

An Introduction to Symplectic Geometry
for Lagrangian Floer Homology

Maxine Elena Calle

A THESIS

in

Mathematics

Presented to the Faculties of the University of Pennsylvania in Partial
Fulfillment of the Requirements for the Degree of Master of Arts

2022

Supervisor of Thesis

Graduate Group Chairman

Contents

1	Some Symplectic Geometry	10
1.1	Symplectic vector spaces	11
1.1.1	Subspaces of a symplectic vector space	14
1.2	Symplectic manifolds	16
1.2.1	The cotangent bundle	19
1.3	Lagrangian submanifolds	24
1.3.1	Generating symplectomorphisms of cotangent bundles	29
1.4	Some Theorems	33
1.4.1	Moser theorems	33
1.4.2	Weinstein Lagrangian Neighborhood Theorem	37
2	Compatible Structures	39
2.1	Riemannian structures	40
2.1.1	Geodesic flow	42
2.2	Almost-complex structures	43

	3
2.2.1	Linear complex structures and compatible triples 44
2.2.2	Almost-complex structures and integrability 48
2.3	Compatible triples 50
3	Hamiltonians 54
3.1	Isotopies, vector fields, and Hamiltonian systems 54
3.1.1	Hamiltonian vector fields 58
3.1.2	Integrable systems 64
3.2	Time-dependent Hamiltonians and periodic solutions 67
3.2.1	The Arnold Conjecture 69
4	Lagrangian Floer Homology 71
4.1	Morse homology 73
4.2	Floer homology for Lagrangian intersections 80
4.2.1	The moduli space of J -holomorphic strips 88
4.2.2	Floer complex for Lagrangian intersections 94
4.3	Floer homology for Hamiltonian diffeomorphisms 103
4.4	Epilogue: Floer homotopy theory 110

Introduction

A natural first question to ask when broaching the subject of symplectic geometry is *what is a symplectic manifold?* Opening up our favorite textbook (such as [dS06, MS95, McD98]) to the first page or so, we will quickly learn that a symplectic structure on a smooth manifold is a closed, non-degenerate 2-form. But what does this mean, geometrically, and why is this a natural structure to study?

Classically, symplectic geometry arises as the natural setting for Hamiltonian mechanics, and investigating this connection can help us understand why symplectic structures are the “correct” abstraction to study. Suppose we have a manifold M and a single particle moving around in M . We can think of M as the configuration space $\text{Conf}_1(M)$ of different possible states that our small 1-particle system could be in. We could further complicate the system by adding in more particles, which would also complicate our space of possible states, that is, the underlying manifold we want to study.

However, the configuration space does not keep track of all the information we want it to. It sees things “discretely” in the sense that it cannot keep track of how

the system evolves over time. That is, picking out a point in our space (i.e. a possible state of our system) does not tell us anything about how the system might look 10 seconds later. In order to record such information, we need to somehow introduce dynamics on our configuration space, which we do mathematically via vector fields.

In classical mechanics, the dynamics typically come from an energy function $H: M \rightarrow \mathbb{R}$. We can use H to generate a vector field X_H , which describes how the energy is changing locally. To reflect this idea, we want X_H to depend only on H and to depend *linearly* on H (meaning $X_{aH_1+H_2} = aX_{H_1} + X_{H_2}$). Furthermore, following the “conservation of energy” rule from physics, we want H to be *constant* along the flow lines of X_H . Abstracting these principles to a mathematical setting, we land in the world of symplectic geometry. The following explanation should not be taken as rigorous proof, but is meant to give the reader an indication of why we might believe that this is the case.

First, we need a way to coherently associate each $H: M \rightarrow \mathbb{R}$ to a vector field X_H . We can reinterpret this as a linear association of X_H (a section of $TM \rightarrow M$) to the differential dH (a section of $T^*M \rightarrow M$). This means our association should be a bundle morphism $TM \rightarrow T^*M$, or equivalently a section of $T^*M \otimes T^*M \rightarrow M$. This is how we arrive at a tensor field ω such that $\omega(X_H, -) = dH$. Moreover, we want ω to be non-degenerate so that we can always solve for X_H . Conservation of energy implies that $0 = (dH)(X_H) = \omega(X_H, X_H)$ and so (modulo many details) ω is alternating, hence $\omega \in \Lambda^2(M)$. Finally, the fact that ω should be closed comes from

the idea that “the laws of physics should not depend on time.”

Symplectic geometry has a reputation for living somewhere between topology and geometry. This reputation is built upon the observation that a symplectic structure will impose more restrictions on a manifold than predicted by the topology alone, but will not be restrictive enough to distinguish between individual points. This latter statement is the content of Darboux’s Theorem, which is a fundamental result for symplectic manifolds. In particular, Darboux’s Theorem implies that a symplectic structure will not yield any local invariants, such as the curvature invariants of Riemannian geometry, and so we will need to develop non-local techniques and tools for studying symplectic manifolds and their symplectomorphisms.

One phenomena which naturally arises in the symplectic setting is that of Lagrangian submanifolds, which are submanifolds upon which ω vanishes and have the maximal dimension where this could possibly happen. Going back to the discussion of configuration spaces and dynamics, we can think of a Lagrangian submanifold as encoding the set of possible initial momenta of a given point in configuration space. To make this a bit more precise, any manifold (symplectic or not) lives inside its cotangent bundle (which admits a canonical symplectic structure) as a Lagrangian submanifold via the zero section. Moreover, given any Lagrangian submanifold N of a symplectic manifold M , we can find a neighborhood of N in M which looks like a neighborhood of N inside of T^*N ; this is the Weinstein neighborhood theorem. In addition to being interesting objects in their own right, Lagrangian submanifolds are

extremely useful tools for approaching problems in symplectic geometry, and perhaps one of the most famous applications of this sort is A. Floer's solution to the Arnold Conjecture for a certain collection of symplectic manifolds.

On the face of it, the Arnold conjecture is a statement about fixed points of Hamiltonian diffeomorphisms, that is, symplectomorphisms which arise as smooth deformations of M under a Hamiltonian flow. However it exists as a sub-problem of bounding intersection points of Lagrangian submanifolds. Floer homology was developed as a tool to address the problem of Lagrangian intersections, and unites ideas from Morse theory, symplectic geometry, and analysis. The idea of Floer homology is to emulate Morse theory as much as possible for a specific function on a specific manifold, although there are analytic difficulties that arise which do not appear in the Morse setting. In fact, it is well understood that Floer homology cannot be defined in general, although it has been thoroughly developed in specific cases.

The special properties of Floer homology that make it particularly useful rely on the interesting way that symplectic structures interact with other geometries, particularly Riemannian and almost complex structures. Remarkably, the choice of any two of these structures determines a unique compatible third one. This is another instance of symplectic geometry being both flexible and rigid: it is flexible enough to coexist with other geometries, but rigid enough that the cohabitation imposes certain restrictions.

The name "symplectic" is said to have been coined by H. Weyl in the 1930's. It

comes from the Latin word *complectere*, which roughly translates to *to weave or braid together*. While we feel this name is quite poetically appropriate, we leave it up to our readers to decide for themselves.

Outline

Our goal is to explore symplectic geometry with an eye towards understanding the geometry behind Floer homology for Lagrangian intersections. Rather than centering our attention on the details of Floer's proof of the Arnold Conjecture (which requires a surprising amount of analysis), we will primarily focus on the geometric background necessary to understand Lagrangian Floer homology, including symplectic manifolds and Lagrangian submanifolds, compatible Riemannian and almost complex structures, and Hamiltonian vector fields. We will not be able to give a complete account of any of these rich areas, but rather aim to give enough of an overview so that a reader who is unfamiliar with the symplectic background (say, a homotopy theorist who is interested in learning about Floer homotopy theory) feels more prepared to approach the field.

The first chapter covers the basics of symplectic geometry, including symplectic linear algebra (Section 1.1), symplectic manifolds and symplectomorphisms (Section 1.2), and a bit of the theory of Lagrangian submanifolds (Section 1.3). We conclude with a brief overview of some of the important theorems, including the Moser theorems and the Weinstein neighborhood theorems (Subsection 1.4.2). In the

second chapter, we investigate how symplectic structures interact with other geometric structures our manifold might possess. In particular, Riemannian (Section 2.1) and almost complex (Section 2.2) structures fit particularly nicely into the symplectic picture. A choice of any two of these structures determines the third, and the three structures together form a compatible triple (Section 2.3). Chapter 3 discusses Hamiltonian vector fields and their flows, with the goal of stating and understanding the Arnold Conjecture (Subsection 3.2.1), which bounds the fixed points of a Hamiltonian diffeomorphism from below. This prepares us for Chapter 4, where we recast the Arnold Conjecture as a Lagrangian intersection problem. We discuss the ideas of Floer homology for Lagrangian intersections (Section 4.2) and Hamiltonian diffeomorphisms (Section 4.3), and we conclude with a very brief discussion of Floer homotopy theory (Section 4.4).

Acknowledgements

I would like to thank my project advisor, Jonathan Block, for his helpful insights in supervising this project. This thesis was completed as part of my Ph.D. preliminary exam at the University of Pennsylvania, and I am grateful to the other members of my exam committee, Mona Merling and Rob Ghrist, as well as the math graduate chair Ron Donagi, for this opportunity to learn about something outside of my research area. I am also supported by the National Science Foundation Graduate Research Fellowship Program (DGE-1845298).

Chapter 1

Some Symplectic Geometry

We begin with a bit of symplectic linear algebra (Section 1.1) to prepare for the symplectic manifold theory. In Subsection 1.2.1, we focus on the cotangent bundle as an important example of a symplectic manifold. Then in Section 1.3, we explore the theory of Lagrangian submanifolds, and in particular how they can help us generate and identify symplectomorphisms. Finally, we conclude with a survey of some foundational theorems of symplectic geometry, including the Moser theorems (Subsection 1.4.1) and the Weinstein Lagrangian neighborhood theorem (Subsection 1.4.2). This material may be found in any standard reference for symplectic geometry, and our references include [dS06, McD98, MS95].

1.1 Symplectic vector spaces

Let V be a finite-dimensional vector space over \mathbb{R} and let $\Omega: V \times V \rightarrow \mathbb{R}$ be a bilinear form on V . By the tensor-hom adjunction, Ω also gives us a map $\iota\Omega: V \rightarrow V^*$. More concretely, for each $v \in V$ we have a map

$$\begin{aligned}\Omega_v: V &\rightarrow \mathbb{R} \\ v' &\mapsto \Omega(v, v')\end{aligned}$$

and $\iota\Omega$ just sends $v \in V$ to $\Omega_v \in V^*$.

Definition 1.1. A *symplectic form* on V is a 2-form $\Omega \in \Lambda^2(V^*)$. That is, $\Omega: V \times V \rightarrow \mathbb{R}$ is a bilinear form which is

- (i) skew-symmetric: $\Omega(v, v') = -\Omega(v', v)$ for all $v, v' \in V$;
- (ii) non-degenerate: for $v \in V$, $\Omega(v, v') = 0$ for all $v' \in V$ if and only if $v = 0$.

An equivalent condition to (ii) is that $\ker \iota\Omega = 0$, which (since V is finite-dimensional) means $\iota\Omega$ is bijective. The map Ω is said to be a *linear symplectic structure* on V and (V, Ω) is called a *symplectic vector space*.

The symplectic linear structure Ω restricts some of the properties of the vector space V . For instance, if (V, Ω) is a symplectic vector space, then $\dim V$ must be even. This is because the matrix representation of Ω with respect to a chosen basis is skew-symmetric and invertible: if we let A denote the matrix representation of Ω ,

then $\det A = \det(-A^t) = (-1)^{\dim V} \det A$. In fact, there is a (not necessarily unique) choice of basis for V called a *symplectic basis* which gives the matrix representation of Ω a particularly nice form.

Definition 1.2. A *symplectic basis* for a symplectic vector space (V, Ω) is a basis $e_1, \dots, e_n, f_1, \dots, f_n$ for V such that $\Omega(e_i, e_j) = 0 = \Omega(f_i, f_j)$ and $\Omega(e_i, f_j) = \delta_{ij}$ for all $1 \leq i, j \leq n$. With respect to this basis, Ω has matrix representation

$$\begin{bmatrix} 0 & \text{id}_n \\ -\text{id}_n & 0 \end{bmatrix}.$$

where id_n is the $n \times n$ identity matrix.

By a skew-symmetric version of the Gram-Schmidt algorithm, we can always find such a symplectic basis. Take any non-zero $e_1 \in V$. By non-degeneracy, there is some $f_1 \in V$ such that $\Omega(e_1, f_1) \neq 0$ and by scaling f_1 we can moreover assume that $\Omega(e_1, f_1) = 1$. Write $V_1 = \text{span}\{e_1, f_1\}$, and decompose $V = V_1 \oplus V_1^\Omega$, where

$$V_1^\Omega := \{v \in V \mid \Omega(v, v') = 0 \text{ for all } v' \in V_1\}.$$

This subspace is called the *symplectic complement* of V_1 in V . Continue inductively, taking non-zero $e_2 \in V_1^\Omega$ and $f_2 \in V_1^\Omega$ with the right properties, and decomposing $V_1^\Omega = V_2 \oplus V_2^\Omega$ for $V_2 = \text{span}\{e_2, f_2\}$. Since $\dim V$ is finite, this process terminates eventually and the resulting $e_1, \dots, e_n, f_1, \dots, f_n$ are our symplectic basis.

By bilinearity, the values of Ω are determined by its values on a symplectic basis. Thus, given a basis of V , we can give V a symplectic structure by defining a form whose symplectic basis is the given one. So every even-dimensional vector space has a symplectic structure, and the next example introduces the canonical one.

Example 1.3. Let $e_1, \dots, e_n, f_1, \dots, f_n$ be the canonical basis for \mathbb{R}^{2n} (so e_i has a 1 in the i^{th} spot and zeros elsewhere and f_j has a 1 in the $(n + j)^{\text{th}}$ spot and zeros elsewhere). This basis determines a symplectic vector space structure $(\mathbb{R}^{2n}, \Omega_0)$. Specifically, if $v = \sum_{i=1}^n a_i e_i + \sum_{j=1}^n b_j f_j$ and $w = \sum_{i=1}^n c_i e_i + \sum_{j=1}^n d_j f_j$, then we have $\Omega_0(v, w) = \sum_{i=1}^n a_i d_i - \sum_{i=1}^n b_i c_i$.

We would like to be able to compare symplectic structures on \mathbb{R}^{2n} , which motivates us to define maps between symplectic vector spaces more generally.

Definition 1.4. A *symplectomorphism* between symplectic vector spaces (V, Ω) and (V', Ω') is a linear isomorphism $\phi: V \rightarrow V'$ which preserves the symplectic form under pullback, $\Omega = \phi^* \Omega'$. That is, $\Omega(v, v') = \Omega'(\phi(v), \phi(v'))$ for all $v, v' \in V$. The symplectic vector spaces (V, Ω) and (V', Ω') are said to be *symplectomorphic*.

In the case that $(V, \Omega) = (V', \Omega')$, a symplectomorphism will actually preserve the symplectic form, and is called a *linear symplectic transformation* of V . The symplectic transformations of V form a Lie group called the *symplectic group* and is denoted by $Sp(V, \Omega)$ (or sometimes just $Sp(V)$ if the form Ω is clear from context). That is,

$$Sp(V) = \{A \in GL_{2n}(\mathbb{R}) \mid \Omega(u, v) = \Omega(Au, Av)\} \leq GL_{2n}(\mathbb{R}).$$

Being symplectomorphic defines an equivalence relation between even dimensional real vector spaces. Moreover every $2n$ -dimensional symplectic vector space is symplectomorphic to $(\mathbb{R}^{2n}, \Omega_0)$ by a change of basis from the given symplectic basis to the standard one.

1.1.1 Subspaces of a symplectic vector space

A linear symplectic structure Ω on V picks out special subspaces of V , based on the behavior of Ω on the subspace. Given a linear subspace $W \subseteq V$, its *symplectic orthogonal* is

$$W^\Omega := \{v \in V \mid \Omega(v, w) = 0 \text{ for all } w \in W\} \subseteq V.$$

That is, W^Ω is the kernel of the map $V \rightarrow W^*$ which sends $v \mapsto \Omega_v|_W$. The symplectic orthogonal enjoys a lot of the same properties as the usual orthogonal complement W^\perp . For example, we have $(W^\Omega)^\Omega = W$. This is because $\dim W + \dim W^\Omega = \dim V$ and the inclusion $W \subseteq (W^\Omega)^\Omega$ is thus an injective map of vector spaces of the same dimension.

Subspaces of V are given special names depending on how they intersect with their symplectic orthogonal.

Definition 1.5. Let W be a subspace of (V, Ω) . Then

- (i) W is *symplectic* if $\Omega|_W$ is non-degenerate, so $(W, \Omega|_W)$ is a symplectic vector space. This is equivalent to saying that $W \cap W^\Omega = 0$.

(ii) W is *isotropic* if $W \subseteq W^\Omega$, so $\Omega|_W \equiv 0$.

(iii) W is *coisotropic* if $W^\Omega \subseteq W$.

(iv) W is *Lagrangian* if it is isotropic and coisotropic, that is, $W = W^\Omega$.

By the equality $\dim V = \dim W + \dim W^\Omega$, every isotropic subspace must have dimension $\leq \frac{1}{2} \dim V$. Moreover, there is equality if and only if W is Lagrangian.

Lagrangian subspaces have some particularly nice properties and are important objects of study in symplectic geometry and topology. For example, any basis of a Lagrangian subspace W can be extended to a symplectic basis of V . Let e_1, \dots, e_n be a basis of W and choose an appropriate $f_1 \in \text{span}\{e_2, \dots, e_n\}^\Omega$. We can find this f_1 by non-degeneracy of Ω , and we know $\Omega(e_1, e_i) = 0$ for all i because W is Lagrangian. Similarly, we can choose $f_2 \in \text{span}\{e_1, e_3, \dots, e_n, f_1\}^\Omega$, and continue until we have a symplectic basis for V .

Example 1.6. We will find examples of subspaces of type (i)–(iv) in $(\mathbb{R}^{2n}, \Omega_0)$. An example of a symplectic subspace is $\text{span}\{e_1, f_1\}$ and an example of an isotropic subspace is $\text{span}\{e_1, e_2\}$ — even though these two subspaces are isomorphic as \mathbb{R} -vector spaces, they are *not* symplectomorphic. The isotropic subspace $\text{span}\{e_1, e_2\}$ cannot be Lagrangian, as the dimension is too small. We need to expand up to $\text{span}\{e_1, \dots, e_n\}$ to get a Lagrangian subspace. Also, $\text{span}\{e_1, \dots, e_n, f_1\}$ is an example of a subspace which is coisotropic but not Lagrangian.

Example 1.7. Any codimension 1 subspace is coisotropic. Suppose $\text{codim}W = 1$ and

choose $v \in V \setminus W$, so $V = W \oplus \text{span}\{v\}$. By non-degeneracy, there must be some other $w \in V$ so that $\Omega(v, w) \neq 0$ and hence in particular $w \notin \text{span}\{v\}$. So $w \in W$ and hence $v \notin W^\Omega$, implying $W^\Omega \subseteq W$.

Example 1.8. If W is a Lagrangian subspace of (V, Ω) , then there is a symplectomorphism between (V, Ω) and $(W \oplus W^*, \Omega')$ for

$$\Omega'(w \oplus f, w' \oplus f') = f'(w) - f(w').$$

Let e_1, \dots, e_n be a basis of W and let e^1, \dots, e^n be the dual basis for W^* . Extend e_1, \dots, e_n to a symplectic basis $e_1, \dots, e_n, f_1, \dots, f_n$ for (V, Ω) and define $\phi: V \rightarrow W \oplus W^*$ by $\phi(e_i) = e_i \oplus 0$ and $\phi(f_j) = 0 \oplus e^j$. The fact that e_1, \dots, e_n is a symplectic basis for (V, Ω) implies that $e_1 \oplus 0, \dots, e_n \oplus 0, 0 \oplus e^1, \dots, 0 \oplus e^n$ is a symplectic basis for $(W \oplus W^*, \Omega')$.

More generally, if W is any vector space, then $W \oplus W^*$ has a canonical symplectic structure given by the formula in the example above. In this sense, every symplectic vector space is the direct sum of a vector space and its dual.

1.2 Symplectic manifolds

A symplectic form on a manifold M is a differential 2-form which satisfies two conditions: closedness and non-degeneracy. The former is an analytic condition and the latter is an algebraic one. Recall that a 2-form $\omega \in \Lambda^2(M)$ assigns each $p \in M$ a

skew-symmetric bilinear form $\omega_p: T_pM \times T_pM \rightarrow \mathbb{R}$ and the assignment $p \mapsto \omega_p$ is smoothly varying. The form ω is *closed* if its exterior derivative is 0 ($d\omega = 0$) and *exact* if it is the image of some 1-form under d ($\omega = d\alpha$ for some $\alpha \in \Lambda(M)$).

Definition 1.9. A *symplectic structure* on a manifold M is a closed 2-form ω on M such that (T_pM, ω_p) is a symplectic vector space for all $p \in M$. A manifold M with a symplectic structure ω is called a *symplectic manifold*. A *symplectomorphism* $\phi: (M, \omega) \rightarrow (M', \omega')$ is a diffeomorphism such that $\omega = \phi^*\omega'$. This pullback condition means that $\omega_p(v, u) = \omega'_{\phi(p)}(d\phi_p(u), d\phi_p(v))$ for all $p \in M$ and $u, v \in T_pM$.

Example 1.10. Let $M = \mathbb{R}^{2n}$ with linear coordinates $x_1, \dots, x_n, y_1, \dots, y_n$. Then

$$\omega_0 := \sum_{i=1}^n dx_i \wedge dy_i$$

gives \mathbb{R}^{2n} the structure of a symplectic manifold. A symplectic basis of T_pM is given by $\left(\frac{\partial}{\partial x_i}\right)_p, \left(\frac{\partial}{\partial y_i}\right)_p$ for $0 \leq i \leq n$.

Just as any n -manifold is locally “the same” as \mathbb{R}^n , any $2n$ -symplectic manifold is locally “the same” as $(\mathbb{R}^{2n}, \omega_0)$.

Theorem 1.11. [*Darboux’s Theorem*] Let (M, ω) be a $2n$ -dimensional symplectic manifold and let $p \in M$. There is a coordinate chart $(U; x_1, \dots, x_n, y_1, \dots, y_n)$ centered at p such that

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

on U .

Such a chart is called a *Darboux chart*. Darboux's Theorem tells us there are no local invariants on symplectic manifolds (up to symplectomorphism). This marks a stark difference between symplectic geometry and some other types of geometry; for instance, Riemannian geometry has local curvature invariants which let us distinguish between points on a manifold.

Although there are no local invariants, the symplectic form ω does impose some global restrictions on the manifold M . For instance, the n^{th} exterior power $\omega^n = \omega \wedge \cdots \wedge \omega$ is a *volume form* (a non-vanishing form of top degree) on the $2n$ -manifold M , and thus gives M a canonical orientation. This means that non-orientable manifolds (e.g. the Möbius band) cannot be given a symplectic structure.

To prove that $\omega^n \neq 0$, it suffices to work locally. For $p \in M$, we can take a Darboux chart upon which ω has the form $\sum_i dx_i \wedge dy_i$. By induction, we can show

$$\omega^k = k! \sum_{1 \leq i_1 \leq \cdots \leq i_k \leq n} dx_{i_1} \wedge dy_{i_1} \wedge \cdots \wedge dx_{i_k} \wedge dy_{i_k},$$

so in particular $\omega^n = n! dx_1 \wedge dy_1 \wedge \cdots \wedge dx_n \wedge dy_n$ is non-zero. The form $\frac{\omega^n}{n!}$ is called the *symplectic volume* of (M, ω) .

Moreover, if M is a closed manifold, then the discussion above implies that the de Rham cohomology class $[\omega^n] \in H^{2n}(M, \mathbb{R})$ is non-zero. Otherwise, if $\omega^n = d\omega'$ for some $(2n - 1)$ -form ω' , then by Stokes' theorem

$$\text{vol}(M) = \int_M \omega^n = \int_{\partial M} \omega' = 0$$

since $\partial M = \emptyset$, which is absurd. Since $[\omega^n] = [\omega]^n$ is non-zero in $H^{2n}(M, \mathbb{R})$, it follows that $[\omega] \in H^2(M, \mathbb{R})$ is also non-zero, which is to say ω is not exact. Thus any compact manifold with trivial H^2 cannot have a symplectic structure, e.g. S^{2n} for $n \geq 2$. A similar argument works to show that a closed $2n$ -manifold with trivial H^{2k} for some $1 \leq k \leq n$ cannot be a symplectic manifold.

Example 1.12. Every orientable surface is symplectic. This is because orientability implies the existence of a volume form, which (because surfaces are 2-dimensional) is also a symplectic form.

1.2.1 The cotangent bundle

Many important examples of symplectic structures come from the cotangent bundle of a manifold. Recall that the cotangent bundle

$$T^*M = \{(p, \xi) \mid p \in M, \xi \in T_p^*M\}$$

inherits the structure of a $2n$ -manifold from the n -manifold M . Let $(U; x_1, \dots, x_n)$ be a chart for M . Then for any $p \in U$, the differentials $(dx_1)_p, \dots, (dx_n)_p$ form a basis for T_p^*U , meaning any $\xi \in T_p^*M$ can be written as $\xi = \sum_{i=1}^n \xi_i (dx_i)_p$, and so $(T^*U; x_1, \dots, x_n, \xi_1, \dots, \xi_n)$ is a chart for T^*M .

Two important forms on T^*M are the *tautological 1-form* (or *Louiville 1-form*) α and the *canonical symplectic form* $\omega = -d\alpha$. Ironically, the tautological nature of α makes it a bit difficult to explain, but we will do our best to define these forms both

with and without coordinates.

Definition 1.13. (Tautological and canonical form) The tautological 1-form α is a section of the bundle $T^*(T^*M) \rightarrow T^*M$, so assigns each $(p, \xi) \in T^*M$ to a linear functional $\alpha_{(p, \xi)}: T_{(p, \xi)}(T^*M) \rightarrow \mathbb{R}$. Recall that the cotangent bundle T^*M comes with a projection $\pi: T^*M \rightarrow M$ which induces a map $d\pi: T(T^*M) \rightarrow TM$. This in turn induces pullback map $(d\pi)^*: T^*M \rightarrow T^*(T^*M)$ and we will use this to define α . At a point $(p, \xi) \in T^*M$, the pullback $(d\pi)^*_{(p, \xi)}$ sends a covector $\zeta \in T^*_p M$ to the covector $(d\pi)^*_{(p, \xi)}\zeta = \zeta \circ (d\pi)_p \in T^*_{(p, \xi)}(T^*M)$. Define

$$\alpha_{(p, \xi)} := (d\pi)^*_{(p, \xi)}(\xi) = \xi \circ (d\pi)_p.$$

Note that this makes sense because $\xi \in T^*_p M$ is in the domain of $(d\pi)^*_{(p, \xi)}$. Given a tangent vector $\nu \in T_{(p, \xi)}(T^*M)$, the value of $\alpha_{(p, \xi)}(\nu) = \xi((d\pi)_p(\nu))$ is computed by projecting ν onto the tangent space $T_p M$ using $d\pi_p$, and then applying ξ to this projection.

The canonical symplectic 2-form ω is defined as $-d\alpha$. To see that ω is in fact symplectic, we just need to know it is non-degenerate. It will be easiest to do this in local coordinates. Let $(T^*U; x_1, \dots, x_n, \xi_1, \dots, \xi_n)$ be a coordinate chart around $(p, \xi) \in T^*M$. Then $dx_1, \dots, dx_n, d\xi_1, \dots, d\xi_n$ forms a basis for $T^*_{(p, \xi)}(T^*M)$ so we can write $\alpha = \sum_{i=1}^n \alpha_i dx_i + \sum_{j=1}^n \tilde{\alpha}_j d\xi_j$ for some smooth $\alpha_i, \tilde{\alpha}_j$. By the coordinate-free

description of α , we see that $\alpha_i = \xi_i$ and $\tilde{\alpha}_j = 0$, so

$$\alpha = \sum_{i=1}^n \xi_i dx_i.$$

Thus the local description of $\omega = -d\alpha$ is

$$\omega = \sum_{i=1}^n dx_i \wedge d\xi_i,$$

which coincides with the standard symplectic form on \mathbb{R}^{2n} .

Proposition 1.14. *The tautological 1-form α is uniquely characterized by the property that $\mu^*\alpha = \mu$ for every $\mu: M \rightarrow T^*M$.*

Proof. We first prove uniqueness. Suppose $\beta \in \Lambda(TM)$ such that $\mu^*\beta = \mu$ for all $\mu \in \Lambda(M)$. Let $(U; x_1, \dots, x_n)$ be a chart of $p \in M$, so we can write $\mu = \sum_i \mu_i dx_i$ in coordinates. Then in the chart $(T^*U; x_1, \dots, x_n, \xi_1, \dots, \xi_n)$ of T^*M , we can also write $\beta = \sum_i \beta_i dx_i + \sum_j \tilde{\beta}_j d\xi_j$. In these coordinates, we have $\mu^*\beta = \sum_i (\beta_i \circ \mu) dx_i + \sum_j (\tilde{\beta}_j \circ \mu) d\mu_j$. If this is equal to μ , then we must have $\beta_i \circ \mu = \mu_i$ and $\tilde{\beta}_j \circ \mu = 0$, meaning that $\beta_i = \xi_i$ and $\tilde{\beta}_j \equiv 0$. But this is precisely the local definition of α , and this computation also shows that α satisfies the pullback condition. \square

It follows immediately from this proposition that the canonical 2-form ω has the feature that $\mu^*\omega = -d\mu$ for any $\mu \in \Lambda(M)$, since $\mu^*(-d\alpha) = -d(\mu^*\alpha)$. Another property of these two forms is that they are natural in the following sense:

Proposition 1.15. *Let $\phi: M \rightarrow M'$ be a diffeomorphism. Then the lift $\tilde{\phi}: T^*M \rightarrow T^*M'$ is a diffeomorphism which pulls α' back to α ,*

$$(\tilde{\phi})^* \alpha' = \alpha.$$

Proof. Recall that the lift $\tilde{\phi}$ is defined pointwise by $\tilde{\phi}(p, \xi) = (\phi(p), \tilde{\xi})$ where $\tilde{\xi}$ is such that $(d\phi)_p^* \tilde{\xi} = \xi$. Focusing on the second component, we note that $\tilde{\phi}|_{T_p^*M}: T_p^*M \rightarrow T_{\phi(p)}^*M$ is just defined to be the inverse of $(d\phi)_p^*: T_{\phi(p)}^*M' \xrightarrow{\cong} T_p^*M$, which (modulo some details) implies that $\tilde{\phi}$ is also a diffeomorphism. We will now show the pullback condition holds at every point $(p, \xi) \in T^*M$. Compute

$$\begin{aligned} ((\tilde{\phi})^* \alpha')(p, \xi) &= (d\tilde{\phi})_{(p, \xi)}^* \alpha'_{(\phi(p), \tilde{\xi})} \\ &= (d\tilde{\phi})_{(p, \xi)}^* (d\pi')_{(\phi(p), \tilde{\xi})}^* \tilde{\xi} && \text{by definition of } \alpha', \\ &= d(\pi' \circ \tilde{\phi})_{(p, \xi)}^* \tilde{\xi} \\ &= d(\phi \circ \pi)_{(p, \xi)}^* \tilde{\xi} && \text{since } \tilde{\phi} \text{ is a lift,} \\ &= (d\pi)_{(p, \xi)}^* (d\phi)_p^* \tilde{\xi} \\ &= (d\pi)_{(p, \xi)}^* \xi && \text{by definition of } \tilde{\xi}, \\ &= \alpha_{(p, \xi)} && \text{by definition.} \end{aligned}$$

□

An immediate corollary of this proposition is that the lift $\tilde{\phi}$ also pulls ω' back to

ω , and hence is a symplectomorphism $T^*M \rightarrow T^*M'$.

In the special case that $M' = M$, we get an injective group homomorphism $\text{Diff}(M) \rightarrow \text{Sym}(T^*M, \omega)$. This homomorphism is not surjective, which is to say that not every symplectomorphism of (T^*M, ω) will arise this way. For example, symplectomorphisms given by translation along cotangent fibers is not induced by any diffeomorphism $M \rightarrow M$. However, any symplectomorphism ψ of T^*M which preserves α is of the form $\psi = \phi^*$ for some diffeomorphism $\phi: M \rightarrow M$. This is because if ψ preserves α , then it must also preserve the cotangent fibration, meaning that we can find a diffeomorphism $\phi: M \rightarrow M$ so that $\pi \circ g = \phi \circ \pi$. Moreover, we can show that $\psi = \tilde{\phi}$ via the same computation as in the proof above, using the fact that $g^*\alpha = \alpha$ and $\pi \circ g = \phi \circ \pi$.

We can also use smooth functions on M to generate symplectomorphisms of (T^*M, ω) which may not preserve α . If $f: M \rightarrow \mathbb{R}$ is any smooth function, define $\tau_f: T^*M \rightarrow T^*M$ by $(x, \xi) \mapsto (x, \xi + df_x)$. That is, τ_f shifts along the cotangent fibers by df . It follows that

$$\tau_f^*\alpha = \alpha + \pi^*df,$$

because

$$\begin{aligned}
(\tau_f^* \alpha)(p, \xi) &= (d\tau_f)_{(p, \xi)}^* \alpha_{(p, \xi + df_p)} \\
&= (d\tau_f)_{(p, \xi)}^* (d\pi)_{(p, \xi + df_p)}^* (\xi + df_p) \\
&= d(\pi \circ \tau_f)_{(p, \xi)}^* (\xi + df_p) \\
&= d\pi_{(p, \xi)}^* \xi + d\pi_{(p, \xi)}^* df_p \\
&= \alpha_{(p, \xi)} + (\pi^* df)_{(p, \xi)}.
\end{aligned}$$

This implies $\tau_f^* \omega = \omega$, and therefore τ_f is a symplectomorphism as claimed.

In the next subsection, we will see how to generate symplectomorphisms between $T^*M \rightarrow T^*M'$ outside of the specific case $M' = M$, using Lagrangian submanifolds.

1.3 Lagrangian submanifolds

Lagrangian submanifolds are one of the phenomena which arise naturally in the symplectic setting. In addition to being interesting objects of study in their own right, Lagrangian submanifolds are incredibly useful for generating and identifying symplectomorphisms. We will also see in Chapter 4 that we can recast the Arnold Conjecture in terms of intersection points of Lagrangian submanifolds.

We saw in Section 1.1 that Lagrangian subspaces are the biggest subspaces upon which the symplectic form vanishes. This definition carries over to manifolds.

Definition 1.16. A submanifold $N \subseteq M$ is *Lagrangian* if for each $p \in N$, $T_p N$ is a

Lagrangian subspace of $T_p M$. Equivalently, if $i: N \rightarrow M$ is the inclusion, then N is Lagrangian if $i^* \omega = 0$ and $\dim N = \frac{1}{2} \dim M$.

If N is a Lagrangian submanifold, we can canonically identify its normal bundle with its cotangent bundle. This is because $\nu_p N := T_p M / T_p N$ is isomorphic to $T_p^* N$, via the canonical non-degenerate bilinear pairing $\Omega': T_p M / T_p N \times T_p N \rightarrow \mathbb{R}$ given by $\omega'_p([v], w) = \omega_p(v, w)$ for $[v] \in \nu_p N$. Non-degeneracy of this bilinear pairing implies $\tilde{\omega}'_p: \nu_p N \rightarrow T_p^* N$ is an isomorphism.¹ This observation, together with the following example, will be important for our statement of the Weinstein tubular neighborhood theorem (Subsection 1.4.2). In fact, we shall see that the Weinstein neighborhood theorem says that, in a certain sense, the following example is the most important (or even only) example of a Lagrangian submanifold.

Example 1.17. Every manifold M is a Lagrangian submanifold of its cotangent bundle $T^* M$, via the zero section. Recall that the zero section of M is

$$z(M) = \{(x, \xi) \mid x \in M, \xi \equiv 0 \in T_x^* M\} \subseteq T^* M,$$

which is isomorphic to M by projection onto the first coordinate. Since $\dim T^* M = 2 \dim M$, M has the correct dimension to be a candidate Lagrangian subspace, so it suffices to show that $z^* \omega = 0$ (where ω is the canonical 2-form on $T^* M$, see

¹This works more generally for any symplectic vector space (V, Ω) and Lagrangian subspace W . The bilinear form $\Omega': V/W \times W \rightarrow \mathbb{R}$ which sends $([v], w)$ to $\Omega(v, w)$ induces a canonical isomorphism $\tilde{\Omega}': V/W \xrightarrow{\cong} W^*$. Note that Ω' is well-defined because if $[v] = [v']$, then $v - v' \in W$ and so $\Omega'([v], w) - \Omega'([v'], w) = \Omega'(v - v', w) = 0$; it is non-degenerate because $W^\Omega = W$.

Definition 1.13). In a chart $(T^*U, x_1, \dots, x_n, \xi_1, \dots, \xi_n)$, any points of $z(M)$ in this chart will have $\xi_1 = \dots = \xi_n = 0$. So $\alpha = \sum_{i=1}^n \xi_i dx_i$ is 0 on $z(M)$, and hence so is $\omega = -d\alpha$. Therefore M is a Lagrangian submanifold of T^*M .

Generalizing this example, given any section $\mu: M \rightarrow T^*M$, we can ask whether the graph

$$M_\mu = \{(x, \mu(x)) \mid x \in M\} \subseteq T^*M$$

is Lagrangian. Remarkably, the answer depends only on whether μ is closed (as a 1-form) or not.

Proposition 1.18. *The submanifold $M_\mu \subseteq T^*M$ is Lagrangian if and only if μ is a closed 1-form.*

Proof. Note that M_μ is diffeomorphic to M under μ , and hence satisfies the dimension requirement for Lagrangian submanifold-ness. Our goal is to show that the inclusion $i: M_\mu \rightarrow T^*M$ pulls ω back to 0 if and only if $d\mu = 0$. To this end, we note that we can decompose $\mu = i \circ \tilde{\mu}$, where $\tilde{\mu}: M \rightarrow M_\mu$ is the diffeomorphism of M onto its image in T^*M .

Recall from Proposition 1.14 that $\mu^*\alpha = \mu$. In particular, this means $d\mu = 0$ if and only if $\mu^*d\alpha = 0$. But

$$\mu^*d\alpha = (i \circ \tilde{\mu})^*d\alpha = \tilde{\mu}^*i^*d\alpha = i^*d\alpha = -i^*\omega,$$

and hence $d\mu = 0$ if and only if $i^*\omega = 0$, which is to say that M_μ is Lagrangian. \square

There is one conspicuous class of closed 1-forms which we can draw from: the exact ones, i.e. those μ such that $\mu = df$ for some $f \in C^\infty(M)$. Such an f is called a *generating function* for M_μ , and two functions generate the same M_μ if and only if they differ by a locally constant function. Note that if $H_{dR}^1(M) = 0$, e.g. when M is simply connected, then every closed 1-form is exact and hence every Lagrangian submanifold of the form M_μ will be generated by some f . However, not every Lagrangian submanifold of T^*M will arise as a graph M_μ , as we shall see in the next example.

Example 1.19. Let N be any k -dimensional submanifold of M . Recall that the *conormal bundle* of N in M is

$$\nu^*N = \{(p, \xi) \in T^*M \mid p \in N, \xi \in \nu_p^*N\}$$

where $\nu_p^*N := \{\xi \in T_p^*M \mid \xi(v) = 0 \text{ for all } v \in T_pN\}$ is the *conormal space* at $p \in N$. The inclusion $i: \nu^*N \rightarrow T^*M$ exhibits ν^*N as a n -submanifold of T^*M , which can be seen most easily using charts adapted to N . Let $(U; x_1, \dots, x_n)$ be a chart centered at $p \in N$ such that $x_{k+1} = \dots = x_n = 0$ on $U \cap N$. If $(T^*U; x_1, \dots, x_n, \xi_1, \dots, \xi_n)$ is the corresponding chart on T^*M , then $\nu^*N \cap T^*U$ is described by $x_{k+1} = \dots = x_n = 0$ and $\xi_1 = \dots = \xi_k = 0$. Thus ν^*N is an n -submanifold of the $2n$ -manifold T^*M , with k degrees of freedom coming from x_1, \dots, x_k and $n - k$ degrees of freedom coming from ξ_{k+1}, \dots, ξ_n .

In fact, ν^*N is a Lagrangian submanifold, which we can show locally. Recall that

α is described locally on T^*U as $\sum_i \xi_i dx_i$, so in particular for $p \in \nu^*N$,

$$(i^*\alpha)_p = \alpha_p|_{T_p(\nu^*N)} = \sum_{i>k} \xi_i dx_i|_{\text{span}\{\frac{\partial}{\partial x_i}\}_{i \leq k}} = 0.$$

Hence $i^*\alpha = 0$ so ν^*N is a Lagrangian submanifold of T^*M .

Note that ν^*N will not be a graph of a section $M \rightarrow T^*M$ in general. However, if we take the extreme example of $N = M$, then we see that the conormal bundle ν^*M is just the zero section $z(M)$ which we saw was Lagrangian in Example 1.17. At the other extreme is $N = \{p\}$, in which case the conormal bundle $\nu^*\{p\}$ is the cotangent fiber T_p^*M .

Lagrangian submanifolds can also help identify when a diffeomorphism $\phi: M_1 \rightarrow M_2$ is a symplectomorphism. The idea is check whether the graph of the candidate symplectomorphism is a Lagrangian submanifold of $M_1 \times M_2$, much like we did with M_μ inside of T^*M .

Let (M_1, ω_1) and (M_2, ω_2) be symplectic manifolds. Then the product $M_1 \times M_2$ admits a symplectic structure via the “twisted” form

$$\tilde{\omega} = (pr_1)^*\omega_1 - (pr_2)^*\omega_2$$

where pr_i is the projection onto the i^{th} coordinate. Note that $\tilde{\omega}$ is closed because the exterior derivative commutes with pullbacks. It is non-degenerate because its

top-power is non-zero,

$$\tilde{\omega}^{2n} = \binom{2n}{n} ((pr_1)^*\omega_1)^n \wedge ((pr_2)^*\omega_2)^n \neq 0.$$

If $\phi: M_1 \rightarrow M_2$ is a diffeomorphism, then we can consider the graph

$$\Gamma_\phi = \{(p, \phi(p)) \mid p \in M_1\}$$

as a $2n$ -submanifold of the $4n$ -manifold $M_1 \times M_2$.

Proposition 1.20. *A diffeomorphism $\phi: M_1 \rightarrow M_2$ is a symplectomorphism if and only if its graph Γ_ϕ is a Lagrangian submanifold of $(M_1 \times M_2, \tilde{\omega})$.*

Proof. Let $i: M_1 \rightarrow M_1 \times M_2$ denote the embedding of Γ_ϕ , i.e. $i(p) = (p, \phi(p))$. It suffices to show that $i^*\tilde{\omega} = 0$. But

$$i^*\tilde{\omega} = i^*pr_1^*\omega_1 - i^*pr_2^*\omega_2 = (pr_1 \circ i)^*\omega_1 - (pr_2 \circ i)^*\omega_2 = \text{id}_M^*\omega_1 - \phi^*\omega_2,$$

and so $i^*\tilde{\omega} = 0$ if and only if $\phi^*\omega_2 = \omega_1$. □

1.3.1 Generating symplectomorphisms of cotangent bundles

Proposition 1.20 combined with Proposition 1.18 gives us a way to construct symplectomorphisms of cotangent bundles. If $f: M_1 \times M_2 \rightarrow \mathbb{R}$ is a smooth function, then $df: M_1 \times M_2 \rightarrow T^*(M_1 \times M_2) \cong T^*M_1 \times T^*M_2$ is a closed 1-form and so its

graph Γ_{df} is a Lagrangian submanifold of $T^*M_1 \times T^*M_2$. If we can somehow construct a diffeomorphism $T^*M_1 \rightarrow T^*M_2$ with the right graph then Proposition 1.20 tells us that we actually have a symplectomorphism.

There is one important issue here which we have to contend with: Proposition 1.18 exhibits Γ_{df} as a Lagrangian submanifold of $T^*(M_1 \times M_2)$ where the symplectic structure is given by $\omega := \text{pr}_1^*\omega_1 + \text{pr}_2^*\omega_2$. But Proposition 1.20 requires a Lagrangian submanifold of $T^*(M_1 \times M_2)$ with the *twisted* symplectic structure $\tilde{\omega} = \text{pr}_1^*\omega_1 - \text{pr}_2^*\omega_2$. This prompts us to introduce the twist map

$$\begin{aligned} \sigma: T^*M_1 \times T^*M_2 &\rightarrow T^*M_1 \times T^*M_2 \\ (x, \xi; y, \zeta) &\mapsto (x, \xi; y, -\zeta). \end{aligned}$$

Note that this map is a smooth involution, and hence a diffeomorphism. Moreover, it is a symplectomorphism because $\sigma^*\tilde{\omega} = \omega$, which can be shown by computing $\sigma^*\alpha = \tilde{\alpha}$ (where α is the tautological 1-form and $\tilde{\alpha}$ is the twisted version) in local

coordinates:

$$\begin{aligned}
\sigma^* \alpha &= \sigma^* \left(\sum_i \xi_i dx_i + \sum_j \zeta_j dy_j \right) \\
&= \sum_i \xi_i dx_i + \sum_j (\zeta_j \circ \sigma_2) d(y_j \circ \sigma_2) \\
&= \sum_i \xi_i dx_i + - \sum_j \zeta_j dy_j \\
&= \tilde{\alpha}.
\end{aligned}$$

Since d commutes with pullback, $\sigma^* \omega = \sigma^*(-d\alpha) = -d\sigma^* \alpha = -d\tilde{\alpha} = \tilde{\omega}$.

Since σ is a symplectomorphism, it will preserve Lagrangian submanifolds. This implies that the twisted graph

$$\Gamma_{df}^\sigma := \sigma(\Gamma_{df}) = \{(x, y, d_x f, -d_y f) \mid (x, y) \in M_1 \times M_2\}$$

will be a Lagrangian submanifold of $(T^*(M_1 \times M_2), \tilde{\omega})$. (Here $d_x f$ is $(df)_{(x,y)}$ projected to $T_x^* M_1 \times \{0\}$, and similarly for $d_y f$.) Our goal now is to construct a diffeomorphism $\phi: T^* M_1 \rightarrow T^* M_2$ whose graph is Γ_{df}^σ . If we can do so, then Proposition 1.20 lets us conclude that ϕ is a symplectomorphism; in this case we call ϕ the *symplectomorphism generated by f* and f the *generating function* of ϕ .

Working locally, finding such a ϕ amounts to solving particular “Hamilton equations” (which we will discuss in more detail in Chapter 3). Asking that $\Gamma_\phi = \Gamma_{df}^\sigma$ is

the same as asking for

$$\phi(x, \xi) = (y, \zeta) \text{ if and only if } \xi = d_x f \text{ and } \zeta = -d_y f.$$

Given a point $(x, \xi) \in T^*M_1$, we want to find (y, ζ) so that $y = \phi_1(x, \xi)$ and $\zeta = \phi_2(x, \xi)$. This means we need to solve

$$\xi_i = \frac{\partial f}{\partial x_i}(x, y) \text{ and } \zeta_j = -\frac{\partial f}{\partial y_j}(x, y).$$

Note that if we can solve the first equation, i.e. write $y = \phi_1(x, \xi)$, then we can also solve the second equation by plugging in this first solution to obtain $\zeta = \phi_2(x, \xi)$. All this to say, it suffices to understand when we can solve $\xi_i = \frac{\partial f}{\partial x_i}(x, y)$.

By the implicit function theorem, we can solve $\xi_i = \frac{\partial f}{\partial x_i}(x, y)$ locally for y in terms of x and ξ , provided that

$$\det \left[\frac{\partial}{\partial y_j} \left(\frac{\partial f}{\partial x_i} \right) \right]_{i,j=1}^n \neq 0$$

holds. This local condition is necessary for f to generate ϕ , and it is also locally sufficient; globally, there still may be a bijectivity issue.

1.4 Some Theorems

Many fundamental results of symplectic geometry have the following flavor: we start with a submanifold N in (M, ω) and find a neighborhood of it in M which is symplectomorphic to a “normal neighborhood” in some other symplectic manifold. We have already seen an example of this when $N = \{p\}$ is a point — Darboux’s Theorem tells us we can find a neighborhood symplectomorphic to $(\mathbb{R}^{2n}, \omega_0)$. More generally, if we have two neighborhoods of N which have different symplectic structures, we can look for a symplectomorphism between them which fixes N . Both the Moser and Weinstein theorems provide results in this vein. Since our focus is developing the necessary background for Lagrangian Floer homology, we will not provide detailed proofs of these important theorems and instead point the reader to their favorite symplectic geometry resource (the ones we use include [dS06], [McD98], and [MS95]).

1.4.1 Moser theorems

The Moser theorem helps us compare two different symplectic structures ω_0 and ω_1 on the same $2n$ -manifold M . Given these two symplectic structures, the natural question to ask is whether they are symplectomorphic, i.e. if there is a diffeomorphism $\phi: M \rightarrow M$ with $\phi^*\omega_1 = \omega_0$. If there is such a ϕ , we can then ask about its homotopical properties — in particular, can we find a ϕ which is homotopic to the identity id_M ? Or, even stronger, can we find a homotopy through diffeomorphisms? We call such a homotopy $\rho: M \times I \rightarrow M$ an *isotopy*; that is, ρ is an isotopy if each ρ_t is a

diffeomorphism and $\rho_0 = \text{id}_M$.

An first observation is that a symplectomorphism ϕ cannot be isotopic to the identity if $[\omega_0] \neq [\omega_1]$ in $H^2(M, \mathbb{R})$. This is because $\phi \sim \text{id}_M$ implies there is a homotopy operator H on the de Rham complex² so that $\text{id}_M^* \omega_1 - \phi^* \omega_1 = dH\omega_1 + Hd\omega_1$. Since ω_1 is closed, we have

$$\omega_1 = \phi^* \omega_1 + dH\omega_1 = \omega_0 + dH\omega_1$$

and so $[\omega_1] = [\omega_0]$ in cohomology. Is this necessary condition sufficient? The Moser theorem tells us *yes*, under certain conditions; in general, the answer is *no* (cf. [MS95, Example 7.23]).

Definition 1.21. We say that (M, ω_0) and (M, ω_1) are

1. *strongly isotopic* if there is an isotopy $\rho_t: M \rightarrow M$ such that $\rho_1^* \omega_1 = \omega_0$;
2. *deformation-equivalent* if ω_0 and ω_1 can be joined by a smooth family ω_t of symplectic forms;
3. *isotopic* if they are deformation-equivalent and the cohomology class $[\omega_t]$ is independent of t .

Note that strongly isotopic implies symplectomorphic ($\phi := \rho_1$), strongly isotopic implies isotopic ($\omega_t := \rho_t^* \omega_1$ joins ω_1 to ω_0 and $[\omega_t] = [\omega_1]$ by homotopy invariance

²Recall that a de Rham *homotopy operator* between ρ_0 and ρ_1 is a linear map $H: \Lambda^k(M) \rightarrow \Lambda^{k-1}(M)$ satisfying the *homotopy formula* $\rho_1^* - \rho_0^* = dH + Hd$.

of de Rham cohomology), and isotopic implies deformation equivalent. The Moser theorem will show that isotopic implies strongly isotopic on compact manifolds, so we have a chain of implications

$$\begin{array}{ccc} \text{strongly isotopic} & \implies & \text{symplectomorphic} \\ \updownarrow & & \\ \text{isotopic} & \implies & \text{deformation equivalent} \end{array}$$

whenever M is compact.

Theorem 1.22. *Let M be a compact manifold with two symplectic forms ω_0 and ω_1 . Suppose that ω_t is a smooth family of closed, non-degenerate 2-forms, $0 \leq t \leq 1$, joining ω_0 and ω_1 , so that $[\omega_t]$ is independent of t , i.e. $\frac{d}{dt}[\omega_t] = [\frac{d}{dt}\omega_t] = 0$. Then there is an isotopy $\rho: M \times \mathbb{R} \rightarrow M$ so that $\rho_t^*\omega_t = \omega_0$ for $0 \leq t \leq 1$.*

Note that we may take $\omega_t = (1 - t)\omega_0 + t\omega_1$ whenever this form is symplectic for all t . This theorem says that if (M, ω_0) and (M, ω_1) are isotopic via ω_t , then they are strongly isotopic via ρ_t . Said another way, when we look at the map which sends a symplectic form on M to its class in $H^2(M, \mathbb{R})$, all symplectic forms living in the same path-connected component of a fiber are symplectomorphic.

The proof of the Moser Theorem involves something called the *Moser trick* [Mos65], which is a method that turns out to be extremely useful in many situations. The idea is to construct the isotopy as the flow of a time-dependent vector field X_t , which can be found by solving a particular equation involving the (non-degenerate!) ω_t . We can apply the same ideas to neighborhoods of submanifolds of M to get a relative

version of the Moser theorem. We will return to the Moser trick later in Section 3.1, in particular in Remark 3.4, after we have discussed flows of time-dependent vector fields in more detail.

Theorem 1.23. *Let M be a manifold, and $i: N \rightarrow M$ the inclusion of a compact submanifold. Suppose ω_0 and ω_1 are symplectic forms in M so that $\omega_0|_p = \omega_1|_p$ for all $p \in N$. Then there exist neighborhoods U_0 and U_1 of N in M and a diffeomorphism $\phi: U_0 \rightarrow U_1$ so that*

$$\begin{array}{ccc} & X & \\ & \swarrow & \searrow \\ U_0 & \xrightarrow{\phi} & U_1 \end{array}$$

commutes and $\phi^\omega_1 = \omega_0$.*

If we take $N = \{p\}$ to be a point, then we can use the relative Moser theorem to prove Darboux's Theorem (Theorem 1.11). That is, take $p \in M$ and let $\phi: U \rightarrow \mathbb{R}^{2n}$ be a local chart centered at p . We want to show that $\phi^*\omega$ is equal to the standard form ω_0 on a neighborhood of p . We saw in Section 1.1 that we can choose ϕ so that $\phi^*\omega = \omega_0$ at p , and hence the result follows by the theorem above.

Both Darboux's Theorem and Moser's Theorem can be viewed as special cases of a more general problem. Suppose we have a $2n$ -manifold M and k -submanifold N , along with two neighborhoods U_0, U_1 of N in M . Given two symplectic forms ω_0, ω_1 on these neighborhoods, can we find a symplectomorphism of M preserving N ? That is, is there a diffeomorphism $U_0 \rightarrow U_1$ with $\phi^*\omega_1 = \omega_0$ $\phi(N) = N$? Darboux's Theorem answers this when N is just a point, and Moser's Theorem lies at the other

extreme when N is all of M .

1.4.2 Weinstein Lagrangian Neighborhood Theorem

The Weinstein neighborhood theorems are like symplectic versions of the tubular neighborhood theorem. The first version of the theorem says that if N is Lagrangian in M in two different ways, then there is a symplectomorphism connecting the two.

Theorem 1.24. *Let $i: N \rightarrow M$ be the inclusion of a compact n -submanifold, and suppose we have two symplectic forms ω_0 and ω_1 on M so that N is a Lagrangian submanifold of both (M, ω_0) and (M, ω_1) (i.e. $i^*\omega_0 = i^*\omega_1 = 0$). Then there exist neighborhoods U_0 and U_1 of N in M and a diffeomorphism $\phi: U_0 \rightarrow U_1$ so that*

$$\begin{array}{ccc} U_0 & \xrightarrow{\phi} & U_1 \\ & \swarrow i & \nearrow i \\ & N & \end{array}$$

and $\phi^*\omega_1 = \omega_0$.

There is a generalization of this theorem which is called the Coisotropic Embedding Theorem, which says that we may conclude the same result if N is instead assumed to be coisotropic with $i^*\omega_0 = i^*\omega_1$. The Weinstein Lagrangian Neighborhood Theorem helps us to classify Lagrangian embeddings via the following theorem.

Theorem 1.25. *Let $i: N \rightarrow M$ be the inclusion of a compact Lagrangian submanifold of (M, ω) . Let $z: N \rightarrow T^*N$ denote the zero section and ω_0 denote the canonical symplectic form on T^*N . Then there are neighborhoods U_0 of N in T^*N and U of N*

in M along with a diffeomorphism $\phi: U_0 \rightarrow U$ so that

$$\begin{array}{ccc}
 U_0 & \xrightarrow{\phi} & U \\
 & \swarrow z & \nearrow i \\
 & X &
 \end{array}$$

and $\phi^*\omega = \omega_0$.

This theorem says that the collection of Lagrangian embeddings is the same (up to symplectomorphism) as the collection of embeddings of manifolds into their cotangent bundles as zero sections. In this sense, all Lagrangian submanifolds look like the zero section from Example 1.17. There are also classifications of isotropic and coisotropic embeddings, due to Weinstein [Wei77, Wei81] and Gotay [Got82], respectively. The isotropic case is related to symplectic vector bundles and the coisotropic case is related to zero sections of a certain bundle which is dependent on the data of the embedding.

Chapter 2

Compatible Structures

As we have mentioned before, symplectic geometry occupies a liminal space between the rigidity of geometry and the flexibility of topology. Despite (or maybe because of) this, imposing additional structure on a symplectic manifold can result in some interesting geometry. Specifically, if we investigate the interaction of a Riemannian structure with our symplectic structure, then we will find an almost complex structure lurking in the shadows.

This interaction is basically happening at the level of linear algebra, because each of these three structures (symplectic, Riemannian, almost complex) involves a smoothly-varying requirement on the tangent spaces. That is, a symplectic structure ω on a manifold amounts to asking for a smoothly-varying symplectic structure on the tangent spaces (along with the global condition that $d\omega = 0$),³ a Riemannian structure

³If we do not require that ω is closed, we get an *almost symplectic* structure.

g is a smooth assignment of an inner product to each of the tangent spaces, and an almost-complex structure J is a smoothly varying complex vector space structure on the tangent spaces. The data of any two of (ω, g, J) determine the third in a compatible way, and form what is called a *compatible triple*. The relations can be summarized in the following table:

Given...	use...	to get...	and ask...
ω, J	$\omega(Ju, Jv) = \omega(u, v)$ $\omega(u, Ju) > 0$ for $u \neq 0$	$g(u, v) := \omega(u, Jv)$ positive inner product	is g flat?
g, J	J is orthogonal, $g(Ju, Jv) = g(u, v)$	$\omega(u, v) := g(Ju, v)$ non-degenerate 2-form	is ω closed?
ω, g	polar decomposition	J almost complex	is J integrable?

We will primarily deal with the third row, where we are given a symplectic manifold, choose a Riemannian structure (which is guaranteed to exist), and therefore determine an almost complex structure. After recalling the definition of Riemannian (Section 2.1) and almost complex structures (Section 2.2), we will describe compatibility at the level of vector spaces (Subsection 2.2.1) and at the level of manifolds (Section 2.3). Our primary reference for this section is [dS06], although this material may be found in any standard textbook.

2.1 Riemannian structures

Riemannian geometry is the geometry of a positive-definite symmetric bilinear form, which (in contrast to symplectic geometry) allows us to study more local notions like distance, angle, and curvature.

Definition 2.1. A *Riemannian metric* g on M is a smooth assignment of a positive

inner product g_p on T_pM for each $p \in M$. Here, smoothness means $p \mapsto g(X_p, X_p)$ is a smooth function on M for each vector field X .

We say (M, g) is a *Riemannian manifold*. The prototypical example is Euclidean space under the usual Euclidean inner product, but it turns out that every smooth manifold admits at least one Riemannian metric. One of the benefits of having such a metric is that we can talk about things like distance.

Definition 2.2. Let $\gamma: [a, b] \rightarrow M$ be a smooth curve on (M, g) . Define the *length* of γ to be

$$L(\gamma) = \int_a^b g_{\gamma(t)} \left(\frac{d\gamma}{dt}, \frac{d\gamma}{dt} \right)^{1/2} dt.$$

The (Riemannian) *distance* $d(p, q)$ between two points $p, q \in M$ is the infimum of $L(\gamma)$ over γ connecting p and q .

The length of a curve is independent of its parametrization. The curve γ is a *geodesic* if it locally minimizes distance and has constant velocity (i.e. no acceleration) — these curves are the analogy of “straight lines” in our (usually not-flat) manifold. A smooth curve joining p and q is a *minimizing geodesic* if it realizes the distance $d(p, q)$. A Riemannian manifold is *geodesically convex* if every two points are joined by a unique minimizing geodesic. Note that there are geodesics which are not minimizing geodesics: for instance, given any two points on S^2 , there are two geodesics which connect them (the two arcs of the great circle they reside on), but in general one of these great arcs will be shorter than the other. However, S^2 is not geodesically convex,

since there are two minimizing geodesics connecting every point with its antipode.

Geodesics are important for many reasons, one of which is that they minimize *energy* in addition to minimizing length. The *energy* (or *action*) of a curve γ is

$$E(\gamma) = \frac{1}{2} \int_a^b g_{\gamma(t)} \left(\frac{d\gamma}{dt}, \frac{d\gamma}{dt} \right) dt.$$

The comparison between $l(\gamma)$ and $E(\gamma)$ is akin to that of x^2 and $|x|$ (forgetting the integral for a moment): both x^2 and $|x|$ have the same minimum at 0, but x^2 is nicer to work with (i.e. its smooth) as opposed to $|x|$. Similarly, curve joining p to q will minimize the energy if and only if it is a minimizing geodesic. Hence if we want to minimize L we may as well minimize E (which is easier to work with analytically, due to the presence of the square).

2.1.1 Geodesic flow

One way that Riemannian structure interacts with symplectic structure is through geodesic flow. Note that any curve γ in M determines a curve in TM via $t \mapsto (\gamma(t), \gamma'(t))$. The vector field on TM whose integral curves are of this form is called the *geodesic field* on TM and its flow is called the *geodesic flow*. The metric g provides an identification of $T_p M$ with $T_p^* M$ by sending v to $g_p(v, -)$, and the corresponding flow on $T^* M$ is sometimes called the *cogeodesic flow*. Remarkably, this flow comes from a certain symplectomorphism of $T^* M$, which is in turn generated by a certain smooth function on $M \times M$.

Let $f: M \times M \rightarrow \mathbb{R}$ be given by $f(x, y) = -\frac{1}{2}d(x, y)^2$, and let $d_x f$ and $d_y f$ denote the components of $df_{(x,y)}$ with respect to $T_{(x,y)}^*(M \times M) \cong T_x^*M \times T_y^*M$. The discussion from Subsection 1.3.1 tells us that Γ_{df} is a Lagrangian submanifold of $T^*(M \times M)$ and that if

$$\Gamma_{df}^\sigma = \{(x, y, df_x, -df_y) \mid x, y \in M\}$$

is also the graph of diffeomorphism $\phi: T^*M \rightarrow T^*M$, then ϕ is actually a symplectomorphism. In this case, $\phi(x, \xi) = (y, \zeta)$ if and only if $\xi = d_x f$ and $\zeta = -d_y f$, meaning that we need to solve

$$\xi = d_x f \quad \text{and} \quad \zeta = -d_y f$$

for (y, ζ) in terms of (x, ξ) . Under the identification of tangent and cotangent spaces, we know $\xi = g_x(v, -)$ and $\zeta = g_y(w, -)$ for some $v \in T_x M$, $w \in T_y M$, so equivalently we want to solve

$$g_x(v, -) = d_x f \quad \text{and} \quad g_y(w, -) = d_y f$$

for (y, w) in terms of (x, v) . If γ is the (unique) geodesic with $\gamma(0) = x$ and $\gamma'(0) = v$ (this is often called $\exp(x, v)$), then $(y, w) = (\gamma(1), \gamma'(1))$ gives us our solution. That is, $\phi(x, g_x(v, -)) = (\gamma(1), g_{\gamma(1)}(\gamma'(1), -))$ is our symplectomorphism of T^*M .

2.2 Almost-complex structures

An almost-complex structure on a manifold basically means that its tangent spaces have the structure of complex vector spaces. The word “almost” is present to indicate

that this structure on tangent spaces is not enough to guarantee that the manifold itself has a complex structure. If the almost complex structure comes from a bonafide complex structure, it is said to be integrable, but this is difficult to check in general.

2.2.1 Linear complex structures and compatible triples

We will briefly review linear complex structures and explore their interactions with linear symplectic and inner product structures. A complex structure essentially means that the vector space has an endomorphism which looks like “multiplication by $\sqrt{-1}$.”

Definition 2.3. A *complex structure* on a (real) vector space V is a linear map $J: V \rightarrow V$ with $J^2 = -\text{id}$. The pair (V, J) is called a *complex vector space*.

Example 2.4. Consider the standard symplectic vector space $(\mathbb{R}^{2n}, \Omega_0)$ with standard symplectic basis $e_1, \dots, e_n, f_1, \dots, f_n$, and define J_0 by $e_i \mapsto f_i, f_j \mapsto -e_j$. As matrices,

$$\Omega_0 = \begin{pmatrix} 0 & \text{id} \\ -\text{id} & 0 \end{pmatrix} \quad \text{and} \quad J_0 = \begin{pmatrix} 0 & -\text{id} \\ \text{id} & 0 \end{pmatrix}.$$

Note that $J_0^2 = -\text{id}$ and $G_0(-, -) := \Omega_0(-, J_0-)$ is a positive inner product on \mathbb{R}^{2n} .

This example gives us a natural isomorphism with \mathbb{C}^n . In general, a complex structure J is equivalent to a \mathbb{C} -vector space structure if we identify J with multiplication by $\sqrt{-1}$. If V also happens to have a symplectic structure Ω , then we can look at how J and Ω interact.

Definition 2.5. Let (V, Ω) be a symplectic vector space. A complex structure J on

V is *compatible* with Ω (or Ω -*compatible*) if

$$G_J(-, -) := \Omega(-, J-)$$

is a positive inner product on V . That is, J is Ω -compatible if and only if

- (symplectomorphism condition) $\Omega(Ju, Jv) = \Omega(u, v)$,
- (taming condition) $\Omega(v, Jv) > 0$ for all $v \neq 0$.

Proposition 2.6. *A complex structure J is Ω -compatible if and only if there is a symplectic basis $e_1, \dots, e_n, f_1, \dots, f_n$ for V with $f_j = Je_j$.*

Proof. First suppose we have a symplectic basis of this form. Recall that this means $\Omega(e_i, e_j) = \Omega(f_i, f_j) = 0$ and $\Omega(e_i, f_j) = \delta_{ij}$. The fact that J is Ω -compatible follows from a straightforward check of the two conditions above. For instance, the taming condition (on basis vectors) follows because $\Omega(e_i, Je_i) = \Omega(e_i, f_i) = 1$ and $\Omega(f_j, Jf_j) = \Omega(f_j, -e_j) = \Omega(e_j, f_j) = 1$. Checking the symplectomorphism condition is similar. □

An immediate corollary is that every symplectic linear structure (V, Ω) admits a compatible complex structure J : take a symplectic basis $e_1, \dots, e_n, f_1, \dots, f_n$ for (V, Ω) and define $J: V \rightarrow V$ by $Je_i = f_i$ and $Jf_j = -e_j$, just as in Example 2.4. Conversely, given a complex structure (V, J) , there is always a symplectic structure Ω so that J is Ω -compatible— Ω is defined by $\Omega(u, v) = G(Ju, v)$ for a positive inner product G such that $J^* = -J$.

Remark 2.7. There is a way to build a Ω -compatible almost-complex structure that does not depend on a choice of basis (which is important if we want to extend these definitions to manifolds). Choose a positive inner product G on V . By non-degeneracy of Ω and G , we have two isomorphisms

$$\tilde{\Omega} : v \mapsto \Omega(v, -) \quad \text{and} \quad \tilde{G} : v \mapsto G(v, -)$$

between V and V^* . Hence there is an invertible linear map $A : V \rightarrow V$ so that $G(Au, v) = \Omega(u, v)$. Moreover, A is skew-symmetric because $G(A^*v, u) = -G(u, Av) = G(Av, u) = \Omega(v, u) = -\Omega(u, v) = -G(Au, v)$. Thus $AA^* = -A^2$ is symmetric and positive-definite ($G(AA^*, u, v) = G(A^*u, A^*u) > 0$ for $u \neq 0$), and hence its diagonalization has positive eigenvalues λ_i . Rescaling the eigenvalues to $\sqrt{\lambda_i}$, we get a new matrix $\sqrt{AA^*}$ which is also symmetric and positive-definite. Define

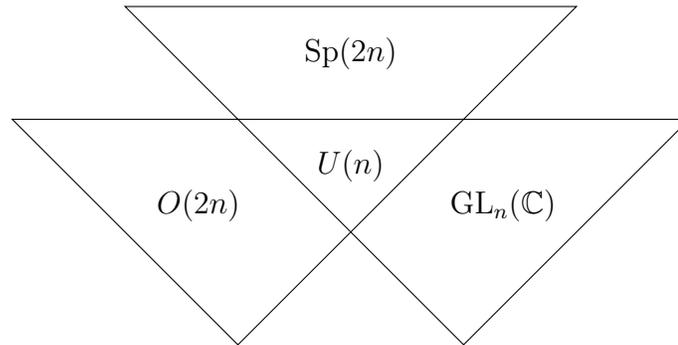
$$J := (\sqrt{AA^*})^{-1}A.$$

The factorization $A = \sqrt{AA^*}J$ is called the *polar decomposition* of A . By the properties of A and $\sqrt{AA^*}$, along with the fact that A commutes with $\sqrt{AA^*}$ (and hence so does J), it follows that J is skew-symmetric, orthogonal, and compatible with Ω .

Note that in general, the inner product defined by $\Omega(u, Jv) = G(\sqrt{AA^*}u, v)$ will be different than G , unless $AA^* = \text{id}$. The important observation is that this construction is canonical after the initial choice of G , which will let us extend these ideas

to manifolds.

Our recurring theme throughout this section is that once we choose any two of the three structures (Ω , G , or J), we determine the third. In the linear setting, this is because of the following picture of subgroups of $\mathrm{GL}_{2n}(\mathbb{R})$:



Recall that $Sp(2n)$ is the symplectic linear group of all linear transformations of \mathbb{R}^{2n} which preserve the standard symplectic structure. We identify an $n \times n$ matrix $A + iB$ with the real $2n \times 2n$ matrix $\begin{pmatrix} A & -B \\ B & A \end{pmatrix}$ in order to think of $\mathrm{GL}_n(\mathbb{C})$ as a subgroup of $\mathrm{GL}_{2n}(\mathbb{R})$. The picture above says that the intersection of any two of $Sp(2n)$, $O(2n)$, $\mathrm{GL}_n(\mathbb{C})$ is $U(n)$. To see this for a $2n \times 2n$ matrix $M = \begin{pmatrix} A & C \\ B & D \end{pmatrix}$, the first step is to unpack the subgroup requirements:

1. $M \in Sp(2n)$ if and only if it commutes with $\begin{pmatrix} 0 & -\mathrm{id}_n \\ \mathrm{id}_n & 0 \end{pmatrix}$;
2. $M \in O(2n)$ if and only if $M^T M = \mathrm{id}_n$;
3. $M \in \mathrm{GL}_n(\mathbb{C})$ if and only if $C = -B$ and $D = A$;

4. $M \in U(n)$ if and only if $M \in \text{GL}_n(\mathbb{C})$ and $(A + iB)^{-1} = (A + iB)^*$,

and then show that any two of 1–3 imply 4. For instance, if we assume 3 holds, then we can write $M^T = \begin{pmatrix} A & B \\ -B & A \end{pmatrix}$ which is precisely the conjugate transpose of $A + iB$. Hence 2 holds if and only if 4 does, i.e. $\text{GL}_n(\mathbb{C}) \cap \text{Sp}(2n) = U(n)$.

The upshot of this observation is that choosing a (compatible) complex structure for a symplectic vector space is equivalent to choosing an inner product structure. We will soon see that a similar situation arises on manifolds.

2.2.2 Almost-complex structures and integrability

Like Riemannian and symplectic structures, an almost complex structure happens on the level of tangent spaces.

Definition 2.8. An *almost-complex* structure on a manifold M is a smooth assignment $J: x \mapsto J_x$ of complex structures on tangent spaces,

$$J_x: T_x M \rightarrow T_x M \quad \text{and} \quad J_x^2 = -\text{id}.$$

The pair (M, J) is called an *almost-complex manifold*.

The name “almost-complex” invites us to ask *what’s the relationship with complex manifolds?* Every complex manifold (meaning a manifold which is locally homeomorphic to \mathbb{C}^n) admits a canonical almost-complex structure.

We will briefly mention how to construct the almost-complex structure locally. Let $(U, \phi: U \rightarrow \mathbb{C}^n)$ be a complex chart for M . Write the components of $\phi = (z_1, \dots, z_n)$ as $z_j = x_j + iy_j$. Then at $p \in U$, the tangent space $T_p M$ can be expressed as the \mathbb{R} -linear span of $\left. \frac{\partial}{\partial x_j} \right|_p, \left. \frac{\partial}{\partial y_j} \right|_p$. As we might guess, the correct definition of a complex structure J_p on $T_p M$ is given by

$$J_p \left(\left. \frac{\partial}{\partial x_j} \right|_p \right) = \left. \frac{\partial}{\partial y_j} \right|_p \quad \text{and} \quad J_p \left(\left. \frac{\partial}{\partial y_j} \right|_p \right) = - \left. \frac{\partial}{\partial x_j} \right|_p.$$

To check that this is globally well-defined, we would need to show that the local constructions agree on the overlap of two charts. We point the interested reader to [dS06, Proposition 15.2] for the complete details.

Although every complex structure induces an almost-complex structure, not every almost-complex structure can be upgraded to a complex one. In the special case that a given almost-complex structure does come from a complex structure, we give it a special name.

Definition 2.9. An almost complex structure is called *integrable* if it is induced by a complex manifold structure.

To detect whether an almost-complex structure (M, J) is integrable, we can use its *Nijenhuis tensor* \mathcal{N} which is defined as

$$\mathcal{N}(V, W) := [JV, JW] - J[V, JW] - J[JV, W] - [V, W]$$

for two vector fields V, W on M . Since \mathcal{N} is a tensor (basically because the bracket $[-, -]$ is), the value of $\mathcal{N}(V, W)$ at $p \in M$ depends only on the vectors $V_p, W_p \in T_pM$. In the 1950's, Newlander and Nirenberg [NN57] showed that J is integrable if and only if \mathcal{N} vanishes everywhere.

Example 2.10. We can use these ideas to show that orientable surface is a complex manifold. Recall from Example 1.12 that every orientable surface Σ is symplectic, and hence has a compatible almost-complex structure J . Since Σ has complex-dimension 1, for any $p \in \Sigma$ we can pick a nowhere-vanishing local vector field V so that $\{V_p, JV_p\}$ is a basis for $T_p\Sigma$. (Note that if $JV_p \in \text{span}\{V_p\}$ then $Jv = av$ for some $a \in \mathbb{R}$ but then $-V_p = J^2V_p = a^2V$ which is impossible.) Then

$$\begin{aligned} \mathcal{N}(V, JV) &= [JV, J^2V] - J[V, J^2V] - J[JV, JV] - [V, JV] \\ &= -[JV, V] - J[V, -V] - J[JV, JV] - [V, JV] \\ &= [V, JV] + J[V, V] - J[JV, JV] - [V, JV] \\ &= 0. \end{aligned}$$

Thus \mathcal{N} vanishes at every $p \in \Sigma$, which implies J is integrable.

2.3 Compatible triples

Just like with linear structures, Riemannian, symplectic, and almost complex structures on manifolds interact in a nice way. As before, a choice of any two of them will

determine a third compatible one.

Definition 2.11. An almost complex structure J on M is *compatible* with ω (or *ω -compatible*) if the assignment of $x \in M$ to $g_x: T_xM \times T_xM \rightarrow \mathbb{R}$ given by

$$g_x(u, v) := \omega_x(u, J_x v)$$

is a Riemannian metric on M . A triple (ω, g, J) is called a *compatible triple* if $g(-, -) = \omega(-, J-)$.

Note that if (ω, g, J) is a compatible triple, then any one of the structures may be expressed in terms of the others. Specifically, we have

$$g(u, v) = \omega(u, Jv)$$

$$\omega(u, v) = g(Ju, w)$$

$$J(u) = \tilde{g}^{-1}(\tilde{\omega}(u)).$$

Recall that $\tilde{g}, \tilde{\omega}: TM \rightarrow T^*M$ are the linear isomorphisms induced by the bilinear forms g, ω . In particular, since Riemannian metrics always exist on manifolds, a symplectic manifold always admits a compatible almost-complex structure.

Remark 2.12. For the bundle-minded reader, we note that the existence of these structures can be phrased in terms of reduction of the structure group of the tangent bundle $TM \rightarrow M$. Specifically, a $2n$ -manifold M is

- Riemannian if and only if TM is a $O(2n)$ -bundle,
- symplectic if and only if TM is a $Sp(2n)$ -bundle,
- almost-complex if and only if TM is a $GL_n(\mathbb{C})$ -bundle.

As discussed at the end of Subsection 2.2.1, any two of these conditions is equivalent to reducing the structure group to $U(n)$ (hence giving us the third condition as well).

Remark 2.13. If a ω -compatible almost-complex structure is integrable, then M is called a *Kähler manifold* and the symplectic form ω is called a *Kähler form*. Main examples of Kähler manifolds include compact Riemann surfaces, complex torii, and complex projective space. We have the following chain of implications:

$$\begin{array}{ccccc}
 \text{Kähler} & \implies & \text{symplectic} & & \\
 \Downarrow & & \Downarrow & & \\
 \text{complex} & \implies & \text{almost complex} & \implies & \text{smooth even-dimensional orientable}
 \end{array}$$

However in general we cannot trace backwards along these arrows. For instance, a manifold can be symplectic and complex without being Kähler (cf. [Thu76]), and there are almost-complex manifolds which are neither complex nor symplectic (cf. [Tau94]).

Given a symplectic manifold (M, ω) , we can investigate the set of compatible almost-complex structures; this set is path-connected in the following way.

Proposition 2.14. *Let (M, ω) be a symplectic manifold and J_0, J_1 two ω -compatible almost-complex structures. Then there is a smooth family $J_t, 0 \leq t \leq 1$, of compatible*

almost-complex structures joining J_0 to J_1 .

Proof. By compatibility, we get two Riemannian metrics $g_0(-, -) := \omega(-, J_0-)$ and $g_1(-, -) := \omega(-, J_1-)$. The linear homotopy $g_t = (1 - t)g_0 + tg_1$ gives us a smooth family of Riemannian metrics, from which we can obtain a smooth family of J_t 's using polar decomposition. \square

Similarly, if we have two symplectic structures ω_0 and ω_1 and an almost-complex structure J which is compatible with both, then we get a smooth family of symplectic⁴ forms $\omega_t = (1 - t)\omega_0 + t\omega_1$; i.e. ω_0 and ω_1 are deformation equivalent. The converse is not true in general, as two deformation equivalent symplectic structures may not share a compatible almost-complex structure (cf. [dS06, §13.3]).

Remark 2.15. Note that we have not relied on ω being closed, and so everything we have discussed may be extended to *almost-symplectic* manifolds (manifolds with a non-degenerate 2-form which is not necessarily closed).

⁴The form ω_t is closed since both ω_0, ω_1 are, and is non-degenerate since $g_t(-, -) = \omega_t(-, J)$ is positive (hence non-degenerate).

Chapter 3

Hamiltonians

In the introduction, we discussed how symplectic geometry is the natural abstract setting for Hamiltonian mechanics. This chapter will explore this connection in more detail (see, in particular, Example 3.7) with the goal of stating the Arnold Conjecture and understanding it as a statement about periodic solutions of a time-dependent Hamiltonian flow. We begin with a discussion of Hamiltonian functions and their corresponding vector fields and isotopies (Section 3.1) and state the Liouville-Arnold Theorem (Theorem 3.14). Then we define time-dependent Hamiltonians and formulate the Arnold Conjecture in Subsection 3.2.1.

3.1 Isotopies, vector fields, and Hamiltonian systems

Recall that an *isotopy* is a map $\rho: M \times \mathbb{R} \rightarrow M$ where each $\rho_t: M \rightarrow M$ is a diffeomorphism and $\rho_0 = \text{id}_M$. We can visualize an isotopy as an infinite cylinder

on M , with M itself sitting in the middle (corresponding to $\rho_0 = \text{id}_M$). Taking a horizontal time slice (i.e. picking a ρ_t) gives us a diffeomorphism of M , whereas taking a vertical time slice (fixing a point $p \in M$ and varying t) gives us a path in the cylinder. If we imagine ρ as defining a flow on M , then these paths are like the integral curves of this flow. This also defines a *time-dependent vector field* $v: M \times \mathbb{R} \rightarrow TM$. Each v_t is defined by the equation $\frac{d\rho_t}{dt} = v_t \circ \rho_t$, i.e.

$$v_t(p) = \left. \frac{d}{ds}(\rho_s \circ \rho_t^{-1})(p) \right|_{s=t}.$$

That is, $v_t(p)$ is the tangent vector of the integral curves defined by $\rho(p): \mathbb{R} \rightarrow M$ at time t . Visually, $v_t(p)$ points towards where in M the point p should flow under ρ .

Conversely, given a time-dependent vector field v , if the vector fields v_t are all compactly supported (for instance, when M is compact), then there exists an isotopy ρ which generates v , i.e. ρ satisfies the equation above. In fact, when M is compact, there is a one-to-one correspondence between time-dependent vector fields on M and isotopies of M .

Even if we cannot solve the differential equation globally for ρ , we can still solve it locally. That is, for any $p \in M$ and sufficiently small t , there is a one-parameter family of local diffeomorphisms ρ_t satisfying $\frac{d\rho_t}{dt} = v_t \circ \rho_t$ and $\rho_0 = \text{id}_M$.

Example 3.1. The nicest sort of time-dependent vector fields are the time-independent ones. If $v_t = v$ is independent of t , then the associated isotopy is called the *exponential map* and denoted $\exp tv$. This is the unique smooth family of diffeomorphisms

satisfying

$$\frac{d}{dt} \exp tv(p) = v(\exp tv(p)) \quad \text{and} \quad \exp tv|_{t=0} = \text{id}_M.$$

When working on a symplectic manifold (M, ω) , we want to understand how a time-dependent vector field will interact with our symplectic form. This motivates the *Lie derivative*, which can be defined more generally, but in our specific context will measure how differential forms change along the flow of v_t .

Definition 3.2. The *Lie derivative* by v_t is the operator $\mathcal{L}_{v_t}: \Lambda^k(M) \rightarrow \Lambda^k(M)$ defined by

$$\mathcal{L}_{v_t}(\omega) = \frac{d}{dt}(\rho_t)^*\omega|_{t=0}.$$

Note that when v_t is independent of t , then $\mathcal{L}_v\omega = \frac{d}{dt}(\exp tv)^*\omega|_{t=0}$.

The Lie derivative is closely related to the exterior derivative, and both capture the idea of taking a derivative, albeit in different ways. This difference is bridged by the *interior product* with respect to a vector field; in our case that vector field will be the time-dependent vector field v .

Definition 3.3. The *interior product* along v_t is $\iota_{v_t}: \Lambda^{k+1}(M) \rightarrow \Lambda^k(M)$ defined by

$$\iota_{v_t}(\omega)(X_1, \dots, X_k) = \omega(v_t, X_1, \dots, X_k).$$

The form $\iota_{v_t}\omega$ is sometimes called the *contraction* of ω with v_t .

That is, ι_{v_t} turns a $(k+1)$ -form into a k -form by sticking the vector field v_t in the

first slot. The identity that links ι_{v_t} , \mathcal{L}_{v_t} , and d is known as *Cartan's magic formula*:

$$\mathcal{L}_v \omega = \iota_v d\omega + d\iota_v \omega.$$

This formula says that the Lie derivative of ω along v_t is the same as a contraction of ω along v_t plus a correction term which takes into account the variation of v_t .

Remark 3.4. (Moser's trick.) Suppose we have a smooth family of k -forms α_t on M , $0 \leq t \leq 1$, and we want to find an isotopy $\phi_t: M \rightarrow M$ so that $\phi_t^* \alpha_t = \alpha_0$. Moser's trick is to try to construct ϕ_t as the time-1 flow of a time-dependent vector field X_t , and then solving $\phi_t^* \alpha_t = \alpha_0$ becomes equivalent to solving another, potentially easier equation. Given a time dependent vector field X_t , its flow satisfies

$$\begin{aligned} \frac{d}{dt} \rho_t^*(\alpha_t) &= \frac{d}{ds} \rho_s^* \alpha_t|_{s=t} + \frac{d}{ds} \rho_t^* \alpha_s|_{s=t} && \text{by the chain rule,} \\ &= \rho_t^*(\mathcal{L}_{X_t} \alpha_t) + \rho^* \frac{d\alpha_t}{dt} \\ &= \rho_t^*(\mathcal{L}_{X_t} \alpha_t + \frac{d\alpha_t}{dt}) \\ &= \rho_t^*(\frac{d\alpha_t}{dt} + d\iota_{X_t} \alpha_t + \iota_{X_t} d\alpha_t) && \text{by Cartan's magic formula.} \end{aligned}$$

Thus to prove that $\phi_t^* \alpha_t = \alpha_0$, it suffices to show that

$$0 = \frac{d}{dt} \phi_t^*(\alpha_t) = \phi_t^*(\frac{d\alpha_t}{dt} + d\iota_{X_t} \alpha_t + \iota_{X_t} d\alpha_t).$$

In many cases, solving $\frac{d\alpha_t}{dt} + d\iota_{X_t} \alpha_t + \iota_{X_t} d\alpha_t = 0$ may be much easier than solving

$\phi_t^* \alpha_t = \alpha_0$, and many important theorems in symplectic geometry are proved using the Moser trick, as discussed in Section 1.4.

Remark 3.5. Given a symplectic manifold (M, ω) we say a vector field X on M is *symplectic* if it preserves ω , meaning that $\mathcal{L}_X \omega = 0$. By Cartan's magic formula, this is equivalent to saying that $\iota_X \omega$ is closed.

3.1.1 Hamiltonian vector fields

The idea of Hamiltonian vector fields is that $\iota_X \omega$ is not only closed, but also exact. That is, our vector field (and corresponding isotopy) are generated by a smooth function, often called the *Hamiltonian* or *energy* functional. Let $H: M \rightarrow \mathbb{R}$ be a smooth function, so dH is a closed 1-form on M . By non-degeneracy of ω , there is a unique vector field X_H so that

$$\iota_{X_H} \omega = \omega(X_H, -) = dH.$$

Note that $p \in M$ is a critical point of H if and only if X_H vanishes at p , and moreover H is constant along the trajectories of X_H since

$$X_H(H) = dH(X_H) = \omega(X_H, X_H) = 0.$$

In this sense, X_H is like the “symplectic gradient” of H . The differential system associated to this vector field is

$$\frac{dx}{dt} = X_H(x(t)).$$

If the flow of X_H is complete (for instance if M is compact), then it generates an isotopy $\rho: M \times \mathbb{R} \rightarrow M$ which solves the differential system. Or equivalently, X_H generates a one-parameter family of diffeomorphisms $\rho_t: M \rightarrow M$ so that

$$\begin{cases} \rho_0 = \text{id}_M, \\ \frac{d\rho_t}{dt} = X_H(\rho_t). \end{cases}$$

Each diffeomorphism ρ_t preserves ω , in the sense that $\rho_t^*\omega = \omega$, because

$$\frac{d}{dt}\rho_t^*\omega = \rho_t^*\mathcal{L}_{X_H}\omega = \rho_t^*(\iota_{X_H}d\omega + d\iota_{X_H}\omega) = \rho_t^*(\iota_{X_H}0 + ddH) = 0$$

and $\rho_0^*\omega = \omega$ at $t = 0$. Thus every function on M yields a family of symplectomorphisms ρ_t . Note that we required both the non-degeneracy and closedness of ω to make this argument work.

Definition 3.6. The function H is called a *Hamiltonian function* and the vector field X_H its *Hamiltonian vector field*. The differential system associated to X_H is called a *Hamiltonian system*.

To say that a vector field X is Hamiltonian is precisely to say that $\iota_X\omega$ is exact.

We also saw in Remark 3.5 that $\iota_X\omega$ is closed if and only if X is symplectic ($\mathcal{L}_X\omega = 0$). This means that every Hamiltonian vector field is symplectic, but a symplectic vector field may not be Hamiltonian. The obstruction is measured by $H_{dR}^1(M)$; in particular, if $H_{dR}^1(M) = 0$, then every symplectic vector field is Hamiltonian.

Example 3.7. We can recover the classical “Hamilton’s equations” when we consider \mathbb{R}^{2n} with its standard symplectic structure ω_0 . In terms of coordinates (q_i, p_i) , a curve $\rho_t = (q(t), p(t))$ is an integral curve for a Hamiltonian vector field X_H precisely if it solves

$$\frac{dq_i}{dt} = \frac{\partial H}{\partial p_i} \quad \text{and} \quad \frac{dp_i}{dt} = -\frac{\partial H}{\partial q_i}.$$

In Euclidean space, we can give this a simple physical interpretation. In particular, if we take $H(q, p) = \frac{1}{2}|p|^2 + V(q)$ for some smooth $V: \mathbb{R}^n \rightarrow \mathbb{R}$ (which we think of as describing the potential energy), then

$$dH = \sum_i p_i dp_i + \frac{\partial V}{\partial q_i} dq_i.$$

The associated Hamiltonian vector field is $X_H = \sum_i -p_i \frac{\partial}{\partial q_i} + \frac{\partial V}{\partial q_i} \frac{\partial}{\partial p_i}$, and the Hamiltonian system is

$$\frac{dq_i}{dt} = p_i \quad \text{and} \quad \frac{dp_i}{dt} = -\frac{\partial V}{\partial q_i}.$$

This describes p as the speed of the particle (of mass $m = 1$) whose position is given by q . We can think of $\mathbb{R}^{2n} \cong T^*\mathbb{R}^n$ as the *phase space* for a system in \mathbb{R}^n : the first n coordinates given by q describe \mathbb{R}^n as a configuration space for the system and the

last n coordinates given by p record the dynamics of the system, i.e. how the system evolves over time.

More specifically, for $n = 3$, the Hamilton equations above are equivalent to Newton's second law, which states that a particle of mass m moving in \mathbb{R}^3 under the influence of a potential energy function $V: \mathbb{R}^3 \rightarrow \mathbb{R}$ will move along a curve $q(t)$ which solves

$$m \frac{d^2 q}{dt^2} = -\nabla V(q).$$

Here $q = (q_1, q_2, q_3)$ provides the coordinates of \mathbb{R}^3 . If we set $p_i = m \frac{dq_i}{dt}$ to be the *momenta* of the particle, then Newton's second law becomes

$$\begin{cases} \frac{dp_i}{dt} = m \frac{d^2 q_i}{dt^2} = -\frac{\partial V}{\partial q_i} = -\frac{\partial H}{\partial q_i} \\ \frac{dp_i}{dt} = \frac{1}{m} p_i = \frac{\partial H}{\partial p_i} \end{cases}$$

for energy function $H(q, p) = \frac{1}{2m} |p|^2 + V(q)$. But these are precisely Hamilton's equations for H on the phase space $T^*\mathbb{R}^3 \cong \mathbb{R}^6$.

Example 3.8. Given any vector field X on M , there is a unique vector field \tilde{X} on T^*M whose flow lifts the flow of X . This lift \tilde{X} is Hamiltonian with Hamilton $H = \iota_{\tilde{X}}\alpha$, which is to say that $d\iota_{\tilde{X}}\alpha = \omega(\tilde{X}, -) = \iota_{\tilde{X}}\omega$. To see this, note that Cartan's magic formula says

$$d\iota_{\tilde{X}}\alpha = \mathcal{L}_{\tilde{X}}\omega - \iota_{\tilde{X}}d\alpha = \mathcal{L}_{\tilde{X}}\omega + \iota_{\tilde{X}}\omega,$$

so it suffices to show $\mathcal{L}_{\tilde{X}}\omega = 0$. But this is true because we know $(\tilde{\rho}_t)^*\alpha = \alpha$ (and

hence $(\tilde{\rho}_t)^*\omega = \omega$) by Proposition 1.15, which implies

$$\mathcal{L}_{\tilde{X}}\omega = \frac{d}{dt}(\tilde{\rho}_t)^*\omega|_{t=0} = \frac{d}{dt}\omega|_{t=0} = 0.$$

Remark 3.9. If (g, ω, J) is a compatible triple on M and $H: M \rightarrow \mathbb{R}$ a Hamiltonian function, we can associate two different vector fields to H , namely, the gradient ∇H and the Hamiltonian vector field X_H . Then

$$\omega(X_H, Y) = dH(Y) = g(\nabla H, Y) = \omega(J\nabla H, Y)$$

for all vector fields Y , which means that $X_H = J\nabla H$ and $\nabla H = -JX_H$.

We can also interpret Hamiltonian vector fields in terms of \mathbb{R} -actions on our manifold. Recall that a (left) *action* of a Lie group G on a manifold M is a group homomorphism $\psi: G \rightarrow \text{Diff}(M)$. The action is said to be *smooth* if the evaluation map

$$\begin{aligned} ev_\psi: G \times M &\rightarrow M \\ (g, p) &\mapsto \psi_g(p) \end{aligned}$$

is a smooth map of manifolds. If M has a symplectic structure ω , then the action is called *symplectic* (or G acts by *symplectomorphisms*) if the smooth action ψ actually lands in $\text{Sym}(M, \omega) \subseteq \text{Diff}(M)$.

Example 3.10. If $G = \mathbb{R}$, then a smooth action of \mathbb{R} on M is equivalent to a complete vector field on M . If X is a complete vector field with flow $\Phi: M \times \mathbb{R} \rightarrow M$, then the smooth action of \mathbb{R} is given by flowing, i.e. $t \mapsto \Phi(-, t)$. Conversely, if ψ is a smooth \mathbb{R} -action, then we can define a complete vector field X by $X_p = \left. \frac{d\psi_t(p)}{dt} \right|_{t=0}$.

This method also gives us a one-to-one correspondence between symplectic actions of \mathbb{R} and complete symplectic vector fields on M . If a symplectic \mathbb{R} -action generates a Hamiltonian vector field, then we say the action itself is *Hamiltonian*. That is, we say ψ is Hamiltonian if its vector field X satisfies $\iota_X \omega = dH$ for some $H: M \rightarrow \mathbb{R}$.

Example 3.11. Consider \mathbb{R}^{2n} with its standard symplectic structure. The vector field $X = -\frac{\partial}{\partial x_1}$ is Hamiltonian (with Hamiltonian $H = x_1$), and so the action generated by X

$$\psi_t: (x_1, \dots, x_n, y_1, \dots, y_n) \mapsto (x_1 - t, \dots, x_n, y_1, \dots, y_n)$$

is Hamiltonian. The orbits of the action ψ look like lines parallel to the x_1 -axis.

This gives us a way to talk about symplectic and Hamiltonian actions of \mathbb{R} , but we would like to extend this to other Lie groups. This is immediately possible for S^1 , which is a quotient of \mathbb{R} , as we can just view an action of S^1 as a 2π -periodic action of \mathbb{R} .

Example 3.12. Consider S^2 with the standard symplectic form for cylindrical coordinates, $\omega = d\theta \wedge dz$. The \mathbb{R} -action $\psi_t(\theta, z) = (\theta + t, z)$, or equivalently the vector field $X = \frac{\partial}{\partial \theta}$, is 2π -periodic and hence descends to an action of S^1 . The action is just rotation by angle t around the z -axis, and so the orbits look like horizontal circles.

This is a Hamiltonian S^1 -action with Hamiltonian function $H = z$.

If G is a product of S^1 's and \mathbb{R} 's, then we can define a G -action to be Hamiltonian if its restriction to each factor is Hamiltonian with a Hamiltonian function which is invariant under the action of the other factors of G . For a more general Lie group, we need to soup up the definition of a Hamiltonian to something called a *moment map*. For more details, see [dS06, §22].

3.1.2 Integrable systems

Integrability of a system roughly means that its dynamics are very well-behaved, in the sense that there are sufficiently many *conserved quantities* (functions in the dependent variables whose value remains constant along trajectories). Although integrability is a non-generic quality, many (but not all) systems studied in physics are integrable.

Suppose (M, ω, H) is a Hamiltonian system. To detect conserve quantities, we want to know when a function $f: M \rightarrow \mathbb{R}$ is constant along trajectories of X_H . If we let ρ_t denote the flow of X_H , then

$$\begin{aligned}
 \frac{d}{dt}(f \circ \rho_t) &= \rho_t^* \mathcal{L}_{X_H} f && \text{by definition of } \mathcal{L}, \\
 &= \rho_t^* \iota_{X_H} df && \text{by Cartan's formula,} \\
 &= \rho_t^* \iota_{X_H} \iota_{X_f} \omega && \text{by definition of } X_f, \\
 &= \rho_t^* \omega(X_f, X_H) && \text{by definition of the interior product.}
 \end{aligned}$$

Thus f is constant along trajectories if and only if $\omega(X_f, X_H) = 0$. The function $\omega(X_f, X_H)$ is called the *Poisson bracket* of f and H , and is usually denoted $\{f, H\}$. The Poisson bracket gives $C^\infty(M)$ the structure of a Lie algebra (cf. [dS06, §18.3]). A function $f: M \rightarrow \mathbb{R}$ for which $\{f, H\} = 0$ is called a *conserved quantity* (or *integral of motion*).

A Hamiltonian system may admit multiple conserved quantities f_j , which will depend on H in general. If f_i and f_j commute with respect to the Poisson bracket, then it must be that $\{f_i, f_j\} = \omega(X_{f_i}, X_{f_j}) = 0$. This means that at each $p \in M$ the Hamiltonian vector fields generate an isotropic subspace of $T_p M$. If M is symplectic, this means there can be at most n independent⁵ conserved quantities.

Definition 3.13. A Hamiltonian system (M, ω, H) is (*completely*) *integrable* if it admits n functionally independent integrals of motion $f_1 = H, f_2, \dots, f_n$ so that $\{f_i, f_j\} = 0$ for all i, j .

As the reader might guess, there is a connection between integrable systems and Lagrangian submanifolds. Specifically, the level set of a regular value of $f := (f_1, \dots, f_n)$ will be a Lagrangian submanifold. The Liouville-Arnold Theorem describes this in more detail.

Theorem 3.14 (Liouville-Arnold Theorem). *Let (M, ω, H) be an integrable Hamiltonian system with integrals of motion $f_1 = H, \dots, f_n$. Define $f = (f_1, \dots, f_n)$ and let $c \in \mathbb{R}^n$ be a regular value of f . Then*

⁵Recall that functions on M are said to be *independent* if their differentials are pointwise linearly independent on an open dense subset of M .

- (i) $f^{-1}(c)$ is a Lagrangian submanifold of (M, ω) ,
- (ii) If the Hamiltonian flows X_{f_1}, \dots, X_{f_n} are complete on a connected component $L \subseteq f^{-1}(c)$, then L is diffeomorphic to $T^k \times \mathbb{R}^{n-k}$ for some k . There are “angle coordinates” $\theta_1, \dots, \theta_n$ on L for which the Hamiltonian equations are

$$\frac{d}{dt}\theta_i = v_i$$

for some $v_i \in \mathbb{R}$. The Hamiltonian flow on L is periodic if and only if $v_i \in \mathbb{Q}$ for all $i = 1, \dots, k$ and $v_i = 0$ for $i = k + 1, \dots, n$.

- (iii) If L is a connected component of $f^{-1}(c)$ then there are “action coordinates” $(\varphi_1, \dots, \varphi_n)$ which are integrals of motion for (M, ω, H) in a neighborhood of L . Along with the angle coordinates, they form a Darboux chart for L .

Note that a consequence of (ii) is that any compact connected component of $f^{-1}(c)$ is diffeomorphic to T^n , and such a torus is called a *Liouville torus*. The proof of (iii) makes use of the Weinstein Lagrangian neighborhood theorem (Theorem 1.25) and part (ii) to find a neighborhood U of L in M and a neighborhood U_0 of L in $T^*(L) \cong T^*(T^k \times \mathbb{R}^{n-k})$ along with a symplectomorphism $\phi: U_0 \rightarrow U$. This map ϕ gives a coordinate chart on U where the coordinates in the base space are the angle coordinates; the coordinates on the fibers are called the action coordinates. This gives a Darboux chart on U_0 which can be mapped to a Darboux chart on U (since ϕ is a symplectomorphism). For more details and proofs of the other parts, see [Arn78,

§50].

3.2 Time-dependent Hamiltonians and periodic solutions

So far we have developed the theory for a fixed Hamiltonian function, but in general we may want to allow our Hamiltonian to change over time. For example, dynamical processes in quantum mechanics may be described by a Hamiltonian that depends on time. Define a *time-dependent* Hamiltonian to be a smooth function $H: M \times \mathbb{R} \rightarrow \mathbb{R}$, and let H_t denote $H(-, t): M \rightarrow \mathbb{R}$. The Hamiltonian vector fields X_t are still defined by $X_t := X_{H_t}$, but the associated differential system is no longer autonomous:

$$\frac{dx}{dt} = X_t(x(t)).$$

However, just as before, the solutions of this system define a one-parameter family of diffeomorphisms ρ_t so that

$$\begin{cases} \rho_0 = \text{id}_M, \\ \frac{d\rho_t}{dt} = X_t(\rho_t). \end{cases}$$

The proof from the time-independent case also carries over to this setting to show that the ρ_t preserve the symplectic form ω .

Suppose that a time-dependent Hamiltonian is 1-periodic, meaning that $H_{t+1} = H_t$ for all $t \in \mathbb{R}$. This implies that X_t (and hence ρ_t) is also 1-periodic. Then there is a one-to-one correspondence between the fixed points of ρ_1 and the period-1 orbits of $\rho: M \times \mathbb{R} \rightarrow \mathbb{R}$. This is because $\rho_1(p) = p$ if and only if $\{x(t) = \rho(t, p) \mid t \in I\}$ is a

closed orbit.⁶ A periodic solution $x(t)$ (with $x(0) = x(1) = p$) is called *non-degenerate* if

$$\det(\text{id} - d\rho_1(p)) \neq 0.$$

This is to say that the differential $d\rho_1: T_pM \rightarrow T_pM$ does not have eigenvalue 1. In local coordinates, the matrix $(d\rho_1)_p$ is the Jacobian matrix $\exp(J_0(d^2\rho_1)_p(H))$. The non-degeneracy of p as a periodic orbit is equivalent to the invertibility of $\exp(J_0(d^2\rho_1)_p(H)) - \text{id}$, i.e. whether or not the Hessian $(d^2\rho_1)_p$ has eigenvalues in $2\pi\mathbb{Z}$.

Remark 3.15. In the special case that $H_t = H$ is time-*independent*, then it is certainly 1-periodic. In particular, all of the (non-degenerate) critical points of H are periodic trajectories and hence (non-degenerate) fixed points of ρ_1 . Moreover, if a critical point of H is non-degenerate as a periodic solution, then it is non-degenerate as a critical point of the function H (cf. [AD14, Proposition II.5.4.5]). Therefore ρ_1 has at least as many (non-degenerate) fixed points as H has (non-degenerate) critical points.

The Arnold Conjecture seeks to provide a similar bound to the fixed points of ρ_1 for a time-dependent 1-periodic Hamiltonian. The remainder of this thesis will be dedicated to stating the Arnold conjecture, reinterpreting it as a Lagrangian intersection problem, and explaining how Floer homology theory can help us approach this problem.

⁶Without assuming that H_t is one-periodic, we may conclude the “if” direction but not necessarily the “only if.”

3.2.1 The Arnold Conjecture

Let (M, ω) be a compact symplectic manifold and $f: M \rightarrow M$ a symplectomorphism. Suppose further that f is the time 1 flow of a 1-periodic Hamiltonian vector field. That is, there is a time-dependent Hamiltonian $H: M \times \mathbb{R} \rightarrow \mathbb{R}$ with $H_{t+1} = H_t$ which generates a one-parameter family of diffeomorphisms $\rho_t: M \rightarrow M$ so that $f = \rho_1$. Such a map f is often called a *Hamiltonian diffeomorphism*.

One formulation of the Arnold conjecture bounds the number of fixed points that f can have from below:

$$\#\{\text{fixed points of } f\} \geq \begin{array}{l} \text{minimal number of critical points} \\ \text{of a smooth function on } M \end{array}$$

and moreover

$$\#\{\text{non-degenerate fixed points of } f\} \geq \begin{array}{l} \text{minimal number of critical points} \\ \text{of a Morse function on } M \end{array}.$$

Using Morse theory (see Section 4.1), we also know that this last term is bounded below by the sum of the Betti numbers of M . The Arnold conjecture is a sharper result than the Lefschetz fixed point theorem, which provides a similar bound in terms of the alternating sum of the Betti numbers. Note that Remark 3.15 has already shown the Arnold conjecture holds for time-independent Hamiltonians.

There are several variants of the Arnold Conjecture, some of which are proven

and some of which remain open. The following version (using rational coefficients), which is now known to hold in full generality (cf. [Sal97, p.4]).

Theorem 3.16. *Let (M, ω) be a compact symplectic manifold and let f be the time-1 flow of a 1-periodic Hamiltonian function. Suppose that all critical points of f are non-degenerate. Then*

$$\#\{\text{fixed points of } f\} \geq \sum_{i=0}^{2n} \dim H_i(M; \mathbb{Q}).$$

It was first proven to hold for Riemannian surfaces by Eliashberg [Eli79] and then for the $2n$ -torus by Conley and Zehnder [CZ83]. The big breakthrough came in a series of papers by Floer [Flo88a, Flo89b, Flo89a], where he adopted techniques of Gromov and used the Morse-Smale-Witten complex to develop “infinite dimensional Morse theory” which became known as Floer homology. He was able to prove that the Arnold conjecture holds for a large class symplectic manifolds, and his ideas have since been extended to the general case by many others (see, for instance, the work of Hofer-Salamon [HS95], Ono [Ono95] and Fukaya-Ono [FO99], and Liu-Tian [LT98]). Most recently, Abouzaid-Blumberg [AB21] have shown that the Arnold conjecture holds for coefficients in a field of characteristic p by constructing a version of Floer homology with coefficients in Morava K -theory. In the next chapter, we will detail some of Floer’s ideas for his proof of the Arnold conjecture (using $\mathbb{Z}/2$ -coefficients) for a particular class of closed symplectic manifolds.

Chapter 4

Lagrangian Floer Homology

Floer homology was built to attack the Arnold conjecture, but its techniques have proven to be useful outside of this specific application. The purpose of this chapter is to outline the main ideas of the Floer complex, with an eye towards establishing the Arnold conjecture for $\mathbb{Z}/2$ -coefficients (under certain assumptions). The first observation is that we can recast the fixed point problem as an intersection problem: namely, the fixed points of a Hamiltonian diffeomorphism $\phi_1: M \rightarrow M$ will be in bijection with the intersection points of the diagonal $\Delta(M)$ and the graph Γ_{ϕ_1} , which are both Lagrangian submanifolds of the product $M \times M$. Thus we can approach the Arnold Conjecture through the lens of the more general problem of bounding $|L_0 \cap L_1|$, where L_0, L_1 are two Lagrangian submanifolds of some ambient symplectic manifold (M, ω) .

The idea of Floer homology for Lagrangian intersections is to construct a complex

generated by the intersection points $L_0 \cap L_1$. The differential is given by “counting” J -holomorphic strips which connect intersection points, where J is a suitably chosen almost complex structure on M . However, the differential is a very delicate thing and may not always be well-defined (in the sense of giving rise to a homology theory). Specifically, we need to rule out J -holomorphic “bubbles” which can pop up on these strips. This is done by imposing some restrictions on M , L_0 , and L_1 . Consequently, Floer homology for Lagrangian intersections may not always be able to be defined. In the case where L_1 is the image of L_0 under a Hamiltonian diffeomorphism, then we can build the Floer complex in a slightly different way, although we will run into many of the same analytic difficulties. The purpose of our exposition is not to delve into these difficulties, but rather provide an overview of the main ideas and storyline.

The inspiration for Floer homology comes from Morse theory; Floer homology is sometimes called “infinite-dimensional Morse homology.” Thus it will do us good to spend some time discussing and understand the ideas of Morse homology, which we will then generalize to the Floer setting in Section 4.2, albeit with some important technical adjustments (see Subsection 4.2.1). After defining the Floer complex in Subsection 4.2.2, we will sketch how Floer homology provides a solution to the Arnold conjecture (with some assumptions on our underlying symplectic manifold) in Section 4.3. Finally, we briefly survey the still-developing area of Floer homotopy theory, which admittedly was one of the author’s motivations for writing this thesis. Our primary references are Floer’s original research papers [Flo88a, Flo88b, Flo88c, Flo89a, Flo89b]

and others' later expositions including [Sal90, Sal97, Aur13, AD14].

4.1 Morse homology

Given a function $f: M \rightarrow \mathbb{R}$ on a manifold M , we can study the critical points (points $p \in M$ such that $(df)_p = 0$) of this function to reveal structural information about the underlying space. If we think of f as describing the topography of the manifold, critical points are the places where manifold is locally “flat.” We denote the collection of critical points of f by $\text{Crit}(f)$. A *Morse function* is $f: M \rightarrow \mathbb{R}$ for which each $p \in \text{Crit}(f)$ is non-degenerate, which means the Hessian d^2f_p is non-singular. We can classify different types of critical points based on the nearby behavior of the space, using the *Morse index*, which describes the number of linearly independent directions in which f is decreasing.

Definition 4.1. The *Morse index* $\text{ind}(p)$ of a critical point p is the index of $(d^2f)_p$, that is, the maximum dimension of a subspace upon which the Hessian at p is negative-definite.

The Morse index can be counted as the number of negative entries in the diagonalization of $(d^2f)_p$. A local minimum is thus a critical point with Morse index 0, while a local maximum has index equal to the dimension of the manifold. In addition to providing information about the manifold around a critical point, the Morse index also completely determines the behavior of f at this point, as seen through the Morse Lemma.

Lemma 4.2 (The Morse Lemma). *Given a critical point $p \in \text{Crit}(f)$, there is a chart $(\Omega(p), \phi)$ around p where $\phi: (\Omega(p), p) \rightarrow (\mathbb{R}, 0)$ with*

$$(f \circ \phi^{-1})(x_1, \dots, x_n) = f(p) + \sum_{i=1}^{\text{ind}(p)} x_i^2 - \sum_{i=\text{ind}(p)+1}^n x_i^2.$$

The neighborhood $\Omega(p)$ that appears in the lemma statement is called a *Morse chart*; these charts are discussed in far more detail in [AD14, §2.1]. An immediate corollary of the lemma is that critical points of a Morse function are isolated, and in particular a Morse function on a compact manifold can only have finitely many critical points.

In order to obtain information about the geometry of M from f , we can examine the gradient vector field.

Definition 4.3. On a Riemannian manifold (M, g) , recall that the *gradient* of $f: M \rightarrow \mathbb{R}$ at $p \in M$ is the unique vector field ∇f determined by

$$g(\nabla_p f, v) = (df)_p(v)$$

for every $v \in T_p M$. The integral curves of $-\nabla f$ are called the *gradient flow lines*, and they satisfy the differential equation

$$\frac{d\varphi}{dt} + \nabla_{\varphi} f = 0.$$

As we then expect, the flow lines tell us how to “descend” along f . Note that if the image of φ does contain a critical point p , then in fact it must be the constant curve $\varphi(t) = p$. Thus there are two kinds of flow lines, the constant flow lines at critical points and the flow lines that stay away from critical points (but may get arbitrarily close) along which f is strictly decreasing.

For any point $x \in M$, there is a unique gradient flow through x , which we call the *minimal* flow for x and denote φ_x . A crucial aspect of the gradient flow is that all the flows “start” and “end” at critical points, meaning that for any gradient flow line $\varphi: \mathbb{R} \rightarrow M$, there exist critical points $a, b \in \text{Crit}(f)$ such that

$$\lim_{t \rightarrow -\infty} \varphi(t) = a \quad \text{and} \quad \lim_{t \rightarrow \infty} \varphi(t) = b.$$

We say that $a =: s(\varphi)$ is the *starting point* and $b =: e(\varphi)$ is the *ending point* of φ . If we fix a critical point a , we can look at the collection of all points whose minimal flow ends at a (or alternatively those points whose minimal flow starts at a) to get the stable (or unstable) manifold of a .

Definition 4.4. Let $a \in \text{Crit}(f)$ and define its *stable manifold* to be

$$W^s(a) = \{x \in M \mid e(\varphi_x) = a\},$$

and its *unstable manifold* to be

$$W^u(a) = \{x \in M \mid s(\varphi_x) = a\}.$$

The stable and unstable manifolds help us formalize the idea of how the flow is “attracted to” and “repelled from” critical points. The stable and unstable manifolds are sometime referred to as the *ascending* and *descending* manifolds of a , respectively. The Stable Manifold Theorem from dynamical systems tells us that the stable and unstable manifolds of $a \in \text{Crit}(f)$ are submanifolds of M that are diffeomorphic to open disks, with

$$\dim(W^u(a)) = \text{codim}(W^s(a)) = \text{ind}(a).$$

Remark 4.5. Since the unstable manifolds of distinct critical points are disjoint and also $\coprod_{a \in \text{Crit}(f)} W^u(a)$ covers M , we can decompose M in terms of these submanifolds. This decomposition roughly resembles a CW complex, with one (open) k -cell for each critical point of index k . This decomposition is guaranteed to be a CW complex only when f is Morse-Smale (defined below).

Definition 4.6. The pair (f, g) (or just f if the choice of metric is clear) is *Morse-Smale* if the stable and unstable manifolds intersect transversely,

$$W^s(a) \pitchfork W^s(b)$$

for all $a, b \in \text{Crit}(f)$.

The Morse-Smale condition is a stability condition which is necessary for defining Morse homology (in particular, to ensure the differential is nilpotent) so we will now assume all of our Morse functions are Morse-Smale. This is not a huge restriction, as every Morse function can be perturbed in a certain sense to a Morse-Smale one (cf. [AD14, Theorem I.2.2.5]).

In the Morse-Smale case, the intersection $W^u(a) \cap W^s(b)$ is actually a submanifold of M which consists of all points on the trajectories connecting a to b ,

$$W(a, b) := \{x \in M \mid s(\varphi_x) = a \text{ and } e(\varphi_x) = b\}.$$

There is an evident \mathbb{R} -action on $W(a, b)$ by translations in time, and this action is free when $a \neq b$. Consequently, we can consider *moduli space of flows*

$$\mathcal{M}(a, b) := W(a, b)/\mathbb{R}.$$

which is a manifold of dimension $\text{ind}(a) - \text{ind}(b) - 1$. Note then that the Morse-Smale condition also prohibits flow lines between points with the same Morse index.

The Morse homology complex is basically formed from the data of the critical points (graded by Morse index) and the number of flows between them, i.e. the cardinality of $\mathcal{M}(a, b)$.

Definition 4.7 (Morse complex). Let $\text{Crit}_k(f)$ denote the set of critical points of f

with index k , and define

$$C_k(f) = \left\{ \sum_{a \in \text{Crit}_k(f)} m_a a \mid m_a \in \mathbb{Z}/2\mathbb{Z} \right\}.$$

The boundary operator $\partial_k : C_k(f) \rightarrow C_{k-1}(f)$ is define on a critical point a by

$$\partial_k(a) = \sum_{b \in \text{Crit}_{k-1}(f)} m_X(a, b)b,$$

where $m(a, b) \in \mathbb{Z}/2\mathbb{Z}$ is the number (mod 2) of trajectories going from a to b . In other words, $m(a, b)$ is the modulo 2 cardinality of $\mathcal{M}(a, b)$.

There are two obvious things that stand in our way of getting a well-defined homology theory: first, that $m(a, b)$ actually makes sense (i.e. that $\mathcal{M}(a, b)$ is finite whenever $\text{ind}(a) - \text{ind}(b) = 1$) and second that $\partial^2 = 0$. For the former, the idea is to show that $\mathcal{M}(a, b)$ is compact and hence must be a finite collection of points whenever $\text{ind}(a) - \text{ind}(b) = 1$ (because $\dim \mathcal{M}(a, b) = \text{ind}(a) - \text{ind}(b) - 1$). For the latter, we can directly calculate

$$\partial_{k-1}(\partial_k(a)) = \sum_{c \in \text{Crit}_{k-2}(f)} \left(\sum_{b \in \text{Crit}_{k-1}(f)} m(a, b)m(b, c) \right) c.$$

Therefore it suffices to show that the inner sum is zero. The main idea is to think of

this number as the cardinality of the disjoint union

$$\coprod_{b \in \text{Crit}_{k-1}(f)} W(a, b) \times W(b, c),$$

and show that this set of points is the boundary of a manifold of dimension 1. Since the boundary of a 1-manifold consists of an even number of points, and we are computing modulo 2, the desired result follows. We point the reader to [AD14, §3.1–3.2] for more details.

Remark 4.8. We can also define a Morse complex with \mathbb{Z} -coefficients, but keeping track of orientations makes things slightly more complicated. If we fix an orientation on the spaces of trajectories (by choosing orientations of the stable manifolds), then $\mathcal{M}(a, b)$ is an oriented compact manifold of dimension 0 whenever $\text{ind}(a) - \text{ind}(b) = 1$, and so is a finite number of points each with some \pm orientation. We now define $m(a, b)$ to be the sum of these signs, noting that this sum modulo 2 is the coefficient in the $\mathbb{Z}/2\mathbb{Z}$ definition, and define the chain complex and boundary operator just as above. Essentially the same proofs work to show that we get a well-defined homology.

The Morse homology groups $\text{HM}_*(M)$ are the homology groups of the Morse complex. It turns out that Morse homology is an ordinary homology theory, meaning it is isomorphic to the cellular (or simplicial or singular) homology of M . This result is rather remarkable, since it implies the Morse homology of M is independent of choice of Morse function f (and metric g). The benefit of working with Morse homology

is that we can pick our favorite Morse function f in such a way as to simplify our calculations, knowing that our final result does not depend on f . Moreover, standard results in homology will tell us something interesting about critical points of Morse functions. For instance, since $\mathrm{HM}_k(M)$ is a subgroup of $C_k(f)$, we know that the number of index k critical points of any Morse function is bounded below by the k^{th} Betti number of M ,

$$\# \mathrm{Crit}_k(f) \geq \dim \mathrm{HM}_k(M).$$

Summing over k , we see that $\# \mathrm{Crit}(f) \geq \sum_k \dim \mathrm{HM}_k(M)$; these are known as the *Morse inequalities*. This tells us, for instance, that any Morse function on S^2 must have at least two critical points (one minimum and one maximum). The Morse inequalities also showed up in our statement of the Arnold Conjecture in Subsection 3.2.1 and we will revisit this connection in Section 4.3.

4.2 Floer homology for Lagrangian intersections

The idea of Floer homology is to do Morse theory for a certain manifold and a certain Morse function on it. However, several complications will arise which we did not have to deal with in the Morse context. For the sake of exposition, and to avoid going into too many technicalities, we will impose some strong conditions on our symplectic manifold M in order to smooth over these new issues (although we will note along the way where we do so).

Let (M, ω) be a symplectic manifold and let L_0, L_1 be two compact Lagrangian

submanifolds which intersect transversally. Consider the path space

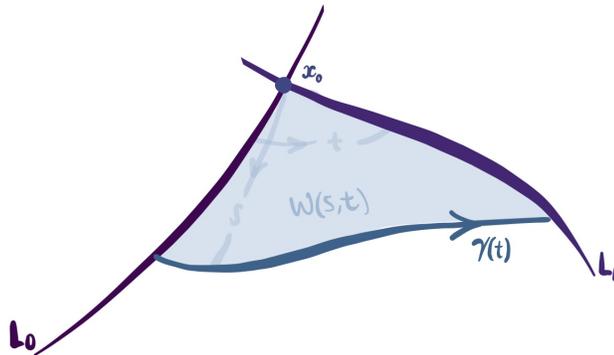
$$\Omega := \Omega(L_0, L_1) = \{\gamma: I \rightarrow M \text{ smooth} \mid \gamma(0) \in L_0, \gamma(1) \in L_1\},$$

and note that the intersection points $L_0 \cap L_1$ are precisely the constant paths in Ω .

We would like to define an action functional whose critical points are the constant paths, and to do so we pass to the universal cover

$$\tilde{\Omega} = \{[\gamma, w] \mid w: I^2 \rightarrow M \text{ homotopy } x_0 \sim \gamma\}$$

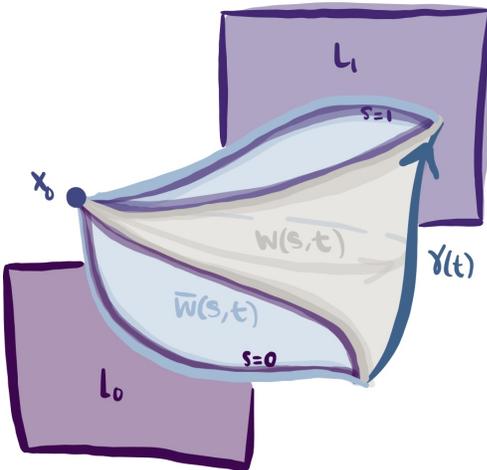
for a chosen basepoint $x_0 \in L_0 \cap L_1$ (viewed as a constant path). That is, w is a path in $\Omega(L_0, L_1)$ from x_0 to γ , so in particular $w(0, -) = x_0$ and $w(1, -) = \gamma$.



Now consider a candidate action functional $A: \tilde{\Omega} \rightarrow \mathbb{R}$ which takes the symplectic area of the square defined by the homotopy w ,

$$A([\gamma, w]) = \int_{w(I^2)} \omega = \int_{I^2} w^* \omega.$$

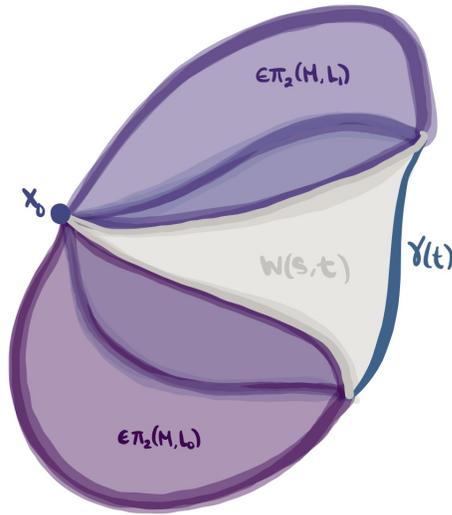
However, A may not be well-defined on $\tilde{\Omega}$. If w' is another homotopy from x_0 to γ , we need w and w' to have the same symplectic area in order for A to be well-defined. By concatenating w' with \bar{w} (which is w with reversed direction), we can define a map $w' \# \bar{w}: S^1 \times I \rightarrow M$, where $(w' \# \bar{w})(-, 0)$ is a loop in L_0 and $(w' \# \bar{w})(-, 1)$ is a loop in L_1 . Showing that w and w' have the same symplectic area is equivalent to showing that $w' \# \bar{w}$ has symplectic area 0; visually, this looks like a (hollow) cylinder in M without a top or bottom which is pinched at one side to the point x_0 .



However, this part of a cylinder cut out by $w' \# \bar{w}$ may not always have symplectic area 0 — in order to ensure this is true we either have to make some assumptions on L_0, L_1 , and M , or pass to a quotient of $\tilde{\Omega}$ called the *Novikov covering*. For the sake of simplifying our exposition, we will take the former approach.

The first thing we will assume is that we can extend $w' \# \bar{w}$ to a map $S^2 \rightarrow M$ without changing the symplectic volume. That is, we want to be able to fill in the loop $(w' \# \bar{w})(-, 0)$ in L_0 with a 2-disk, and similarly for the loop $(w' \# \bar{w})(-, 1)$ in

L_1 , but we do not want these disks to contribute anything to the integral. We can guarantee this to be possible by assuming that L_0, L_1 are simply connected or that $[\omega] \cdot \pi_2(M, L_i) = 0$.



We want to ensure the resulting map $S^2 \rightarrow M$ has zero symplectic area, which basically amounts to assuming that M is *symplectically aspherical*, i.e. that

$$\int_{S^2} f^*\omega = 0$$

for all smooth representative of $[f] \in \pi_2(M)$. Note that this holds trivially if M is aspherical, i.e. $\pi_2(M) = 0$.

Under these assumptions, we have

$$\begin{aligned}
 0 &= \int_{S^1 \times I} (w' \# \bar{w})^* \omega \\
 &= \int_{I^2} (w')^* \omega + \int_{I^2} \bar{w}^* \omega \\
 &= \int_{I^2} (w')^* \omega - \int_{I^2} w^* \omega,
 \end{aligned}$$

so A is well-defined on $\tilde{\Omega}$. In local coordinates, the action functional is given by

$$A([\gamma, w]) = \int_0^1 \int_0^1 \omega \left(\frac{\partial w}{\partial s}, \frac{\partial w}{\partial t} \right) ds dt.$$

In fact, this shows that A is well-defined on Ω (more precisely, the component of Ω containing x_0), rather than $\tilde{\Omega}$, since we may choose any homotopy w for our path γ . Thus we may develop our theory on Ω rather than $\tilde{\Omega}$ (see also Remark 4.10 for a justification that we may make all our definitions on Ω). Just to recap, our assumptions are:

Now, to follow Morse theory, we need to identify the critical points of A and equip Ω with a metric so that we can define gradient flow. The differential of A at γ is a map $T_\gamma \Omega \rightarrow \mathbb{R}$, where $T_\gamma \Omega$ consists of a path of tangent vectors $\xi: I \rightarrow T_{\gamma(t)} M$ so that $\xi(0) \in T_{\gamma(0)} L_0$ and $\xi(1) \in T_{\gamma(1)} L_1$.

Proposition 4.9. *The differential of the action functional is*

$$dA_\gamma(\xi) = \int_0^1 \omega(\xi(t), \gamma'(t)) dt$$

and γ is a critical point if and only if γ is a constant path.

Proof. We can compute $dA_\gamma(\xi)$ by taking a variation $\tilde{\gamma}_\varepsilon: I \rightarrow M$ of paths in Ω (where ε is defined in a small neighborhood of 0) so that $\tilde{\gamma}_0 = \gamma$ and $\frac{d}{d\varepsilon} \tilde{\gamma}_\varepsilon(t)|_{\varepsilon=0} = \xi(t)$. Then $dA_\gamma(\xi) = \frac{d}{d\varepsilon} A(\tilde{\gamma}_\varepsilon(t))|_{\varepsilon=0}$, and so to compute $A(\tilde{\gamma}_\varepsilon(t))$ we choose a lift $w(s, t)$ of γ and an extension $\tilde{w}_\varepsilon(s, t)$ so that $\tilde{w}_0 = w$ and $[\tilde{\gamma}_\varepsilon, \tilde{w}_\varepsilon]$ is in $\tilde{\Omega}$. Finally, we extend ξ to a vector field on w by defining $\xi(s, t) := \frac{d}{d\varepsilon} \tilde{w}_\varepsilon(s, t)|_{\varepsilon=0}$; note that this extension of ξ satisfies $\frac{d}{d\varepsilon} (\tilde{w}_\varepsilon)^* \omega|_{\varepsilon=0} = w^* \mathcal{L}_\xi \omega$. Now we can compute

$$\begin{aligned} dA_\gamma(\xi) &= \frac{d}{d\varepsilon} A(\tilde{\gamma}_\varepsilon(t))|_{\varepsilon=0} \\ &= \int_{I^2} \frac{d}{d\varepsilon} (\tilde{w}_\varepsilon)^* \omega|_{\varepsilon=0} \\ &= \int_{I^2} w^* (\mathcal{L}_\xi \omega) \\ &= \int_{I^2} w^* (d\iota_\xi \omega) && \text{by Cartan's magic formula,} \\ &= \int_{\partial I^2} w^* \iota_\xi \omega && \text{by Stokes' Theorem,} \\ &= \int_0^1 (w(0, t))^* \iota_{\xi(t)} \omega dt && \text{since } L_0, L_1 \text{ are Lagrangian,} \\ &= \int_0^1 \gamma(t)^* \iota_{\xi(t)} \omega dt \\ &= \int_0^1 \omega(\xi(t), \gamma'(t)) dt. \end{aligned}$$

Now, γ is a critical point if $dA_\gamma(\xi) \equiv 0$ for all ξ . Clearly this holds whenever γ is constant (since $\gamma'(t) \equiv 0$), and for the converse direction, if $dA_\gamma(\xi) = 0$ then $\omega(\xi(t), \gamma'(t)) \equiv 0$ and so by non-degeneracy of ω , the path γ is constant. \square

Remark 4.10. We can define $\alpha(\xi) := \int_0^1 \omega(\xi(t), \gamma'(t))dt$ for $\alpha \in T_\gamma\Omega$ on all of Ω , not just the component containing x_0 . However, α is only *locally* closed (with $\alpha = dA$), not globally.

To define the gradient of A we need to choose a Riemannian structure on Ω . Recall from Section 2.3 that a choice of almost-complex structure J on M defines a compatible Riemannian metric $g(-, -) := \omega(-, J-)$ on M . If J_t is a smooth family of almost complex structures, $0 \leq t \leq 1$, then we get a smooth family of metrics g_t on M . We can use g_t to define a metric on Ω given by

$$\langle \xi_1, \xi_2 \rangle := \int_0^1 g_t(\xi_1(t), \xi_2(t))dt$$

which we then use to define the “ L^2 -gradient” of A .

Proposition 4.11. *The gradient of A with respect to $\langle -, - \rangle$ is given by*

$$-\nabla_\gamma A = J_t \frac{d\gamma}{dt}.$$

Proof. Recall that ∇A is the vector field defined by

$$\langle \nabla_\gamma A, \xi \rangle = (dA)_\gamma(\xi)$$

for any $\xi \in T_\gamma\Omega$. Building on our previous computation,

$$\begin{aligned}
(dA)_\gamma(\xi) &= \int_0^1 \omega(\xi, \gamma') \\
&= \int_0^1 \omega(\xi, J_t(-J_t(\gamma'))) dt \\
&= \int_0^1 g_t(\xi, -J(\gamma')) dt \\
&= \langle \xi, -J_t(\gamma') \rangle \\
&= \langle -J_t(\gamma'), \xi \rangle
\end{aligned}$$

so by non-degeneracy of the inner product, $\nabla_\gamma A = -J_t \frac{d\gamma}{dt}$. \square

Since each J_t is an automorphism of TM for each t , the gradient vanishes if and only if $\frac{d\gamma}{dt} = 0$, i.e. γ is constant. We can compute the Hessian at a critical point $\gamma(t) \equiv x \in L_0 \cap L_1$ to be

$$(d^2A)_\gamma(\xi_1, \xi_2) = \int_0^1 \omega(\xi_1(t), \xi_2'(t)) dt$$

for $\xi_1, \xi_2: I \rightarrow T_pM$ vector fields along γ . A non-degenerate critical point of A corresponds precisely to a transverse intersection point of L_0 and L_1 , so if we ask that $L_0 \pitchfork L_1$ then A is Morse. However, there are some serious issues that arise when we try to continue to do Morse theory directly.

- The Hessian $(d^2A)_\gamma$ at a critical point has infinite Morse index, that is, it is a quadratic form with infinite-dimensional negative (and positive) subspaces.

Hence the critical points of A cannot be expected to tell us anything about the topology of Ω .

- We want to study the gradient flow of A on Ω , but $\nabla_\gamma A = J_t \gamma'$ may not be tangent to Ω . That is, $J_0 \gamma'(0)$ may not be in $T_{\gamma(0)} L_0$, and $J_1 \gamma'(1)$ may not be in $T_{\gamma(1)} L_1$. Since $\langle \xi_1, \xi_2 \rangle$ is an L^2 inner product, we only know that $\nabla_\gamma A$ lies in the L^2 -completion of $T_\gamma \Omega$.

To deal with the second issue, Floer's insight is to consider a PDE (rather than an ODE) to define the gradient flow using “ J -holomorphic strips.” This then lets him define a relative index depending on a path between critical points. We will discuss Floer's approach to these technical issues in the next subsection, and then define the Floer complex in Subsection 4.2.2.

4.2.1 The moduli space of J -holomorphic strips

Suppose we had a reasonable definition of a gradient flow line $\phi: \mathbb{R} \rightarrow \Omega$ so that $\frac{d\phi}{ds} = \nabla_\phi A = -J_t \frac{d\phi_s}{dt}$. Since ϕ is a path of paths in M , we can instead consider the map $u: \mathbb{R} \times I \rightarrow M$ so that $u(s, -) = \phi(s)$. The requirement that ϕ land in Ω means that $u(s, 0) \in L_0$ and $u(s, 1) \in L_1$ for all $s \in \mathbb{R}$. The gradient flow ODE now becomes a PDE

$$\frac{\partial u(s, t)}{\partial s} + J_t(u(s, t)) \frac{\partial u(s, t)}{\partial t} = 0.$$

This equation is in fact the Cauchy-Riemannian equation with respect to the ω -compatible almost-complex structure J_t on M and is often abbreviated as $\bar{\partial}_{J_t} u = 0$.

A map u which satisfies $\bar{\partial}_J u = 0$ is called a *J-holomorphic strip*,⁷ and it is these maps which we will use as our gradient flows. Additionally, in order to continue our Morse theory analogy, we will specifically look at *J-holomorphic strips* which connect critical points of A , i.e. constant paths.

Definition 4.12. A *gradient flow line* of A is a J_t -holomorphic strip in M , that is, a smooth map $u: \mathbb{R} \times I \rightarrow M$ such that

$$\frac{\partial u}{\partial s} + J_t \frac{\partial u}{\partial t} = 0, \quad \lim_{s \rightarrow +\infty} u(s, -) = y, \quad \lim_{s \rightarrow -\infty} u(s, -) = x, \quad (*)$$

for $x, y \in L_0 \cap L_1$.

Remark 4.13. If (M, ω) is exact ($\omega = d\alpha$ for some α) but not compact, then we impose the additional condition of *finite energy*,

$$\int_{\mathbb{R} \times I} u^* \omega < \infty.$$

On a compact manifold, a holomorphic strip u will have finite energy if and only if it satisfies the limit conditions in (*) (cf. [RS01]).

Definition 4.14. For $x, y \in L_0 \cap L_1$, the *moduli space* of *J-holomorphic strips* from x to y is the space $\mathcal{M}_J(x, y)$ which consists of maps $u: \mathbb{R} \times I \rightarrow M$ so that

- (i) (*J-holomorphic condition*) $\bar{\partial}_J u = 0$;

⁷If J is independent of t , then $\mathbb{R} \times i[0, 1] \subseteq \mathbb{C}$ is conformally equivalent to the unit disk in \mathbb{C} minus two points on the boundary. For this reason, a map u as above is sometimes called a *J-holomorphic (Whitney) disk*.

(ii) (boundary condition) $u(s, 0) \in L_0$ and $u(s, 1) \in L_1$ for all $s \in \mathbb{R}$;

(iii) (asymptotic condition) $\lim_{s \rightarrow \infty} u(s, -) = y$ and $\lim_{s \rightarrow -\infty} u(s, -) = x$.

Note that the equations $(*)$ are invariant under time translation in the first coordinate. That is, if $u(s, t)$ satisfies $(*)$ then so does $u(s + s_0, t)$ for any $s_0 \in \mathbb{R}$. To account for this overlap, we will consider $\mathcal{M}_J(x, y)$ quotiented by this \mathbb{R} -action

$$\tilde{\mathcal{M}}_J(x, y) := \mathcal{M}(x, y)/\mathbb{R}.$$

To view $\tilde{\mathcal{M}}_J(x, y)$ as a “space of solutions” rather than just a set, we think of it as the zero locus of a certain map. Following the summary in [Flo88a, Proposition 2.1], we will outline the main ideas but leave the explicit details for [Flo88c] (or see [AD14, Chapter II.8]).

Proposition 4.15. *For each $x, y \in L_0 \cap L_1$, there is a smooth Banach manifold $\mathcal{P}(x, y)$ of paths $u: \mathbb{R} \rightarrow \Omega$ so that $u \mapsto \bar{\partial}_J u$ is a smooth section of a Banach space bundle over $\mathcal{P}(x, y)$.*

We see then that $\bar{\partial}_J^{-1}(0) = \tilde{\mathcal{M}}(x, y)$. Properties of the operator $\bar{\partial}_J$ ensure it is a *Fredholm map*, which means that its linearization $D_u \bar{\partial}_J$ at u is a bounded linear operator between Banach spaces with finite-dimensional kernel and finite-codimensional image. In particular, this means the cokernel is finite-dimensional, and we define the

index of $\bar{\partial}_J$ at u to be the Fredholm index of $D_u\bar{\partial}_J$,

$$\text{ind}_u(\bar{\partial}_J) := \dim(\ker D_u\bar{\partial}_J) - \dim(\text{coker } D_u\bar{\partial}_J) = \dim(\ker D_u\bar{\partial}_J) - \text{codim}(\text{im } D_u\bar{\partial}_J).$$

One of the nice things about Fredholm maps is they have an excellent theory of differential topology. For example, the pre-image of any regular value (one where the linearization is surjective) of a Fredholm map will be a smooth manifold of dimension equal to the Fredholm index. In our setting, we want 0 to be a regular value of $\bar{\partial}_J$ so that we may conclude that $\tilde{\mathcal{M}}$ is a smooth finite-dimensional manifold. In [Flo88c], Floer proves the following result.

Proposition 4.16. *There is a dense set $\mathcal{J}(L_0, L_1)$ of smooth compatible almost-complex structures on M so that if $J \in \mathcal{J}(L_0, L_1)$, then $D_u\bar{\partial}_J$ is surjective Fredholm operator for all $u \in \tilde{\mathcal{M}}_J(x, y)$.*

An almost-complex structure that belongs to this set is called *regular*. This proposition says that any compatible almost-complex structure may be perturbed (in an arbitrarily small way) to a regular one, so that the space of “trajectories” connecting x to y is a finite-dimensional manifold.

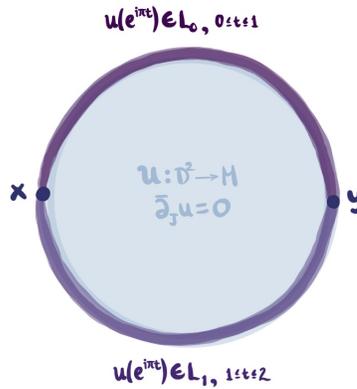
There is an evident dependency on choice of u connecting x and y , which turns out to only depend on the homotopy class of u ; this motivates us to define $\tilde{\mathcal{M}}_J(x, y; [u])$ to be the collection of all $\tilde{u} \in \tilde{\mathcal{M}}_J(x, y)$ with $[\tilde{u}] = [u]$. From Fredholm theory, we know that $\dim \tilde{\mathcal{M}}_J(x, y; [u])$ is the Fredholm index $\text{ind}_u(\bar{\partial}_J)$. Remarkably, the dimension of

$\tilde{\mathcal{M}}_J(x, y; [u])$ can also be computed in terms of a “relative Morse index,”

$$\dim \tilde{\mathcal{M}}(x, y; [u]) = \text{ind}_u(\bar{\partial}_J) = \mu(x, y; [u]).$$

This index μ is called the *Maslov index* or the *Conley-Zehnder index*, depending on the context (cf. [AD14, Chapter II.7]). Note that, unlike the Morse case, this index depends on *two* critical points x, y as well as the homotopy class of a chosen path u connecting x and y .

Briefly, the idea of the Maslov index is to assign a loop of Lagrangian subspaces of \mathbb{R}^{2n} to an integer. Let Λ_n denote the Grassmannian of Lagrangian subspaces of $(\mathbb{R}^{2n}, \omega_0)$, where ω_0 is the standard symplectic form (Example 1.10). We can use the data of a J -holomorphic strip u which connects x to y to define a closed loop in Λ_n . Since $\pi_1(\Lambda_n) \cong \mathbb{Z}$, this construction assigns u to an integer and this integer is called the *Maslov index* of u . Regard u as a map $D^2 \rightarrow M$ so that $u(\partial D^2) \in L_0 \cup L_1$ as follows:

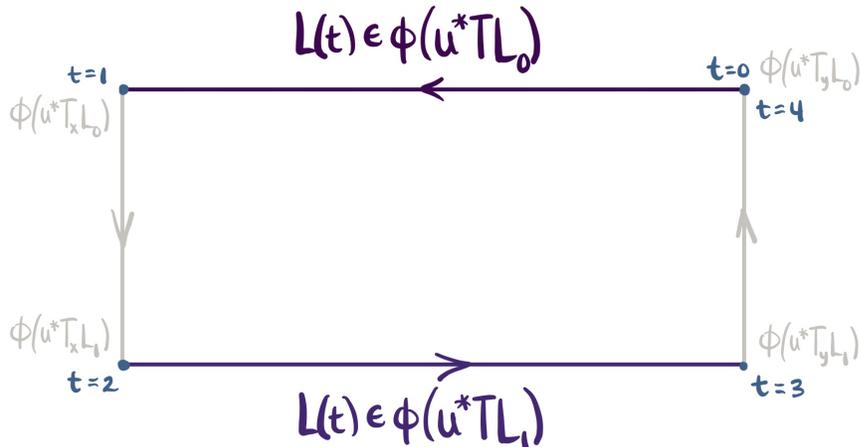


$$\begin{aligned}
u(1) &= y, & u(e^{i\pi t}) &\in L_0 \text{ for } t \in [0, 1], \\
u(-1) &= x, & u(e^{i\pi t}) &\in L_1 \text{ for } t \in [1, 2].
\end{aligned}$$

Then the pullback u^*TM is a symplectic bundle over D^2 , and hence it can be trivialized $u^*TM \cong \mathbb{R}^{2n} \times D^2$. We also get subbundles u^*TL_0 and u^*TL_1 over 1 and -1 , respectively. Using a chosen trivialization ϕ (which ultimately does not matter), we can define a loop $L: [0, 4] \rightarrow \Lambda_n$ so that

$$L(t) = \begin{cases} \phi(u^*T_{u(e^{i\pi t})}L_0) & t \in [0, 1], \\ \phi(u^*T_{u(e^{i\pi(t-1)})}L_1) & t \in [2, 3]. \end{cases}$$

To complete the rest of the loop, we want to find two paths in u^*TM , once connecting $u^*T_xL_0$ to $u^*T_xL_1$ and one connecting $u^*T_yL_0$ to $u^*T_yL_1$. Here is one approach to this (see [McD98, §3]): choose any path from $L(1)$ to $L(2)$, find a symplectic matrix $A \in \text{Sp}(n)$ so that $A(L(1)) = L(0)$ and $A(L(2)) = L(3)$, and then run the chosen path through A , i.e. set $L(3+s) = A(L(2-s))$ for $0 \leq s \leq 1$.



The Maslov index $\mu(x, y; [u])$ is defined to be the element of $\pi_1(\Lambda_n) \cong \mathbb{Z}$ represented by the homotopy class of L . Intuitively, the Maslov index is the (signed) count of the instances t at which $L(t)$ and $L(2+t)$ are *not* transverse to each other, for $0 \leq t \leq 1$. Conley and Zehnder [CZ84] generalize the Maslov index to paths in the symplectic group.

4.2.2 Floer complex for Lagrangian intersections

Unlike Morse homology, the definition of Floer homology is highly dependent on the specific attributes of (M, ω) and the selected Lagrangian submanifolds. Floer homology has been fully developed when the cohomology classes $[\omega]$ and c_1 (the first Chern class) satisfy

$$\int_{S^2} v^* c_1 = \tau \int_{S^2} v^* \omega$$

for all $v: S^2 \rightarrow M$ and some $\tau \in \mathbb{R}$ (cf. [Sal97]). The cases $\tau > 0$, $\tau = 0$, and $\tau < 0$ roughly correspond to conditions of positive, zero, and negative curvature. We will give an overview of the $\tau > 0$ (treated by Floer [Flo88a]) and $\tau = 0$ (treated by Hofer-Salamon [HS95]) cases. More general cases (including when $\tau < 0$) have been more recently resolved by [Ono95], [FO99], [LT98], among others. It is understood that Floer homology may not always be able to be defined in general.

Positive case. Floer developed his theory for *monotone* symplectic manifolds (when $\tau > 0$), which is by far the simplest setting. In this case, Floer shows in [Flo88b] that

there is an index function $\mu: L_0 \cap L_1 \rightarrow \mathbb{Z}$ so that

$$\mu(x, y; [u]) = \mu(x) - \mu(y).$$

This allows us to emulate Morse homology rather directly, acquiring a \mathbb{Z} -graded homology theory. In particular, following [Flo88a], we can define the Floer complex $\text{CF}(L_0, L_1)$ to be the free $\mathbb{Z}/2$ -module on $L_0 \cap L_1$; since L_0 and L_1 are compact and intersect transversely, this generating set is finite. The index function gives us a grading

$$\text{CF}(L_0, L_1) = \bigoplus_k \text{CF}_k(L_0, L_1)$$

where $\text{CF}_k(L_0, L_1)$ is generated by $x \in L_0 \cap L_1$ with $\mu(x) = k$. The differential is defined as in Morse theory,

$$\partial_J(x) = \sum_{\mu(y)=\mu(x)-1} m(x, y)y$$

where $m(x, y)$ is the mod-2 cardinality of $\tilde{\mathcal{M}}_J(x, y)$. The Floer homology is the \mathbb{Z} -graded homology of the chain complex $(\text{CF}(L_0, L_1), \partial_J)$,

$$\text{HF}_*(L_0, L_1) = \ker \partial / \text{im } \partial.$$

Remark 4.17. As with Morse theory, we could also develop Floer homology with \mathbb{Z} -coefficients (or \mathbb{Q} -coefficients), although we would need to take orientations of the

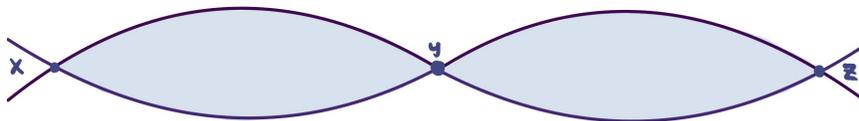
moduli spaces $\mathcal{M}_J(x, y)$ into account, which introduces new technical difficulties.

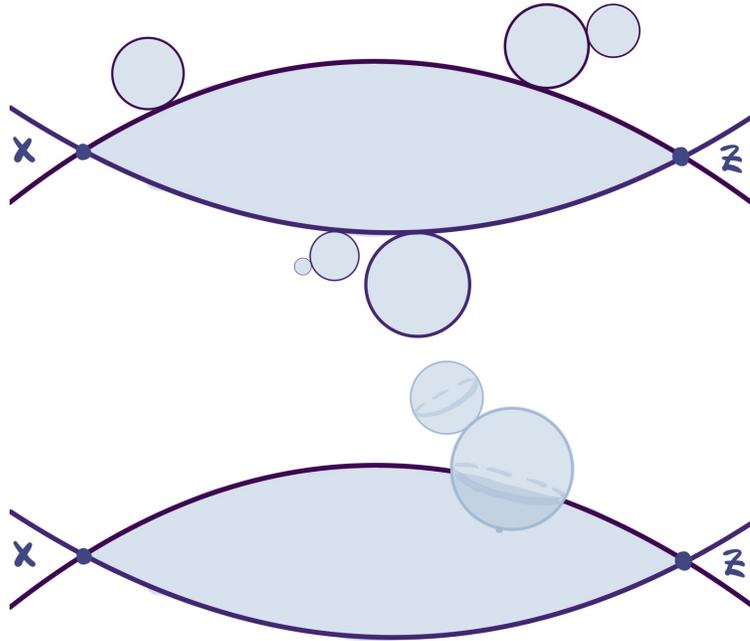
Although the monotone case most closely resembles Morse homology, there are new analytic difficulties which were not present before. The main concerns being compactness of $\tilde{\mathcal{M}}_J(x, y)$, nilpotency of ∂_J , and independence from choice of J . In [Flo88a], Floer proves the following results.

Theorem 4.18 (Floer homology). *Let M be a compact symplectic manifold with compact Lagrangian submanifolds L_0 and L_1 that intersect transversely. Suppose further that $\pi_2(M) = 0$ and $\pi_2(P, L_i) = 0$. Then Floer homology is well-defined:*

- Whenever $\mu(x) = \mu(y) + 1$, the moduli space $\tilde{\mathcal{M}}(x, y)$ is a finite set.
- The differential ∂_J is nilpotent, i.e. $\partial^2 = 0$.
- The homology theory $\mathrm{HF}_*(L_0, L_1)$ is independent of choice of regular almost-complex structure $J = \{J_t\}_t \in \mathcal{J}(L_0, L_1)$.

We will only give the ideas behind the first two points, leaving the details to [Flo88a, §3]. Compactness of the moduli spaces is due to the fact that the assumption $[\omega] \cdot \pi_2(M, L) = 0$ rules out any “bubbling” effects, wherein a sequence of J -holomorphic strips could converge to a J -holomorphic strip with J -holomorphic spheres or disks attached.





The three pictures above correspond to the three possible limiting behaviors of a converging sequence of J -holomorphic strips from x to z : to a trajectory “broken” at $y \in L_0 \cap L_1$, a trajectory with holomorphic disk bubbles (elements of $\pi_2(M, L_0)$ and $\pi_2(M, L_1)$), or a trajectory with holomorphic sphere bubbles (elements of $\pi_2(M)$). Our assumption means we only have to deal with the top picture, which corresponds to the broken trajectories of Morse homology. Proposition 2.2 of [Flo88a] says that if u_n is a sequence in $\mathcal{M}_J(x, y)$, then it contains a subsequence u_{n_i} so that there is a sequence $a_i \in \mathbb{R}$ such that the translated strips $u_{n_i}(s - a_i, t)$ converge to a limiting “broken” J -holomorphic strip. In other words, after we quotient by time translations, $\tilde{\mathcal{M}}_J(x, y) = \mathcal{M}_J(x, y)/\mathbb{R}$ is compact.

Eliminating the bubbling phenomenon is also key to showing that $\partial_J^2 = 0$. The proof is conceptually similar to the proof for the Morse differential, the idea being that

we can show the relevant coefficients count points in the boundary of a 1-manifold.

By definition,

$$\partial_J^2(x) = \sum_{\mu(y)=\mu(x)+1} \left(\sum_{\mu(z)=\mu(y)+1} m(x,y)m(y,z) \right) z$$

so it suffices to prove that the inner sum is 0 mod 2. A “gluing theorem” in [Flo88a, §4] says that for $\mu(x) - \mu(z) = 2$, the boundary of the moduli space can be decomposed into broken strips through some y with relative index 1,

$$\partial \tilde{\mathcal{M}}(x, z) = \coprod_{\substack{y \in L_0 \cap L_1 \\ \mu(y)=\mu(x)-1=\mu(z)+1}} \tilde{\mathcal{M}}_J(x, y) \times \tilde{\mathcal{M}}_J(y, z).$$

The cardinality of the right side is $m(x, y)m(y, z)$, whereas the cardinality of the left side is the cardinality of a smooth 1-manifold with boundary, and hence must be even. If the boundary of the moduli space contains limits with bubbling behavior, then the equality displayed above will not hold.

Zero case. When $\tau = 0$, then the Floer complex becomes more complicated. In general there may not be a \mathbb{Z} -graded homology, although we can always guarantee a $\mathbb{Z}/2$ -grading (if L_0, L_1 are oriented). In particular, we may no longer have an index $\mu(x)$ for each intersection point, but instead have the relative index $\mu(x, y; [u])$ which also depends on the homotopy class of a trajectory connecting x and y . Thus rather than having $\partial(x)$ associate a $\mathbb{Z}/2$ -coefficient to y (with $\mu(x, y; [u]) = 1$), we need to associate a coefficient which somehow lets us keep track of the possible homotopy

classes of chosen u . To make this meaningful, we need to work over the *Novikov field* for $\mathbb{Z}/2$,

$$\Lambda := \left\{ \sum_{j \geq 0} a_j T^{\lambda_j} \mid a_j \in \mathbb{Z}/2, \lambda_j \in \mathbb{R}, \lim_{j \rightarrow \infty} \lambda_j = \infty \right\}.$$

We then define $\text{CF}(L_0, L_1)$ to be the free Λ -module on intersection points $L_0 \cap L_1$.

The differential is defined to be

$$\partial_J(x) = \sum_{\substack{y \in L_0 \cap L_1 \\ \mu(x, y; [u]) = 1}} m(x, y; [u]) T^{\omega(u)} y$$

where now $m(x, y; [u])$ is the mod-2 count of $\tilde{\mathcal{M}}_J(x, y; [u])$ and $\omega(u)$ is the symplectic area $\int_{D^2} u^* \omega$ of u . In general, this sum could be infinite, but Gromov compactness (see [Sal97, §4] or [AD14, Theorem II.6.5.4]) ensures that the sum is well-defined in this case. As before, we need to know that $m(x, y; [u])$ is well-defined, $\partial_J^2 = 0$, and the resulting homology is independent of choice of J . The proofs for this proceed very similarly as in the monotone case, but of course with new subtleties introduced by working in each $\mathcal{M}_J(x, y; [u])$ individually and working over the Novikov field Λ . For more details, see [Aur13] or [Sal97].

To use Lagrangian intersection Floer homology to solve the Arnold Conjecture, one also has to show that Floer homology is invariant under Hamiltonian diffeomorphisms. This is known to be true for both $\mathbb{Z}/2$ coefficients and the Novikov field for $\mathbb{Z}/2$.

Theorem 4.19. *Let L_0 and L_1 be Lagrangian submanifolds of M which intersect transversely. Suppose $\phi: M \rightarrow M$ is a Hamiltonian diffeomorphism so that $L_0 \pitchfork$*

$\phi(L_1)$. Then

$$\mathrm{HF}(L_0, L_1) \cong \mathrm{HF}(L_0, \phi_1(L_1)),$$

as Λ -modules.

The idea is to show that for any two Hamiltonian diffeomorphisms ϕ_1, ψ_1 , there are chain maps $\mathrm{CF}(L_0, \phi_1(L_1)) \cong \mathrm{CF}(L_0, \psi_1(L_1))$ which induce isomorphisms on homology. See [Sal90, §3.4] for the explanation. The upshot is that Floer homology is well-defined even if we consider one Lagrangian manifold. That is, we can define $\mathrm{HF}(L, L) = \mathrm{HF}(L, \phi(L))$ where $\phi: M \rightarrow M$ is any Hamiltonian diffeomorphism so that $L \pitchfork \phi(L)$. Moreover, we can explicitly compute this group in terms of the ordinary homology groups of L .

Theorem 4.20. *There is an isomorphism of Λ -modules*

$$\mathrm{HF}(L, L) \cong H_*(L; \mathbb{Z}/2) \otimes_{\mathbb{Z}/2} \Lambda.$$

Consequently,

$$|L \cap \phi_1(L)| \geq \sum_{j=0}^n H_j(L; \mathbb{Z}/2).$$

Proof sketch. By the invariance of Floer homology under choice of Hamiltonian diffeomorphism, we may choose $\phi: M \rightarrow M$ small enough so that $L \pitchfork \phi(L)$ and $\phi(L)$ lies in a tubular neighborhood of L . By Weinstein's Lagrangian neighborhood theorem (Theorem 1.25), this neighborhood is symplectomorphic to a neighborhood of the zero

section of T^*L with its canonical symplectic structure ω . This symplectomorphism respects all the necessary structure (e.g. takes holomorphic strips to holomorphic strips), so it suffices to prove the isomorphism holds for L inside of (T^*L, ω) .

Now the idea is to choose a sufficiently nice Hamiltonian H on T^*L which is generated by a Morse function $f: L \rightarrow \mathbb{R}$ so that the Floer homology of H coincides with the Morse homology of L . Let f be a Morse function on L and define $H := -f \circ \pi$, where $\pi: T^*L \rightarrow L$ is the projection. The minus sign is present because if X_H is the Hamiltonian vector field so that $\omega(X_H, -) = dH$, then X_H has local representation

$$X_H = \frac{\partial f}{\partial x_1} \frac{\partial}{\partial \xi_1} + \cdots + \frac{\partial f}{\partial x_n} \frac{\partial}{\partial \xi_n},$$

where (x_1, \dots, x_n) are local coordinates for L and $(x_1, \dots, x_n, \xi_1, \dots, \xi_n)$ are the corresponding coordinates for T^*L . The diffeomorphisms $\phi(t)$ are locally given by

$$\phi_t(x_1, \dots, x_n, \xi_1, \dots, \xi_n) = (x_1, \dots, x_n, \xi_1 + t \frac{\partial f}{\partial x_1}(x), \dots, \xi_n + t \frac{\partial f}{\partial x_n}(x)),$$

so in particular the time 1 flow takes the zero section L to the graph Γ_{df} . The intersection points $L \cap \phi_1(L) = L \cap \Gamma_{df}$ are precisely the points for which $(x, 0) = (x, (df)_x)$, i.e. the critical points of f . This establishes a Λ -linear bijection $\text{CF}(L, \phi_1(L)) \rightarrow \text{Crit}(f) \otimes_{\mathbb{Z}/2} \Lambda$, which gives rise to an isomorphism on homology. To establish this fully, one needs to show that it is a chain map, i.e. there is bijection between $\mathcal{M}(x, y)$ (trajectories of $-\nabla f$) and $\mathcal{M}_J(x, y)$ (J -holomorphic strips). This is done by choosing

f and J in an appropriate way, see [Flo88a].

We have the following chain of inequalities:

$$\begin{aligned}
\#(L \cap \phi_1(L)) &\geq \dim(\mathrm{HF}(L, \phi_1(L))) \\
&= \dim(\mathrm{HF}(L, L)) \\
&= \dim(H_*(L; \mathbb{Z}/2) \otimes_{\mathbb{Z}/2} \Lambda) \\
&= \sum_{j=0}^n H_j(L; \mathbb{Z}/2).
\end{aligned}$$

□

In particular, in the monotone case when we get a \mathbb{Z} -graded homology theory with $\mathbb{Z}/2$ -coefficients, Floer homology is isomorphic to ordinary homology. As a corollary, we get Arnold conjecture in special cases: if $f: M \rightarrow M$ is a Hamiltonian diffeomorphism with non-degenerate fixed points, then we see that

$$\#\{\text{fixed points of } f\} \geq \sum_{j=0}^{2n} \dim H_j(M; \mathbb{Z}/2)$$

by applying the theorem to $\Delta(M)$ and Γ_f inside of $M \times M$ (recall from Section 1.3 that these are both Lagrangian submanifolds of $(M \times M, \tilde{\omega})$). However, we need to know that we have the right conditions to build Lagrangian intersection Floer homology, namely that $M \times M$ is symplectically aspherical (which holds if M is) and $\pi_2(M \times M, \Delta(M)) = \pi_2(M \times M, \Gamma_f) = 0$. When this condition is not satisfied, we

need to use slightly different techniques to build Floer homology. The next section discusses how to do this using the data of a Hamiltonian diffeomorphism.

4.3 Floer homology for Hamiltonian diffeomorphisms

Recall from Section 3.2 that the fixed points of a Hamiltonian diffeomorphism $\phi_1: M \rightarrow M$ are in bijection with periodic trajectories of a particular vector field. Let H_t denote the 1-periodic Hamiltonian on M that generates ϕ_t and let X_t denote the corresponding vector field. A fixed point of ϕ_1 will be an intersection point $x \in L \cap \phi_1(L)$, and $\phi_t(x)$ defines a closed orbit in M . This motivates us to study the (free) loop space of M and construct a function whose critical points are precisely the 1-periodic closed orbits of X_t . Of course, the loop space of M is not connected in general, so we will need to restrict to a particular connected component. A natural choice is the component containing all the *contractible* loops in M , since this component contains all the constant loops (which correspond to critical points in the time-independent case).

We start out with the action functional A from Section 4.2, but now defined on the space of contractible loops

$$\mathcal{L}M := \{\gamma: S^1 \rightarrow M \mid [\gamma] = 0 \in \pi_1(M)\}$$

Then, we perturb it slightly using our given Hamiltonian data,

$$A_H(\gamma) = \int_{D^2} w^* \omega - \int_0^1 H_t(\gamma(t)) dt$$

where $w: D^2 \rightarrow M$ is an extension of the loop γ to the disk. (Such a w is guaranteed to exist since γ is trivial in $\pi_1(M)$.)

Remark 4.21. As before, to ensure A_H is well-defined, we need to assume M is symplectically aspherical. Additionally, we replace the assumption for Lagrangian intersections that $[\omega] \cdot \pi_2(M, L_i) = 0$ with the assumption that for every smooth $f: S^2 \rightarrow M$, there is a symplectic trivialization of f^*TM , i.e. $c_1(f^*TM) = 0$. As before, this will eliminate any “bubbling” behavior. Both of these conditions are manifestly satisfied if $\pi_2(M) = 0$.

The critical points of A_H are precisely the 1-periodic closed orbits of X_t .

Proposition 4.22. *The differential of A_H is*

$$(dA_H)_\gamma(\xi) = \int_0^1 \omega(\xi(t), \gamma'(t) - X_t(\gamma(t))) dt$$

for $\xi \in T_\gamma \mathcal{LM}$, and dA_H vanishes if and only if $\gamma(t)$ is a periodic solution of the Hamiltonian system $\frac{d\gamma}{dt} = X_t(\gamma(t))$.

Proof. The proof is very similar to that of Proposition 4.9, and in fact by the proof

of Proposition 4.9 it suffices to differentiate $\int_0^1 H_t(\tilde{\gamma}_\varepsilon(t))dt$ at $\varepsilon = 0$. But this is

$$\begin{aligned} \int_0^1 \frac{\partial}{\partial \varepsilon} H_t(\tilde{\gamma}_\varepsilon(t))|_{\varepsilon=0} dt &= \int_0^1 (dH_t)_{\tilde{\gamma}_0(t)}(\xi(t)) dt \\ &= \int_0^1 \omega_{\gamma(t)}(X_t(\gamma(t)), \xi(t)) dt \\ &= \int_0^1 \omega(\xi(t), X_t(\gamma(t))) dt, \end{aligned}$$

and hence by linearity

$$(dA_H)_\gamma(\xi) = \int_0^1 \omega(\xi(t), \gamma'(t) - X_t(\gamma(t))) dt.$$

The differential is zero if and only if $\gamma'(t) = X_t(\gamma(t))$, as desired. \square

Proceeding as before, we choose a family of compatible almost-complex structures J_t which induce Riemannian metrics g_t . We define the inner product $\langle -, - \rangle$ as before and compute the L^2 -gradient of A_H to be

$$(\nabla_\gamma A)(t) = -J_t \frac{d\gamma}{dt} - \nabla_{\gamma(t)} H_t,$$

recalling from Remark 3.9 that $-J_t X_t(\gamma) = \nabla_\gamma H_t$. The “gradient flow lines” are thus paths $u: \mathbb{R} \rightarrow \mathcal{L}M$ which solve

$$\frac{du}{ds} = -J_t \frac{d\gamma}{dt} - \nabla_{\gamma(t)} H_t,$$

or rather $u(s, t): \mathbb{R} \times I \rightarrow \mathcal{L}M$ which solve

$$\frac{\partial u(s, t)}{\partial s} + J_t(u(s, t)) \frac{\partial u(s, t)}{\partial t} + \nabla H_t(u(s, t)) = 0. \quad (*_H)$$

The equation $(*_H)$ is often called the *Floer equation*, and it is just a Hamiltonian perturbation of the Cauchy-Riemannian equation $(*)$ from Section 4.2. In particular, if $H_t \equiv 0$, then we recover $(*)$. If $H_t = H$ does not depend on t , then the solutions u which also do not depend on t are exactly the trajectories of ∇H , since they satisfy

$$\frac{du}{ds} + \nabla H(u(s)) = 0.$$

Finally, if the solution u does not depend on s (meaning it is stationary), then it satisfies

$$\frac{du}{dt} = J_t(u) \nabla H_t(u) = X_t(u),$$

which supports our expectation that stationary trajectories are periodic solutions of the Hamiltonian system.

With this set up, building the Floer complex proceeds quite similarly as in the Lagrangian intersection setting. The outline is:

- Show that trajectories $u(s, t)$ with finite energy connect two critical points as $s \rightarrow \pm\infty$. Build the moduli space $\mathcal{M}_J(x, y)$ of these trajectories and show that it is (generically) a smooth manifold by describing it as the set of zeros of a

section (the *Floer operator*) of a bundle on a Banach manifold $\mathcal{P}(x, y)$.

- Adapt the Maslov index to this setting. In this case, we can assign each critical point (periodic trajectory) a Maslov index and hence get a \mathbb{Z} -graded complex. As before, $\dim \mathcal{M}(x, y) = \mu(x) - \mu(y)$.
- Define the differential by “counting” trajectories (modulo 2) connecting x to y (with $\mu(x) - \mu(y) = 1$). In particular, show that $\mathcal{M}(x, y)$ is compact and that the assumptions on M eliminate “bubbling,” so $\partial^2 = 0$.
- Show that the resulting homology (with $\mathbb{Z}/2$ -coefficients) is invariant of the almost-complex structure *and* the Hamiltonian H_t (equivalently, the vector field X_t).

We point the reader to [Sal97] or [AD14] for the details. Now, to prove the Arnold Conjecture, it just remains to relate the Floer homology of M to the ordinary homology of M . This is done by showing that Floer homology coincides with Morse homology for a nicely chosen Hamiltonian function, much like ??????. Specifically, we want to choose H so that all of the 1-periodic orbits are actually constant solutions. It turns out that this is true if H is time-independent and “sufficiently C^2 -small,” meaning that the norm of its Hessian is less than 2π .

Proposition 4.23. *Let M be a compact symplectic manifold and $H: M \rightarrow \mathbb{R}$ a Hamiltonian function. If H is sufficiently C^2 -small, then the only 1-periodic solutions its Hamiltonian system are the constant solutions.*

Proof. If H is sufficiently C^2 -small, then there is a finite cover of M by precompact Darboux charts so that every 1-periodic orbit is contained in a chart and $\|dX_H\| < 2\pi$ on each chart. It thus suffices to show the claim holds on a single chart.

So suppose H is a function on \mathbb{R}^{2n} with corresponding vector field X_H so that $\|dX_H\|_{L^2} < 2\pi$. Let $x(t)$ be a solution to the Hamiltonian system. We want to show $x'(t) = 0$. Since $x''(t) = (dH_X)_x \cdot x'$, the bound $\|dX_H\|_{L^2} < 2\pi$ gives

$$\|x''\|_{L^2} < 2\pi \|x'\|_{L^2}$$

if $x'(t) \neq 0$. We will use derive a contradiction by studying the Fourier expansion of x , $\frac{dx}{dt}$, and $\frac{d^2x}{dt^2}$:

$$\begin{aligned} x(t) &= \sum_n c_n(x) e^{2\pi i n t}, \\ x'(t) &= \sum_n 2\pi i n c_n(x) e^{2\pi i n t}, \\ x''(t) &= \sum_n -4\pi^2 n^2 c_n(x) e^{2\pi i n t}. \end{aligned}$$

An application of Parseval's identity gives us

$$\|x''\|_{L^2}^2 = \sum_n |-4\pi^2 n^2 c_n(x')|^2 \geq 4\pi^2 \sum_{n \neq 0} |c_n(x')|^2 = 4\pi^2 \|x'\|_{L^2}^2,$$

since $c_0(x') = 0$. Thus $2\pi \|x'\|_{L^2} \leq \|x''\|_{L^2}$, a contradiction. Hence $x'(t) \equiv 0$, so x is constant. \square

Consequently the periodic orbits of such an H are in 1-to-1 correspondence with its critical points. To connect the Morse homology of H with the Floer homology of A_H (assuming both are well-defined), the idea is to show that the complexes are the same after a shift,

$$\mathrm{CF}_*(H, J; \mathbb{Z}/2) = \mathrm{CM}_{*+n}(H, -JX_H; \mathbb{Z}/2).$$

It is straightforward to show that $\mathrm{Crit}(H) = \mathrm{Crit}(A_H)$, and with a bit more work one can show that the (Maslov) index $\mu(x)$ is the shifted (Morse) index $\mathrm{ind}(x) + n$. The most difficult part is defining the differentials of the two complexes and showing that they coincide. The idea is to pick a compatible almost-complex structure J so that $\nabla H = -JX_H$ is Morse-Smale (and the Floer complex can be defined) and find a relation between the gradient flow lines and the solutions to the Floer equation, that is, the solutions of

$$\frac{du}{ds} + JX_H(u) = 0 \quad \text{and} \quad \frac{\partial u}{\partial s} + J(u)\frac{\partial u}{\partial t} + \nabla H(u) = 0.$$

By taking H/k for k sufficiently large, the solutions of the Floer equation which connect critical points of relative index 2 are independent of t , and are exactly the trajectories of $-JX_H$. In other words, the data of the Floer complex is the same as the data of the Morse complex. Thus we can apply the Morse inequalities to finish the proof of the Arnold Conjecture (for $\mathbb{Z}/2$ -coefficients).

4.4 Epilogue: Floer homotopy theory

Recently, researchers have been interested in studying Floer theory from the perspective of homotopy theory. The beginning of this movement is widely attributed to the work of Cohen-Jones-Segal in the 1990s [CJS95]. One of the main goals of Floer homotopy theory is to associate a (stable) homotopy type to the geometric data of Floer homology. Their proposal was to use a “framed flow category” (built out of similar data as the Floer homology complex) and use it to build a spectrum with the appropriate ordinary homology. There are several variants of the idea of a framed flow category in the literature, and sometimes the same definition will appear under a different name, or different definitions will appear under the same name. Rather than giving a precise definition, we will just highlight some of the common themes, pointing the reader to the references we have cited for specific details (see, in particular, the *compact smooth categories/Morse-Smale categories* in [Coh19, Definition 6] or the *equivariant flow categories* in [AB21, §2.1] or the *flow categories* in [LS14, §3.2]).

“Definition”. A *framed flow category* is a topologically enriched category \mathcal{C} whose homspaces are compact, smooth, framed manifolds with corners. In most cases, the objects come with grading $gr: \text{Ob } \mathcal{C} \rightarrow \mathbb{Z}$, and the homspaces are subject to certain conditions based on the grading, such as:

- $\text{Hom}(x, x) = \{\text{id}_x\}$ for all $x \in \text{Ob } \mathcal{C}$.

- For objects $x \neq y$, if $gr(x) \leq gr(y)$ then $\text{Hom}(x, y)$ is empty, and otherwise $\text{Hom}(x, y)$ is $(gr(x) - gr(y) - 1)$ -dimensional.
- Composition $\circ: \text{Hom}(x, y) \times \text{Hom}(y, z) \rightarrow \text{Hom}(x, z)$ is an embedding into the boundary of $\text{Hom}(y, z)$ in a special way. Furthermore, each point in the boundary of $\text{Hom}(x, y)$ is in the image of the composition map.

These conditions can be made more precise by describing the homspaces as $\langle k \rangle$ -*manifolds* (where k is the dimension of the homspace)⁸. The “framed” part of the definition comes from so-called *neat embeddings* of the homspaces into Euclidean spaces with corners. The homspaces play the role of moduli spaces of trajectories, and the grading on the objects is something like the Morse, Maslov, or Conley-Zehnder index. The idea is that this information should be enough to recover the interesting topological information.

For instance, the prototypical example of a framed flow category comes from Morse theory, as was also developed by Cohen, Jones, and Segal in their unpublished preprint “Morse Theory and Classifying Spaces.” Let $f: M \rightarrow \mathbb{R}$ be a Morse-Smale function on a smooth, closed, finite-dimensional Riemannian manifold M , the *flow category* of f is a prototypical example of a framed flow category. The *flow category* of f is a topologically enriched category \mathcal{C}_f whose objects are the discrete space of critical points and whose homspaces are the *moduli space of broken gradient flows* between them. The grading is given by the Morse index and composition is concatenation

⁸Roughly, $\text{Hom}(x, y)$ being a $\langle k \rangle$ -manifold means that the boundary can be split up into k pieces which intersect at the corners of the manifold. For a precise definition, see [LS14, §3.1].

of the broken flows. The Morse-Smale condition on f ensures that the compactified moduli spaces have the extra structure necessary to make \mathcal{C}_f into a framed flow category (although this is not easy to prove, and in fact was part of the reason that the original preprint was never published — see [Coh19, Remark on p.16] for more details). Using this structure, Cohen-Jones-Segal show in the preprint that for a Morse-Smale function $f: M \rightarrow \mathbb{R}$, the classifying space $B\mathcal{C}_f$ is homeomorphic to the underlying manifold M .

Some of the ideas of Floer homotopy theory have found applications in symplectic geometry and low dimensional topology, particularly in the work of Lipschitz-Sarkar on Khovanov homology [LS14]; more applications are summarized in [Coh19]. Most recently, Abouzaid-Blumberg [AB21] have proved the Arnold Conjecture for coefficients in a field of characteristic p (among many other things, including another proof of the Cohen-Jones-Segal theorem for $B\mathcal{C}_f$ in Appendix D). Rather than trying to construct a Floer complex with characteristic p coefficients, they construct a version of Floer homology with coefficients in Morava K -theory. Lying in the intersection of many areas of topology and geometry, Floer homotopy theory is in its early stages of development, but many interesting things are surely still to come.

References

- [AB21] Mohammed Abouzaid and Andrew Blumberg. Arnold conjecture and Morava K -theory, 2021. arXiv:2103.01507 [math.SG].
- [AD14] Michèle Audin and Mihai Damian. *Morse Theory and Floer Homology*. Universitext. Springer, London, 2014.
- [Arn78] V. Arnold. *Mathematical Methods of Classical Mechanics*. Number 60 in Graduate Texts in Mathematics. Springer-Verlag, 1978.
- [Aur13] Denis Auroux. A beginner’s introduction to Fukaya categories, 2013. arXiv:1301.7056 [math.SG].
- [CJS95] Ralph Cohen, John Jones, and Graeme Segal. Floer’s infinite-dimensional Morse theory and homotopy theory. In H. Hofer, C. Taubes, A. Weinstein, and E. Zehnder, editors, *The Floer Memorial Volume*, volume 133, pages 297–325. Birkhäuser, 1995.

- [Coh19] Ralph Cohen. Floer homotopy theory, revisited, 2019. arXiv:1901.08694 [math.AT].
- [CZ83] C. Conley and E. Zehnder. The Birkhoff-Lewis fixed point theorem and a conjecture of V.I. Arnold. *Inventiones Mathematicae*, 73:33–49, 1983.
- [CZ84] C. Conley and E. Zehnder. Morse-type index theory for flows and periodic solutions of Hamiltonian equations. *Communications in Pure and Applied Mathematics*, 37:207–253, 1984.
- [dS06] Ana Cannas da Silva. *Lectures on Symplectic Geometry*. Number 1764 in Lecture Notes in Mathematics. Springer-Verlag, 2006.
- [Eli79] Y. Eliashberg. Estimates on the number of fixed points of area preserving transformations, 1979. Syktyvar University preprint.
- [Flo88a] Andreas Floer. Morse theory for Langrangian intersections. *Journal of Differential Geometry*, 28:513–547, 1988.
- [Flo88b] Andreas Floer. A relative Morse index for the symplectic action. *Communications on Pure and Applied Mathematics*, 41:393–407, 1988.
- [Flo88c] Andreas Floer. The unregularized gradient flow of the symplectic action. *Communications on Pure and Applied Mathematics*, 41:775–813, 1988.
- [Flo89a] Andreas Floer. Symplectic fixed points and holomorphic spheres. *Communications in Mathematical Physics*, 120:575–611, 1989.

- [Flo89b] Andreas Floer. Witten's complex and infinite dimensional Morse theory. *Journal of Differential Geometry*, 30:207–221, 1989.
- [FO99] F. Fukaya and K. Ono. Arnold conjecture and Gromov-Witten invariants for general symplectic manifolds. *Topology*, 38:993–1048, 1999.
- [Got82] M. Gotay. On coisotropic embeddings of presymplectic manifolds. *Proceedings of the American Mathematical Society*, 84:111–114, 1982.
- [HS95] H. Hofer and D.A. Salamon. Floer homology and Novikov rings. In H. Hofer, C. Taubes, A. Weinstein, and E. Zehnder, editors, *The Floer Memorial Volume*, pages 483–524. Birkhäuser, 1995.
- [LS14] Robert Lipshitz and Sucharit Sarkar. A Khovanov stable homotopy type. *Journal of the American Mathematical Society*, 27:983–1042, 2014.
- [LT98] Gang Liu and Gang Tian. Floer homology and Arnold conjecture. *Journal of Differential Geometry*, 49:1–74, 1998.
- [McD98] Dusa McDuff. Lectures on symplectic topology. *IAS/Park City Math. Series*, 7, 1998.
- [Mos65] J.K. Moser. On the volume elements of manifolds. *Transactions of the American Mathematical Society*, 120:280–296, 1965.

- [MS95] Dusa McDuff and Dietmar Salamon. *Introduction to Symplectic Topology*. Oxford Mathematical Monographs. Oxford University Press, New York, 1995.
- [NN57] A. Newlander and L. Nirenberg. Complex analytic coordinates in almost complex manifolds. *Annals of Mathematics*, 65:391–404, 1957.
- [Ono95] K. Ono. The Arnold conjecture for weakly monotone symplectic manifolds. *Inventiones Mathematicae*, 119:519–537, 1995.
- [RS01] Joel .W. Robbin and Dietmar A. Salamon. Asymptotic behavior of holomorphic strips. *Annales de L’Institut Henri Poincaré — Analyse Non-linéaire*, 18:573–612, 2001.
- [Sal90] Dietmar A. Salamon. Morse theory, the Conley index, and Floer homology. *Bulletin of the London Mathematical Society*, 22:113–140, 1990.
- [Sal97] Dietmar Salamon. Lectures on Floer homology, 1997. Notes for lectures held at the IAS/Park City Graduate Summer School.
- [Tau94] C. Taubes. The Seiberg-Witten invariants and symplectic forms. *Mathematical Research Letters*, 1:809–822, 1994.
- [Thu76] W. Thurston. Some examples of symplectic manifolds. *Proceedings of the American Mathematical Society*, 55:467–468, 1976.

- [Wei77] A. Weinstein. *Lectures on Symplectic Manifolds*. Number 29 in Regional Conference Series in Mathematics. American Mathematical Society, 1977.
- [Wei81] A. Weinstein. Neighborhood classification of isotropic embeddings. *Journal of Differential Geometry*, 16:125–128, 1981.