ_i is a collection of disjoint simple loops, we obtain a singular surface $\bar{S}$ by collapsing each γ_i to a point p_i. Let $\hat{S}$ be the end completion of $\bar{S} - \cup p_i = S - \cup \gamma_i$. Then $\hat{S}$ is a closed surface, and there is a canonical map $\hat{S} \to \bar{S}$ which is the identity on $\bar{S} - \cup p_i$, and is 2–1 on the preimage of each p_i. A complex structure on $\hat{S}$ makes $\bar{S}$ into a *nodal Riemann surface*. Equivalently, a complex structure on $\bar{S}$ is just a complex structure on $S - \cup\gamma_i$ for which the modulus of every annular end is infinite.

If W is a smooth manifold with almost-complex structure J and if $\bar{S}$ is a nodal Riemann surface, then $\bar{u} : \bar{S} \to W$ is (J-)holomorphic if it is holomorphic on $\bar{S} - \cup p_i$ in the usual sense. Such holomorphic curves are called *cusp curves*.

Definition 3.12. A sequence of holomorphic curves $u_n : S_n \to W$ is said to *converge weakly* to a cusp curve $\bar{u} : \bar{S} \to W$ if

(1) the areas of $u_n(S_n)$ converge to the area of $\bar{u}(\bar{S})$; and
(2) there are families of disjoint loops $\gamma_{n,i} \subset S_n$ and diffeomorphisms $\phi_n : \bar{S} - \cup p_i \to S_n - \cup\gamma_{n,i}$ so that the maps $u_n\phi_n$ converge uniformly to $\bar{u}$ on compact subsets of $\bar{S} - \cup p_i$.

We may now state Gromov's Compactness Theorem ([13], Thm. 1.5.B):

Theorem 3.13 (Compactness). *Let W, ω be a smooth closed symplectic manifold with a compatible almost-complex structure J (and associated metric g). Then any sequence of holomorphic curves of uniformly bounded genus and uniformly bounded area has a subsequence which converges weakly to a cusp curve.*

The proof of this theorem shall take up the rest of this subsection and the next. First we observe that there is a uniform upper bound on the curvature of a holomorphic curve.

Lemma 3.14. *For any holomorphic curve S the sectional curvature K_S satisfies $K_S \le K$ pointwise where K is the maximum of the sectional curvature of W.*

Proof. Holomorphic curves are minimal for a compatible metric. Gauss's equation for a minimal surface says $K_S = K_W - |\mathrm{I\!I}|^2/2$ where K_W is the sectional curvature of W along the tangent plane to S, and $\mathrm{I\!I}$ is the second fundamental form. □

Surfaces with an upper curvature bound satisfy an isoperimetric inequality that can be expressed in terms of the geometry of *comparison surfaces*.

Definition 3.15 (Comparison Surface). Let S be a Riemannian surface. A rotationally symmetric Riemannian surface S_0 is a *comparison surface* for S if, for any domain $E \subset S$, if $E_0 \subset S_0$ is a rotationally symmetric disk with the same area, then $\text{length}(\partial E) \ge \text{length}(\partial E_0)$.

Example 3.16. If S has curvature bounded above by K and area bounded above by $4\pi/K$, the sphere S^2_K of constant curvature K is a comparison surface for S.

Lemma 3.17. *Let S be a holomorphic curve, and Let S_0 be a comparison surface for S. Let $u : \mathbb{D} \to S$ be an injective holomorphic map with area A and let $u_0 : \mathbb{D} \to S_0$ be a* comparison map *— i.e. a conformal map onto a symmetric disk with the same area. Then $|du(0)| \le |du_0(0)|$.*

Proof. For any $r \le 1$ let $\mathbb{D}_r \subset \mathbb{D}$ be the disk of radius r. Let $A(r) := \text{area}(u(\mathbb{D}_r))$ and $A_0(r) := \text{area}(u_0(\mathbb{D}_r))$ and likewise $L(r) := \text{length}(u(\partial\mathbb{D}_r))$ and $L_0(r) := \text{length}(u_0(\partial\mathbb{D}_r))$.

By the Cauchy–Schwarz inequality,

$$A'(r) = \int_{\partial \mathbb{D}_r} |du|^2 \ge \frac{1}{2\pi r}\left(\int_{\partial \mathbb{D}_r} |du|\right)^2 = \frac{L(r)^2}{2\pi r}$$

whereas for u_0 we have (by the same calculation) equality: $A_0'(r) = L_0^2(r)/2\pi r$.

Let $\alpha : [0,1] \to [0,1]$ be the function for which $A(r) = A_0(\alpha(r))$ so that $A'(r) = \alpha'(r)A_0'(\alpha(r))$ for any r. Then the isoperimetric inequality gives $L(r) \ge L_0(\alpha(r))$ so

$$A'(r) \ge \frac{L_0^2(\alpha(r))}{2\pi r} = \frac{\alpha(r)}{r}\frac{L_0^2(\alpha(r))}{2\pi\alpha(r)} = \frac{\alpha(r)}{r}A_0'(\alpha(r)) = \frac{\alpha(r)}{r\alpha'(r)}A'(r)$$

so we get $r\alpha'(r) \ge \alpha(r)$ and by integrating $\log(r) - \log(\epsilon) \le \log(\alpha(r)) - \log(\alpha(\epsilon))$ for any positive ϵ. But $\alpha(1) = 1$ so $\alpha(\epsilon) \le \epsilon$ for any $\epsilon \le 1$ and therefore $\alpha'(0) \le 1$. Since $|du(0)| = \alpha'(0)|du_0(0)|$ the lemma follows. □

This Lemma and the isoperimetric inequality in Example 3.16 lets us control the derivative at 0 of a holomorphic disk with area less than $4\pi/K$. Remarkably, it is possible to obtain similar control only from a *diameter* bound. This is the so-called Gromov–Schwarz Lemma:

Theorem 3.18 (Gromov–Schwarz Lemma)**.** *There exist constants $\epsilon > 0$ and $C > 0$ so that any holomorphic map $u : \mathbb{D} \to W$ whose image is contained some ball $B_\epsilon(p)$ of radius ϵ is C-Lipschitz with respect to the hyperbolic metric on $\mathbb{D}$ and the g-metric on W.*

Proof. By compactness of W, there are positive constants $C_1 > 0$ and $\epsilon > 0$ so that on any ball of radius ϵ we can find a 1-form β with $\|\beta\| \le C_1$ for which $d\beta = \omega$. For any $u : \mathbb{D} \to B_\epsilon(p)$ we have

$$\text{area}(u(\mathbb{D})) = \int_{u(\mathbb{D})} \omega = \int_{u(\partial\mathbb{D})} \beta \le C_1 \text{length}(u(\partial\mathbb{D}))$$

It follows that there is a comparison surface S_0 with linear isoperimetric profile; i.e. with the property that a symmetric subdisk E_0 of S_0 satisfies $\text{area}(E_0) = C_1\text{length}(\partial E_0)$. Outside a compact subset, the metric on S_0 has curvature arbitrarily close to $-1/C_1^2$. The surface S_0 typically has a cone point at the origin, but this may be smoothed by inserting a tiny bubble of positive curvature K (as in Example 3.16) to obtain a smooth surface, which is still a comparison surface, and for which the (conformal!) uniformizing map $S_0 \to \mathbb{D}$ is bilipschitz.

By Lemma 3.17 a comparison map $u_0 : \mathbb{D} \to S_0$ has $|du(0)| \le |du_0(0)|$. Composing with the bilipschitz uniformizing map $S_0 \to \mathbb{D}$ and applying the (usual) Schwarz Lemma we get a uniform bound on $|du(0)|$ independent of u. □

3.7. **Completion of the proof.** The Gromov–Schwarz Lemma and a bootstrap argument allows us to promote C^0 convergence of holomorphic curves to C^∞ convergence (this is Lemma 3.19). The bootstrap argument treats the (higher) jets of a holomorphic curve as holomorphic curves in their own right in the space of *bundle maps*.

This requires a few words of explanation. If we fix S and W, there is a complex vector bundle E over $S\times W$ with fiber over a point (p,q) equal to $\text{Hom}_{\mathbb{C}}(T_pS, T_qW)$. The (almost)-complex structures on S and W determine an almost complex structure J_E on E, and if

$u : S \to W$ is a holomorphic curve in W then evidently the graph of (u, du) (i.e. the 1-jet of u) is a holomorphic curve in E with respect to J_E. The symplectic structure on W and an area form on S together give a symplectic structure on $S \times W$; if we pull this back to E and add a fiberwise symplectic form compatible with the fiberwise complex structure we obtain a symplectic structure on E compatible with J_E. For details see e.g. [16], Chapter 3.

Lemma 3.19. *If $u_n : S \to W$ is a sequence of holomorphic maps converging in the C^0 topology to $u : S \to W$ then in fact the u_n converge in the C^∞ topology and the limit is a holomorphic map.*

Proof. Let $p \in S$ and let D be a neighborhood of p in S for which $u_n(D)$ is contained in the ϵ-neighborhood of $u(p)$. Then the u_n are uniformly Lipschitz on D in the hyperbolic metric. Thus the u_n take a sufficiently small neighborhood of the zero section in TD to a small neighborhood of the zero section in $TB_\epsilon(u(p))$.

Thus the 1-jets of u_n have relatively compact image in the space of bundle maps. Applying the Gromov–Schwarz Lemma to the 1-jets of the u_n gives control on the norms of the 2-jets, and by induction on all higher derivatives. □

It is also important to be able to extend holomorphic curves over punctures:

Lemma 3.20. *Let $u : S - p \to W$ be holomorphic with relatively compact image. If either*

(1) $\text{area}(u(S - p))$ is finite; or
(2) $u(S - p)$ is contained in $B_\epsilon(q)$ for some $q \in W$

then u extends to a holomorphic map $\bar{u} : S \to W$.

Proof. By restricting to a subset of S if necessary we may assume $S = \mathbb{D}$ and $p = 0$. In the hyperbolic metric on $\mathbb{D} - 0$ a neighborhood of 0 has arbitrarily small area, so by Gromov–Schwarz the second case reduces to the first.

As in the proof of Lemma 3.17 we use the notation $A(r) := \text{area}(u(\mathbb{D}_r - 0))$ and $L(r) := \text{length}(u(\partial \mathbb{D}_r))$. Since $\text{area}(u(\mathbb{D} - 0))$ is finite and $u : S \to W$ is conformal it follows exactly as in Lemma 3.17 that

$$\infty > A(r) \geq \int_0^r \frac{L(s)^2}{2\pi s} ds$$

and therefore there is a sequence of radii $r_j \to 0$ with $L(r_j) \to 0$. By relative compactness of the image in W we may extract a subsequence of radii so that $u(\partial \mathbb{D}_{r_j}) \to q$. Suppose that there is a sequence of points $w_j \in \mathbb{D}_{r_j} - \mathbb{D}_{r_{j+1}}$ with a subsequence converging to $q' \neq q$ and let $\epsilon = d(q', q)/2$. Then for each index j in the subsequence, $u(\mathbb{D}_{r_j} - \mathbb{D}_{r_{j+1}})$ contains a point $u(w_j)$ arbitrarily close to q', but its boundary lies outside $B_\epsilon(q')$. By the monotonicity formula for minimal surfaces there is a uniform positive lower bound on the area of each $u(\mathbb{D}_{r_j} - \mathbb{D}_{r_{j+1}})$. But this implies $\text{area}(u(\mathbb{D} - 0))$ is infinite, contrary to assumption.

Thus u extends to a continuous map $\bar{u}$, and by Gromov–Schwarz u is uniformly Lipschitz in the hyperbolic metric. Thus the map on 1-jets has finite area and relatively compact image on $\mathbb{D}_r - 0$ and by induction $\bar{u}$ has continuous partial derivatives of all orders and is therefore holomorphic. □

We are now ready to conclude the proof of Theorem 3.13.

Proof. Fix small constants $A < 2\pi/K$ and $\epsilon > 0$ so that every minimal surface F in W intersects every ball $B_\epsilon(p)$ with $p \in F$ in a subsurface of area at least A. If $u : S \to W$ is holomorphic we can find a maximal subset of points $Q \subset S$ so that the ϵ-balls about the points of $u(Q)$ are disjoint (and therefore also the 2ϵ-balls about $u(Q)$ cover $u(S)$). Then $|Q| \le \text{area}(u(S))/A$ and (if we take ϵ small enough) also $|Q| \ge 3$. Thus under the assumptions of the theorem $S - Q$ is hyperbolic of uniformly bounded area and by Gromov–Schwarz, the norm of du in the hyperbolic metric is uniformly bounded on $S-Q$, independent of u.

If $u_n : S_n \to W$ is a sequence of holomorphic curves, and $Q_n \subset S_n$ points as above, then either the hyperbolic metrics on $S_n - Q_n$ have a convergent subsequence in some moduli space, or there is a subsequence for which these metrics degenerate by stretching necks centered at boundedly many essential simple closed loops $\cup\gamma_{n,i}$. Then $u_n : S_n - Q_n - \cup\gamma_{n,i}$ converge on some subsequence to $u : \bar{S} - Q - P$ for some nodal Riemann surface $\bar{S}$ and for some finite collections of points Q, P. By Lemma 3.20 this map extends to a cusp curve $\bar{u} : \bar{S} \to W$, and by the equicontinuity of the u_n in the hyperbolic metrics on $S_n - Q_n$ the areas of the $u_n(S_n)$ converge to the area of $\bar{u}(\bar{S})$. □

This completes the proof of the Compactness theorem.

One immediate application is as follows.

Corollary 3.21. *Let W, ω be symplectic, and let $A \in H_2(W)$ be a spherical class (i.e. a class in the image of $\pi_2(W)$). Suppose there is no spherical class B with $0 < \omega(B) < \omega(A)$. Then for any ω-compatible almost-complex structure J, the unparameterized moduli space $\overline{\mathcal{M}}(S^2, A, J)$ is compact.*

Proof. Any cusp curve $\bar{u} : \bar{S} \to W$ compactifying $\overline{\mathcal{M}}(S^2, A, J)$ has domain a nodal Riemann surface $\bar{S}$ whose irreducible components S_j (of which there are at least 2) are 2-spheres. But $\omega(A) = \text{area}(\bar{u}(\bar{S})) = \sum_j \text{area}(\bar{u}(\bar{S}_j)) = \sum_j \omega(A_j)$ where A_j is the (spherical) homology class represented by $\bar{u} : S_j \to W$. Since every $\omega(A_j) > 0$ we violate the hypothesis. □

3.8. **Dimension 4.** Holomorphic curves in dimension 4 satisfy extra rigidity properties arising from properties of the intersection form on homology. A compact symplectic 4-manifold W^4 is oriented, and there is a symmetric nondegenerate unimodular intersection form on $H_2(W)$ Poincaré dual to the cup product on $H^2(W)$. If $A, B \in H_2(X)$ are represented by smooth oriented surfaces S_A, S_B in general position, then $A \cdot B$ is equal to the signed count of intersections of S_A with S_B.

Holomorphic curves are typically neither nonsingular nor in general position with respect to each other, and therefore we must invoke Theorem 3.9 which describes the local structure of (possibly singular) holomorphic curves and their (possibly self-) intersections. Let's make the following definitions:

Definition 3.22 (Local intersection number)**.** Let $u : \Sigma \to W^4$ be holomorphic (not necessarily connected) with $p_j \in u^{-1}(x)$ for $j = 1, 2$. Write $\psi u \phi_j^{-1}(z) = (z^{Q_j}, f_j(z))$ as in Theorem 3.9. Let Q be the least common multiple of Q_1, Q_2 and write $Q = m_j Q_j$. Then we may define the *local intersection number*

$$\delta(p_1, p_2) := \frac{1}{m_1 m_2} \sum_{\nu^Q = 1} \text{ord}_0 \left(f_1(\nu z^{m_1}) - f_2(z^{m_2}) \right)$$

where ord_0 means order of vanishing at 0.

Likewise we may define:

Definition 3.23 (Local degree)**.** With notation as above let x be an (isolated) singular point of u, and $p \in u^{-1}(x)$ a preimage. Write $\psi u \phi^{-1}(z) = (z^Q, f(z))$ for some $Q > 1$ where $f(z)$ vanishes at 0 to order $Q' \geq Q$. Define the *local degree*

$$\delta(p) := \sum_{\nu^Q = 1, \nu \neq 1} \mathrm{ord}_0 \left(\frac{f(\nu z) - f(z)}{z} \right)$$

Example 3.24. For a single curve $u : \Sigma \to W$ we may define

$$\delta(u) = \sum_{(p_1, p_2)} \delta(p_1, p_2) + \sum_p \delta(p)$$

where the sum is taken over all singularities and multiple points. The quantity $\delta(u)$ measures the difference $\chi(u(\Sigma)) - \chi(\Sigma')$ where Σ' is a suitable 'desingularization' of $u(\Sigma)$.

This is easiest to explain in the integrable case. Consider the (projective) elliptic curve given in an affine chart by the equation $y^2 = x(x-a)(x-b)$. When $0, a, b$ are distinct, this elliptic curve is a nonsingular torus and has $\chi = 0$. The curve $y^2 = x^3 + x^2$ is genus zero and has a transverse double point at 0. Thus $\delta(u) = 1$ and $\chi = 1$. The curve $y^2 = x^3$ is genus zero and embedded (though not smoothly at the cusp point 0) so $\delta(u) = 0$ and $\chi = 0$.

In every case $\delta(p_1, p_2)$ is a positive integer $\geq Q_1 Q_2$. Likewise $\delta(p)$ is always $\geq (Q-1)(Q'-1)$. In fact, since $\delta(p)$ counts the contribution to χ from a change of genus, it is always *even*.

Using these definitions we may give formulae for the intersection product in terms of geometry for holomorphic curves.

Lemma 3.25 (Intersection Formula)**.** *Let $u_j : \Sigma_j \to W$ be simple holomorphic curves in homology classes $A_1, A_2 \in H_2(X)$. Then either the u_j both have the same image, or*

$$A_1 \cdot A_2 = \sum_{(p_1, p_2)} \delta(p_1, p_2)$$

In particular, $A_1 \cdot A_2 \geq 0$ with equality if and only if the curves are disjoint.

Proof. If the curves are nonsingular and the intersections are transverse, this is equivalent to saying that every intersection is positive. This is because the tangent spaces are complex subspaces of TW. The general case follows from Theorem 3.9. ∎

Lemma 3.26 (Adjunction Formula)**.** *Let $u : \Sigma \to W^4$ be a simple holomorphic curve in the homology class $A \in H_2(W)$. Then*

$$c_1(W)(A) = \chi(\Sigma) + A \cdot A - 2 \sum_{(p_1, p_2)} \delta(p_1, p_2) - 2 \sum_p \delta(p)$$

Proof. For simplicity we assume the singularities of u are transverse double points. We have $u^* TW = T\Sigma \oplus \nu$ where ν is the normal bundle, so $c_1(u^* TW) = c_1(T\Sigma) + c_1(\nu)$ in $H^2(\Sigma)$. Thus $c_1(W)(A) = \chi(\Sigma) + c_1(\nu)[\Sigma]$.

Now, $c_1(\nu)[\Sigma]$ is the number of intersections of Σ with a push-off of itself in $u^*\nu$, whereas A is the number of intersections of $u(\Sigma)$ with a push-off of itself; the difference between these numbers is twice the number of double points.

The general case follows from Theorem 3.9. □

For brevity in the sequel we denote the contributions δ from all singularities of u by $\delta(u)$.

As corollaries of the Adjunction Formula we obtain:

Corollary 3.27 (Dimension count)**.** *Let $A\cdot A \le -2$. Then for regular J the space $\mathcal{M}(A,J)$ is empty. Likewise, if $A\cdot A = -1$, for regular J the space $\overline{\mathcal{M}}(A,J)$ is 0-dimensional and its points correspond to embedded smooth curves.*

Proof. By the Adjunction Formula, the dimension of $\mathcal{M}(A,J)$ for regular J is

$$\dim \mathcal{M}(A,J) = 4 + 2c_1(X)(A) = 8 + 2A\cdot A - 4\delta(u)$$

where $\delta(u) > 0$ unless u is embedded and nonsingular. Since $\dim \overline{\mathcal{M}} = \dim \mathcal{M} - 6$ the proof follows. □

Corollary 3.28 (Smooth embeddedness persists)**.** *If u, u' are J, J' holomorphic spheres in the same homology class A, then $\delta(u) = \delta(u')$. In particular*

(1) if for some J there is some u which is a smooth embedding then for every J' every u' is a smooth embedding;
(2) if for some J there is some u which is smooth with exactly one transverse double point then for every J' every u is smooth with exactly one transverse double point.

Proof. By the Adjunction Formula $\delta(u) = \delta(u')$ because the other terms depend only on the homology class A. □

Another phenomenon special to dimension 4 is that one can show $\mathcal{M}(A,J)$ is a smooth manifold near a nonsingular embedded curve u under purely homological conditions on A:

Lemma 3.29 (Automatic Regularity)**.** *Let $u : S^2 \to W^4$ be a smoothly immersed simple holomorphic curve. Then D_u is onto (so that $\mathcal{M}$ is a smooth oriented manifold of the correct dimension near u) if and only if $c_1(W^4)(A) \ge 1$.*

Proof. First let's suppose that J is integrable. Then u^*WX is a holomorphic $\mathbb{C}^2$ bundle over S^2, which splits (because u is nonsingular) as a sum of holomorphic line bundles $TS^2 \oplus \nu$. The cokernel of D_u is therefore isomorphic to the space of holomorphic sections of $\nu^* \otimes K$. But $c_1(W^4)(A) = c_1(TS^2) + c_1(\nu)$ so $c_1(\nu) \ge -1$ so $c_1(\nu^* \otimes K) \le -1$ and therefore D_u is onto.

If J is not integrable there is an argument due to Hofer–Lizan–Sikorav [14] that we summarize. Showing that D_u is onto is equivalent to showing that the adjoint D_u^* has trivial kernel. Ignoring the tangential part of u^*TW, we may write $D_u^* = \bar\partial + a$ for some $a \in \Omega^{0,1}(\mathrm{End}_{\mathbb{R}}\nu^* \otimes K)$ and we want to show that if $\bar\partial f + af = 0$ then (under the homological condition on c_1) we have $f = 0$.

The operator D_u^* is not necessarily complex linear because a isn't. Thus the first step is to replace a by some $b \in \Omega^{0,1}$ so that $Lf = 0$ where $L = \bar\partial + b$. This is elementary: for $z \in S^2$ and $v \in T_zS^2$ we may just take $b(z)v = (a(z)v)f(z)/f(z)$ where $f(z) \ne 0$ and $b(z)v = 0$ where $f(z) = 0$. Now the operator L is complex-linear, and therefore defines a complex connection on $\nu^* \otimes K$. Any complex connection on a bundle over a Riemann surface is integrable, so L defines a new *holomorphic* structure on $\nu^* \otimes K$. Anything in the kernel of L would be a holomorphic section of this bundle, which must vanish because the degree is negative. Thus the kernel of L is trivial and so therefore is the kernel of D_u^*. More work is needed to apply this argument in the weaker regularity in the setting of § **??**. □

4. Floer Homology

4.1. **Morse Theory.** Let M be a smooth compact n-manifold. A smooth function $f : M \to \mathbb{R}$ is *Morse* if the critical points (i.e. points p where $df(p)$ is the zero map from T_pM to $T_{f(p)}\mathbb{R} = \mathbb{R}$) are nondegenerate. This means there is a neighborhood U of p and smooth local coordinates x_i on U vanishing at p such that throughout U,

$$f(x) = f(p) - x_1^2 - \cdots - x_i^2 + x_{i+1}^2 + \cdots + x_n^2$$

for some i, the *index* of the critical point p.

A Morse function is *self-indexing* if $f(p) = i$ for every critical point of index i. Local minima resp. maxima of f are critical points of index 0 resp. n so for a self-indexing Morse function, $f(M) = [0, n]$. It is a fact that every smooth compact manifold admits (many) self-indexing Morse functions.

Let f be such a function. For each $t \in \mathbb{R}$ define $M_t := f^{-1}(-\infty, t]$. Then M_t is empty for $t < 0$ and is equal to M for $t > n$. Furthermore, for any t not equal to $0, 1, \cdots, n$, the space M_t is a smooth manifold with boundary $\partial M_t := f^{-1}(t)$.

Choose a smooth metric g on M. The vector field $\mathrm{grad}(f)$ is defined by $\langle \mathrm{grad}(f), Y\rangle = df(Y)$ for all vectors Y. It is perpendicular to the level sets of f and has magnitude at every point equal to that of df. Thus, $\mathrm{grad}(f)$ vanishes only at the critical points, and the flow it generates gives a diffeomorphism from M_s to M_t whenever $i < s < t < t + 1$.

When we transition from M_s to M_t for $s < i < t$ the topology changes in a neighborhood of each critical point by attaching an i-handle $D^i \times D^{n-i}$. An i-handle is thought of as a thickened neighborhood of its core $D^i \times 0$, which is attached to ∂M_s along its boundary $\partial D^i \times 0$. The cocore is $0 \times D^{n-i}$ with boundary $0 \times \partial D^{n-i}$. The core of the i-handle associated to a critical point p is the closure of the set of flowlines of $\mathrm{grad}(f)$ (in a neighborhood of p) asymptotic to p in the future, while the cocore is the closure of the set of flowlines asymptotic to p in the past.

Thus each M_i has the homotopy type of an i-dimensional CW complex, which is the closure of the union of the flowlines of $\mathrm{grad}(f)$ asymptotic in the future to some critical point of index $\le i$. For each pair of critical points p, q of indexes $i > j$ let $F(p, q)$ denote the space of flowlines of $\mathrm{grad}(f)$ whose closures run from q to p. For generic f, the space $F(p, q)$ is an open oriented manifold of dimension $i - j - 1$. It may be compactified to a compact oriented manifold with corners $\overline{F}(p, q)$ by adding products $F(p, r_1) \times F(r_1, r_2) \times \cdots \times F(r_k, q)$ for intermediate critical points $r_1, \cdots, r_k$.

When $i - j = 1$ the space $F(p, q)$ is a finite set of points, and when $i - j = 2$ the space $F(p, q)$ is a finite union of circles, and open intervals that may be compactified by points of the form $F(p, r) \times F(r, q)$ where the index of r is $i - 1 = j + 1$.

The manifolds $F(p, q)$ may be oriented by thinking of them as intersections of the (oriented) manifolds of *all* flowlines asymptotic to p (resp. q) in the future (resp. past); thus for $i-j=1$ the set $F(p,q)$ is a finite set of *signed* points, and counting with sign gives an integer $n(p,q)$.

We may define a graded chain complex generated in dimension i by critical points of index i, and differential

$$\partial p := \sum_{\mathrm{index}(q)=i-1} n(p,q)q$$

Thus

$$\partial\partial p = \sum_q \sum_r n(p,r)n(r,q)q$$

But for each q, this sum $\sum n(p,r)n(r,q)$ is equal to the number of boundary points of the compact 1-manifold $\overline{F}(p,q)$, counted with sign. Thus $\partial\partial = 0$ and ∂ is the differential of a chain complex, which by construction has homology isomorphic to $H_*(M;\mathbb{Z})$.

4.2. **Floer Homology.** Now let (W^{2n},ω) denote a symplectic manifold with Lagrangian submanifolds L_0, L_1 and for simplicity let's suppose they intersect in general position. Let Ω denote the space of smooth maps $z: I \to M$ with $z(0) \in L_0$ and $z(1) \in L_1$.

The idea of Floer Homology is to define a smooth function a on Ω and compute the "Morse homology" associated to the critical points of a and the gradient flowlines joining them. The critical points will be the intersections $L_0 \cap L_1$ and the gradient flowlines will be holomorphic bigons with edges on $L_0 \cup L_1$ running between two intersection points.

If $\omega = d\lambda$ for some 1-form λ we could define a to be the action $a(z) := \int_z -\lambda$ (compare Example **??**). In fact to do Morse theory we do not really need the function a as such; rather we need its derivative da. If we choose a basepoint z_0 and a sufficiently small neighborhood $z_0 \in U \subset \Omega$ then for any other $z \in U$ we can join z_0 to z by a path z_t in Ω which sweeps out a rectangle $Z: I \times I \to W$ with left and right edges on L_0 and L_1 respectively. We may then define $a(z) := \int_Z \omega$. If y_t were another path sweeping out another rectangle Y then we could sew Z and Y together to make an annulus with boundary curves on L_0, L_1. If these curves were sufficiently small, we could cap them off with small disks in the L_j to make a (null-homotopic) sphere S, and then

$$\int_Z \omega - \int_Y \omega = \int_S \omega = 0$$

The global indeterminacy of a comes from the periods of ω on cylinders whose boundaries are loops in L_1 and L_2. Such cylinders (homotopically) are determined by intersections of conjugacy classes $\pi_1(L_0) \cap g\pi_1(L_1)g^{-1}$ together with the action of $\pi_2(W)$. Thus (for example) if $H_1(L_0) \cap H_1(L_1) = 0$ and $\pi_2(W) = 0$ then a is globally defined.

In any case da is well-defined. The tangent space $T_z\Omega$ is the space of vector fields ξ along z with $\xi(z(j)) \in T_(z(j))L_j$ and

$$da(\xi) = \int_0^1 \omega(z'(t),\xi)dt$$

Thus the critical points of a are precisely the constant maps; i.e. the points of $L_0 \cap L_1$.

A compatible almost-complex structure J determines a metric on Ω by

$$\langle \xi_1, \xi_2 \rangle := \int_0^1 \omega(\xi_1, J\xi_2)$$

and therefore

$$da(\xi) = \int_0^1 \omega(z'(t), \xi) = \int_0^1 \omega(Jz'(t), J\xi) = \langle Jz', \xi \rangle$$

In other words (at least formally) $\text{grad}(a) = Jz'$ and flowlines of $\text{grad}(a)$ (up to sign) are maps $u : I \times I \to W$ with left and right edges on the L_j, satisfying

$$\frac{\partial u}{\partial s} - J\frac{\partial u}{\partial t} = 0$$

where the horizontal factor on $I \times I$ has coordinate t, and the vertical factor s.

In words: gradient flowlines of a are holomorphic rectangles with edges on the L_j. If $x, y \in L_0 \cap L_1$ are critical points of da, then the flowlines from x to y are holomorphic maps $u : \mathbb{D} \to M$ sending $-i$ to x and i to y, and so that the arcs of $\partial\mathbb{D}$ with positive (resp. negative) real part maps to L_1 (resp. L_0).

Let $\mathcal{M}(x, y)$ denote the space of such *holomorphic Whitney disks*. The automorphism group of $\mathbb{D}$ fixing i and $-i$ is $\mathbb{R}$; this acts freely on $\mathcal{M}(x, y)$ (at least when $x \neq y$ so that $\mathcal{M}(x, y)$ contains no constant maps) and we denote the quotient space $\overline{\mathcal{M}}(x, y)$.

The space $\mathcal{M}(x, y)$ might have many components, and the disks u in $\mathcal{M}$ might lie in different homology classes. For a residual set of almost complex structures J the space $\mathcal{M}(x, y)$ is a smooth manifold of dimension that may be computed by the index formula. Since $\mathbb{D}$ is contractible the pullback u^*TW has a (symplectic) trivialization as $\mathbb{D} \times \mathbb{R}^{2n}$. The circle $S^1 = \partial\mathbb{D}$ factorizes as the union of two arcs $\alpha_0 \cup \alpha_1$ where $u : \alpha_j \to L_j$, each oriented to run from x to y. Relative to the trivialization each arc $\alpha_j^* TL_j$ may be thought of as a path in $\mathcal{L}_n$. Join these two paths at the endpoints by paths that do not cross some fixed train. The resulting loop in $\mathcal{L}_n$ has Maslov index $\mu(u)$ and Viterbo [24] shows that the formal dimension of $\mathcal{M}(x, y)$ in the component containing u is $\mu(u)$.

We may now define the Floer Homology of the pair L_0, L_1 as follows. The chain group is the free abelian group generated by intersections $L_0 \cap L_1$. For each x define

$$\partial x := \sum_y n(x, y)y$$

where $n(x, y)$ is the signed count of points in the components of the unparameterized moduli spaces $\overline{\mathcal{M}}(x, y)$ of formal dimension 0 (i.e. for which the corresponding components of $\mathcal{M}(x, y)$ are represented by holomorphic Whitney disks u with Maslov index 1).

This satisfies $\partial^2 = 0$; the homology of this complex is the Floer Homology.

5. Acknowledgments

References

[1] S. Altschuler, *A geometric heat flow for one-forms on three-dimensional manifolds*, Illinois J. Math. **39** (1995), 98–118

[2] V. Arnol'd, *The Sturm Theorems and Symplectic Geometry*, Func. Anal. Appl. **19** (1985), no. 4, 251–259

[3] S. Donaldson, *An application of gauge theory to the topology of 4-manifolds*, J. Diff. Geom **18** (1983), 269–316

[4] Y. Eliashberg, *A few remarks about symplectic filling*, Geom. Top. **8** (2004), 277–293

[5] Y. Eliashberg and W. Thurston, *Confoliations*, Univ. Lect. Ser. **13**, AMS, Providence, RI, 1998.

[6] J. Etnyre, *Lectures on open book decompositions and contact structures*, Clay Math. Proc. 5, AMS, Providence RI, 2006, 103–141
[7] J. Etnyre, *On symplectic fillings*, Alg. Geom. Top. **4** (2004), 73–80
[8] A. Floer, *An Instanton Invariant for 3-Manifolds*, Comm. Math. Phys. **118** (1988), 215–240
[9] A. Floer, *Morse theory for Lagrangian Intersections*, J. Diff. Geom. **28** (1988), 513–547
[10] A. Floer, *The Unregularized Gradient Flow of the Symplectic Action*, Comm. Pure. Appl. Math. **41** (1988), 775–813
[11] E. Giroux, ICM talk
[12] W. Goldman, *The Symplectic Nature of the Fundamental Groups of Surfaces*, Adv. Math. **54** (1984) 200–225
[13] M. Gromov, *Pseudo holomorphic curves in symplectic manifolds*, Invent. Math. **82** (1985), no. 2, 307–347
[14] H. Hofer, V. Lizan and J.-C. Sikorav, *On Genericity for Holomorphic Curves in Four-Dimensional Almost-Complex Manifolds*,
[15] K. Honda, W. Kazez and G. Matić, *Right-veering diffeomorphisms of compact surfaces with boundary*, Invent. Math. **169** (2007), no. 2, 427–449
[16] C. Hummel, *Gromov's Compactness Theorem for Pseudo-holomorphic Curves*, Prog. Math. **151**, Birkhauser, Basel, 1997.
[17] D. McDuff and D. Salamon, *Introduction to Symplectic Topology*, Oxford Math. Monographs, Oxford University Press, Oxford, 1998.
[18] D. McDuff and D. Salamon, *J-holomorphic Curves and Quantum Cohomology*, AMS University Lecture Series Vol. 6, AMS, Providence, 1994.
[19] M. Micallef and B. White, *The structure of branch points in minimal surfaces and in pseudoholomorphic curves*, Ann. Math. (2) **141** (1995), no. 1, 35–85
[20] Y. Mitsumatsu, *Foliations and contact structures on 3-manifolds*, Proceedings of Foliations: Geometry and Dynamics, P. Walczak Ed. 75–125, World Scientific, Singapore 2002.
[21] R. Sacksteder, *Foliations and pseudogroups*, Amer. J. Math. **87** (1965), 79–102
[22] C. Taubes, *Self-dual Yang-Mills connections on non-self-dual 4-manifolds*, J. Diff. Geom. **17** (1982), no. 1, 139–170
[23] C. Taubes, *Casson's Invariant and Gauge Theory*, J. Diff. Geom. **31** (1990) 547–599
[24] C. Viterbo, *Intersection de sous-variétés lagrangiennes, fonctionnelles d'action et indice des systèmes hamiltoniens*, Bull. SMF **115** (1987) 361–390

University of Chicago, Chicago, Ill 60637 USA
E-mail address: dannyc@math.uchicago.edu