Introduction

First impressions. A Fukaya category is an algebraic structure associated to a symplectic manifold. It encodes information about Lagrangian submanifolds, the way in which any two of them intersect, and also about pseudo-holomorphic discs (or polygons). The amount of data packed into a single object makes Fukaya categories fascinating, but also somewhat intimidating, especially from a symplectic topology viewpoint. Before going on, we would like to elaborate a little further on this observation. This will be an informal discussion, addressed to non-specialist readers, and intended to convey just a general idea of what the situation looks like; it contains many statements that are not rigorous, and it is also somewhat more wide-ranging than the rest of book.

Fix a symplectic manifold M^{2n}. Roughly speaking, objects of the Fukaya category are Lagrangian submanifolds $L^n \subset M$. The space of such submanifolds is infinite-dimensional (locally modelled on the space of closed one-forms on L). However, Lagrangian submanifolds which are Hamiltonian isotopic give rise to isomorphic objects in the Fukaya category, which cuts down the local degrees of freedom to the first Betti number $b_1(L)$. Indeed, the classification up to Hamiltonian isotopy is completely understood in some simple cases: curves on surfaces, where the issue is essentially one of topology; less obviously, Lagrangian spheres in $S^2 \times S^2$ [75], and a few related examples. There are other (general) effects which work in one's favour: not all Lagrangian submanifolds actually appear in the Fukaya category, due to obstructions in the sense of [60]; others may occur but be trivial (the zero object); and isomorphism in the Fukaya category may be weaker than Lagrangian isotopy. For an example of the first point, one can simply look at circles in S^2, where the necessary cancellation of obstructions only happens when the two hemispheres in the complement have equal areas; for more interesting cases see [31], [32]. The second phenomenon occurs for the real locus $\mathbb{R}\mathrm{P}^n \subset \mathbb{C}\mathrm{P}^n$ for $n > 1$, assuming that that is Spin ($n = 4k-1$), and that the Fukaya category is taken with $\mathbb{Q}$-coefficients. For the third point, there are some candidate examples in the form of Lagrangian tori in $\mathbb{C}\mathrm{P}^2$, but as far as this author knows, none of them have been rigorously verified. With all this at hand, it becomes feasible in principle to compute the entire Fukaya category for manifolds whose symplectic topology is comparatively simple; the papers quoted above (and others, for instance [132], [102], [101], [59]) can be viewed as part of a wider ongoing effort in that direction. Still, for any reasonably complicated symplectic manifold (for instance, a complex hypersurface in $\mathbb{C}\mathrm{P}^3$ of degree ≥ 4), a direct classification of the objects in the Fukaya category appears to be out of reach.

Maybe we have just been setting the goalposts too high. After all, even if one considers a purely algebraic context, such as modules over a finite-dimensional algebra, the cases where a complete classification can be achieved form a tiny minority [41]. There is, however, one significant difference: for two explicitly given modules, there is an obvious algorithm which decides whether they are isomorphic. More generally, given a finite collection of modules, it is straightforward to determine the morphisms between them, and the composition maps of morphisms. In contrast, in a Fukaya category, the morphisms are Floer cochain groups, and the composition maps operations on such groups, both defined in terms of suitable moduli spaces of pseudo-holomorphic maps. There are some situations where these spaces can be dealt with geometrically: when M is a cotangent bundle and the Lagrangian submanifolds $L = \mathrm{graph}(df)$ are sections, an "adiabatic limit" method reduces the question to one in finite-dimensional Morse theory [58] (variations of this method have been successful in several related cases); when M is an algebraic variety defined over the real numbers, and one considers only the real locus $L = M_{\mathbb{R}}$, algebro-geometric tools can be applied; finally, in the case where M is a surface, the study of pseudo-holomorphic curves becomes combinatorial to a large extent. However, beyond these situations, Floer theory computations traditionally rely on case-by-case methods, which offer no a priori guarantee of success. This book is an attempt to address that problem. Together with other developments, it will hopefully help to make Fukaya categories into a more manageable tool.

The strategy. In view of the difficulties described above, it is clear that if one wants to make any headway, some drastic simplifications are needed. To begin with, we consider only exact symplectic manifolds M, namely those where the symplectic form ø is the exterior derivative of some one-form θ. This of course implies that M cannot be closed, but we impose standard convexity conditions, which make this deficiency invisible to finite area pseudo-holomorphic curves. Affine algebraic varieties are a good source of examples. The simplest class of Lagrangian submanifolds $L \subset M$ are the exact ones, which means that $\theta | L$ is the derivative of some function; we will only deal with these, and exclude all others a priori. This restriction has many useful technical consequences: for instance, exact Lagrangian submanifolds cannot be obstructed, and each of them is a nonzero object of the Fukaya category, whose endomorphism ring is the classical cohomology ring (no quantum corrections). Beyond that, the more fundamental advantage of exactness is that it eliminates the dependence of the Fukaya category on the symplectic area (or Novikov ring) parameter. Mirror symmetry suggests that this dependence is often given by nontrivial transcendental functions, and this is confirmed by sample computations, such as [108], [57]. In our case, the parameter is set equal to zero (or infinity, depending on one's interpretation), which makes the situation less interesting from an enumerative point of view, but more amenable to computation. We should mention that ultimately, one can try using the Fukaya categories of affine varieties as a stepping-stone towards those of their projective completions [129].

To bring things closer to classical homological algebra, we will also impose a Calabi–Yau type condition ($2c_1(M) = 0$, to be precise), together with a corresponding assumption on the Lagrangian submanifolds, which allows us to equip all Floer groups with $\mathbb{Z}$-gradings. The resulting setup can be summarized as follows. Objects of the Fukaya category $\mathcal{F}(M)$, which we call exact Lagrangian branes, are triples $L^\# = (L, \alpha^\#, P^\#)$, where L is a closed exact Lagrangian submanifold of M, $\alpha^\#$ is a real-valued function on L called the grading, and $P^\#$ is a Pin structure on L. The morphisms spaces are Floer cochain groups $CF^*(L_0^\#, L_1^\#)$, which are finite-dimensional graded vector spaces (over some fixed, but arbitrary, coefficient field $\mathbb{K}$; the Pin structures are really only necessary if $\mathrm{char}(\mathbb{K}) \neq 2$). These carry multilinear composition maps μ^d, $d \geq 1$, which form an A_∞-structure. The definition requires auxiliary choices, but the outcome is independent of those choices up to quasi-isomorphism.

Our second main idea, which comes from Kontsevich's work, is to take formal enlargements of the Fukaya category. These are purely algebraic constructions, which can be applied to any A_∞-category $\mathcal{A}$. The first step is to introduce analogues of chain complexes, so-called twisted complexes, which form another A_∞-category $Tw\,\mathcal{A}$ containing $\mathcal{A}$. This contains mapping cones of morphisms, hence is what we call a triangulated A_∞-category; the resulting cohomology level category $H^0(Tw\,\mathcal{A})$, also called the derived category $D(\mathcal{A})$, is a triangulated category in the classical sense. As a second useful step, one can take the split-closure (or Karoubi completion) $D^\pi(\mathcal{A})$ of the derived category, which introduces formal direct summands of all idempotent endomorphisms; there is also a corresponding construction on the A_∞ level, denoted by $\Pi(Tw\,\mathcal{A})$. An apparently paradoxical statement, but one whose truth will be evident to many readers, is that each enlargement makes the category easier to describe. As an illustration, consider the notion of split-generators: we say that a subset of objects split-generates $\Pi(Tw\,\mathcal{A})$ if one can construct all objects out of them by repeatedly applying mapping cones and splitting off direct summands. Whenever $\mathcal{B} \subset \Pi(Tw\,\mathcal{A})$ is the full A_∞-subcategory formed by such a set of split-generators, $\Pi(Tw\,\mathcal{B})$ is quasi-equivalent to $\Pi(Tw\,\mathcal{A})$. Hence, if one is only interested in $\Pi(Tw\,\mathcal{A})$, it is actually sufficient to determine one such $\mathcal{B}$. Let us keep in mind that split-generation is a very weak property; it does not, for instance, imply that the Grothendieck group $K_0(D^\pi(\mathcal{A}))$ is spanned by the classes of split-generators. Hence, it is not unreasonable to hope that a finite set of Lagrangian submanifolds can be found which split-generate $\Pi(Tw\,\mathcal{F}(M))$, and we will see that this is indeed the case in favourable circumstances.

The use of the word "computation" in connection with Fukaya categories means that the underlying geometric structures (symplectic manifolds, Lagrangian submanifolds) have to be somehow encoded in a combinatorial way. For that, we will make systematic use of Picard–Lefschetz theory. In general, the importance of this approach in symplectic geometry largely comes from Donaldson's work [38]; here, we only use the most elementary aspects of the theory, namely the symplectic nature of monodromy maps, as well as the Lagrangian nature of Lefschetz thimbles and vanishing cycles. To take advantage of these properties, we will introduce (following an idea of Kontsevich) Fukaya categories of Lefschetz fibrations, which in a sense stand between the ordinary Fukaya category of the total space and that of a smooth fibre. These will then be used to build a machine, first proposed conjecturally in [126], for doing computations by dimensional induction.

Computing Fukaya categories. To simplify formulations, we will limit ourselves to the case of algebraic varieties. Namely, let $X \subset \mathbb{C}^N$ be an $(n+1)$-dimensional affine algebraic variety which is smooth, as well as smooth at infinity. The latter condition means that its closure $\bar{X} \subset \mathbb{C}\mathrm{P}^N$ is again smooth, and intersects the hyperplane at infinity $\mathbb{C}\mathrm{P}^{N-1} = \mathbb{C}\mathrm{P}^N \setminus \mathbb{C}^N$ transversally. We equip X with the restriction of the Fubini–Study Kähler form, which turns it into an exact symplectic manifold. Additionally, we will assume that the canonical bundle $\mathcal{K}_{\bar{X}}$ is isomorphic to $\mathcal{O}_{\bar{X}}(-d)$ for some $d \in \mathbb{Z}$, which makes X Calabi–Yau. All these properties are also inherited by a generic affine hyperplane section $X^1 = X \cap \mathbb{C}^{N-1}$. Take the Fukaya categories $\mathcal{F}(X)$ and $\mathcal{F}(X^1)$, with coefficients in some field $\mathbb{K}$. We begin by discussing X^1, since our results there are somewhat better:

Theorem A. *Suppose that $d \neq 2$, $n \geq 1$, and* $\mathrm{char}(\mathbb{K}) \neq 2$. *Then $\Pi(Tw\,\mathcal{F}(X^1))$ is computable from the combinational data obtained by applying Picard–Lefschetz theory to X.*

The statement requires some explanation. First of all, by saying that $\Pi(Tw\,\mathcal{A})$ is computable for some A_∞-category $\mathcal{A}$, we mean the following. One can explicitly write down another A_∞-category $\mathcal{B}$, which is finite in a strong sense (finitely many objects, the morphisms form finite-dimensional vector spaces, and composition maps $\mu^d_{\mathcal{B}}$ of sufficiently high order $d \gg 0$ all vanish), such that $\Pi(Tw\,\mathcal{B})$ and $\Pi(Tw\,\mathcal{A})$ are quasi-equivalent A_∞-categories. In particular, this means that $D^\pi(\mathcal{B}) \cong D^\pi(\mathcal{A})$ as (classical) triangulated categories. Note that for a reasonable choice of coefficient field $\mathbb{K}$ (a finite field, or $\mathbb{Q}$, for instance), $\mathcal{B}$ really contains a finite amount of information. One can take that information and enter it into a computer program which, by explicitly constructing twisted complexes and their idempotent endomorphisms, will generate the full (infinite) list of isomorphism classes of objects in $D^\pi(\mathcal{B})$, together with morphisms, composition maps, and exact triangles. This, of course, does not by itself solve any of the big structural questions about $D^\pi(\mathcal{B})$. As for the meaning of “combinational data produced by Picard–Lefschetz theory”, we leave the precise explanation to Section 19, but the rough idea is to look at linear maps from X to $\mathbb{C}$ and $\mathbb{C}^2$, and encode the branch data of those maps in terms of braid monodromy. The same process can be applied to X^1, and repeated until the dimension is reduced to zero. Given explicit equations defining X, these data can be extracted algorithmically using elimination theory (even though in all but the simplest examples, the complexity tends to be prohibitive).

We also need to look briefly at the limitations of the theorem. $n \geq 1$ excludes the lowest nontrivial dimension, where X^1 is a Riemann surface. This is really for convenience only: the Riemann surface case can be treated with similar (in fact, somewhat simpler) methods, but requires some particular adjustments, which would complicate the discussion. The $\mathrm{char}(\mathbb{K}) \neq 2$ condition is a technical artifact of our proof, which uses double branched covers and $\mathbb{Z}/2$-actions; quite probably, other methods exist which would allow one to lift this restriction. The other requirement $d \neq 2$ is somewhat more essential, and we postpone its discussion to a later point (see Remark 19.8).

As mentioned before, the results for $\mathcal{F}(X)$ itself are somewhat weaker. This is not unexpected: for instance, there is no known method for deciding whether X contains any closed exact Lagrangian submanifold, hence whether $\mathcal{F}(X)$ (and its derived categories) are nontrivial. This is in contrast to the case of X^1, where we could use vanishing cycles as a natural source of objects. The best we can do is this:

Theorem B. *Assume that* $\mathrm{char}(\mathbb{K}) \neq 2$. *Then, from the combinatorial data obtained by applying Picard–Lefschetz theory to X, one can construct an A_∞-category which contains a full subcategory quasi-equivalent to $\mathcal{F}(X)$.*

The A_∞-category in question is of the form $\mathit{Tw}\,\mathcal{B}$, where $\mathcal{B}$ is finite in the same sense as before. In fact, by construction $\mathcal{B}$ will be directed, and this has some noteworthy consequences. First of all, $\mathit{Tw}\,\mathcal{B}$ is already split-closed, so it will contain not only $\mathcal{F}(X)$, but the whole of $\Pi(\mathit{Tw}\,\mathcal{F}(X))$. Secondly, directed A_∞-categories are quasi-isomorphic to dg categories with finite-dimensional morphism spaces. Since that property is inherited by the category of twisted complexes, we have:

Corollary C. *In the situation of Theorem B, $\mathcal{F}(X)$ is quasi-isomorphic to a dg category with finite-dimensional hom spaces.*

The practical usefulness of Theorem B depends on how well one understands the geometric meaning of the map $\mathcal{F}(X) \rightarrow \mathit{Tw}\,\mathcal{B}$. For the simplest class of Lagrangian submanifolds in X, namely Lagrangian spheres which can be represented as matching cycles for a Lefschetz pencil of hyperplane sections, these issues can be addressed easily, so that one indeed gets an algorithm for computing Floer cohomology groups (strictly speaking, the case $n = 1$ ought to be excluded here, for the same reason as in Theorem A). For more general Lagrangian submanifolds, the statement remains more of a theoretical nature, at least for now.

The proofs of the results stated above will be given in Section 19; they involve most of the material introduced throughout the book. As a natural by-product, we establish a general relation between Hurwitz moves of vanishing cycles, and mutations of directed Fukaya categories. In fact, we will give two proofs of this (Theorem 17.20 and Theorem 18.24; the proof originally envisaged in [127] is different from either one, and remains unpublished at present). Another interesting consequence is the existence of a spectral sequence for the Floer cohomology of Lagrangian submanifolds lying in the total space of a Lefschetz fibration (Corollary 18.27). Of course, like any spectral sequence, this merely decomposes the problem into a visible part (the starting term) and hidden higher order information (the differentials). In this particular case, the differentials express the way in which the Lagrangian submanifolds are built up from Lefschetz thimbles, where the word "built" is interpreted algebraically, as a Postnikov decomposition in the Fukaya category.

With one exception at the end, this book does not deal with examples or applications. Luckily, the existing literature makes up for that shortcoming (for some pointers, see the beginning of Chapter III). Beyond that, any reader seriously interested in Fukaya categories will eventually have to tackle the groundbreaking [60], which discusses many geometric and algebraic issues far outstripping the framework set up here.