Homological Mirror Symmetry

Given a smooth, compact complex variety X, we can study it through its bounded derived category of coherent sheaves, $D^{\mathrm{b}}\mathrm{Coh}(X)$, and homological mirror symmetry lets us turn the study of that category into symplectic geometry. We do this by looking for a *mirror*, a symplectic manifold $X^\vee$ and an equivalence $D^{\mathrm{b}}\mathrm{Coh}(X) \cong \mathrm{Fuk}(X^\vee)$. There are very many technicalities involved which we will ignore today, this lecture is meant to be an introduction that gives you a rough picture of what's going on.

Symplectic manifolds

Definition. *A* symplectic manifold *is an even dimensional smooth manifold X equipped with a closed 2-form ω such that the map $T_pX \to T_pX$ given by $v \mapsto \omega(v, \_)$ is an isomorphism.*

Here are some examples:

- We can take $\mathbb{R}^2$ with symplectic form $dx \wedge dy$.
- More generally, $\mathbb{R}^{2n}$ with form $\sum_{i=1}^n dx_i \wedge dy_i$.
- Even more generally, given any n dimensional manifold N, the cotangent bundle T^*N is a symplectic manifold via a symplectic form given in local coordinates by $\sum_{i=1}^n dp_i \wedge dq_i$, where $p_1, \dots, p_n$ are local coordinates in N over a coordinate patch U where the cotangent bundle is trivial, and $q_1, \dots, q_n$ are the coordinates on $T * N \mid_U$ giving the trivialization.

Symplectic manifolds arise naturally in classical mechanics. Say you have a system where the space of possible positions of the particles is a manifold N, then T^*N parametrizes both position and momentum of the particles in the system, and is called the phase space. The evolution of the system can be described on T^*N: there is a function $H : T^*N \to \mathbb{R}$, called the energy or Hamiltonian of the system that determines the evolution in the following way: under the isomorphism given by ω, dH will correspond to some vector field X_H called a Hamiltonian vector field, and integrating this field produces a flow $\phi_H(t) : Y \to Y$, which turns out to be a symplectomorphism, i.e., $\phi_H(T)^*\omega = \omega$ at any time t.

Note. This way of setting up classical mechanics only uses the symplectic structure on T^*N, not the fact that it is the cotangent bundle, and so makes sense on other symplectic manifolds. There are indeed natural examples of phase spaces which are not of the form T^*N.

Symplectic geometry is soft

Generally speaking you should think of symplectic manifolds as *soft* o *floppy*, for example:

- A general symplectic manifold Y has a infinite dimensional space of symplectomorphisms (we already saw that a Hamiltonian produces a 1-parameter family of symplectomorphisms). Contrast this with the much more rigid case of algebraic varieties which can easily have only finitely many automorphisms.

- Symplectic manifolds are all locally symplectomorphic to the example of $R^{\nvDash\ltimes}$ with 2-form $\sum_{i=1}^{n} dx_i \wedge dy_i$. So unlike, say, the case of Riemannian geometry where there are interesting local invariants such as curvature, symplectic geometry is global.

Lagrangian submanifolds

Definition. *A* Lagrangian *in a symplectic manifold Y of dimension $2n$ is an n-dimensional submanifold L such that $\omega|_L = 0$. (This means $\omega(v,w) = 0$ if both v and w are tangent to L.)*

Some examples:

- In $\mathbb{R}^2$ we can take the subset given by $x = 0$, or more generally, the graph $\Gamma_f = \{(x, f(x)) : x \in \mathbb{R}\}$ of any smooth function $f : \mathbb{R} \to \mathbb{R}$.

- In T^*N, the zero section is a Lagrangian.

- Given two symplectic manifolds (Y, ω), (Y', ω'), the manifold $Y \times Y'$ can be given the structure of a symplectic manifold by using the 2-form $\omega \times (-\omega')$. Given any diffeomorphism[1] $f : Y \to Y'$, we have that f is a symplectomorphism if and only if Γ_f is a Lagrangian in $Y \times Y'$.

[1] the existence of which forces Y and Y' to have the same dimension, of course

Two Lagrangians inside the same symplectic manifold Y have complementary dimensions and thus, if they are positioned generically will have a discrete set of points of intersection. These points of intersection often have geometric or dynamic meaning, for example, if H is a Hamiltonian on Y, we might be interested in finding all of the orbits of the Hamiltonian flow $\phi_H(t)$ that have period, say, 1. These orbits are given as the intersection $\Gamma_{\phi_H(1)} \cap \Delta$ inside $Y \times Y$ (again, with symplectic structure given by changing the sign of one of the ω's) of the graph of the time 1 flow with the diagonal. In the case of a time independent Hamiltonian, this intersection is *not* transverse and does not consist of a discrete set of points, but rather consists of whole period 1 orbits. This example works better for a time dependent Hamiltonian, but it was only meant to illustrate that intersection points of Lagrangians, here $\Gamma_{\phi_H(t)}$ and Δ, can have intrinsic interest.

Because of this we want some neat way to package Lagrangians and their intersection points, and this is what the Fukaya category is meant to achieve. There are many technicalities involved in defining the Fukaya category and here we will only give an impressionistic account through a series of increasingly detailed definitions.

Succesive approximations to the Fukaya category

Fukaya category, v0.1. The first, roughest approximation to the Fukaya category of a symplectic manifold Y is that the objects are Lagrangians in Y and morphisms $L_1 \to L_2$ are given by paths from some point on L_1 to some point on L_2. Notice that the points of intersection $L_1 \cap L_2$ correspond to the paths of length 0. Notice also that you can't always compose paths, only when their endpoints match appropriately. After this talk you should forget this construction as it is completely wrong.

Fukaya category, v0.2. We can think of the length of a path as a Morse function μ on the space of paths from a point on L_1 to a point on L_2 (or else replace it by a Morse function), and define $\hom(L_1, L_2)$ to the homology $H^\bullet(C_\mu(L_1, L_2))$ of the Morse complex instead. The critical points of this Morse function will be, as mentioned before, the points of intersection of L_1 and L_2. Before going further we should take a short detour:

A 15 second introduction to Morse theory

(If it's supposed to be possible to give an hour and a half introduction to the Fukaya category, it should be possible to give a 15 second introduction to Morse Theory!)

Let M be a smooth manifold and $\mu : M \to \mathbb{R}$ a smooth function. We will assume the critical points of μ are isolated and that we are in the generic situation: the critical points of μ are non-degenerate. The Morse complex of μ is a chain complex that allows us to compute the homology of M with coefficients in $\mathbb{K}$. As a $\mathbb{K}$-module it is given by $\bigoplus_{p\in \mathrm{Crit}(\mu)} \mathbb{K}\langle p\rangle$. The grading is obtained by looking at the gradient flow for μ as follows: for each critical point p of μ, we can consider the union of the flow lines going into p, called the stable locus, and out of p, the unstable locus of p; we define the degree of p to be $\deg(p) = \dim(\mathrm{Unstable}_p)$. If we have two critical points p and q whose degrees differ by one, $\deg(q) = \deg(p) - 1$, there will be a finite number c_{pq} of flow lines going from p to q, and we define the differential by $d_\mu(p) = \sum_{\deg(q)=\deg(p)-1} c_{pq}q$.

Now back to Fukaya v0.2: As we said before, we want the Morse function μ on the space of paths from a point on L_1 to a point on L_2 to be something like length, in that the critical points, in the case of transverse intersection should be given simply by $L_1 \cap L_2$. The Morse differential d_μ will count holomorphic disks[2] in Y: given $p, q \in L_1 \cap L_2$, the coefficient of q in $d_\mu(p)$ will be the number of holomorphic maps from a disk into Y that send two specified points a and b on the boundary to p and q, and such that the two arcs of the boundary of the disk going between a and b are sent to L_1 and L_2 respectively. We want a finite dimensional space of such holomorphic disks which will happen if $\deg(q) = \deg(p) - 1$.

Composition in this category is given by counting holomorphic triangles. Given three Lagrangians L_1, L_2 and L_3 and points of intersection p_{12}, p_{23}, p_{13}, the coefficient of p_{13} in the composition of p_{12} and p_{23} is the number of holomorphic maps from a triangle into Y, with the corner going to p_{12}, p_{23}, p_{13} and the sides of the triangle going to the appropriate Lagrangian (e.g., the side joining p_{12} with p_{23} should be contained in L_2, etc. Notice that this means that composition is more or less induced by composition of paths (cf v0.1).

[2]So at this point we need to assume Y has a complex structure or, at least, an almost holomorphic structure, compatible with the synplectic structure

To cut down on the dimension of the hom spaces, we'd like that when two Lagrangians are related by a time-dependent Hamiltonian, $\phi_H(t)(L_1) = L_2$, they correspond to isomorphic objects in the Fukaya category. If we had that requirement, how many non-isomorphic objects would we have left? Let's bound the number of non-isomorphic Lagrangians *near* a given Lagrangian L. Recall that L has a tubular neighborhood that is symplectomorphic to T^*L; we say L' is near L if it is a section of this embedded copy of T^*L. We can write L' as the graph Γ_θ of a 1-form θ on L and one can check that L' being Lagrangian corresponds exactly to θ being closed. Also, if θ is exact, say $\theta = df$, one can check that L' is Hamiltonian isotopic to L, which we want to imply that L' is isomorphic to L in the Fukaya category. Therefore, $\dim H^1(L, \mathbb{R})$ is an upper bound for the dimension of the space of non-isomorphic Lagrangians near L.

Fukaya category, v0.3. For this version, we actually change the objects: now we will take pairs (L, ξ) of a Lagrnagian and a flat $U(1)$ connection on L. (Even this is not the final form of the objects, more decorations are missing.) For the morphisms, we twist the Morse complex using the connections: $\hom((L_1, \xi_1), (L_2, \xi_2)) = \bigoplus_{p \in L_1 \cap L_2} \mathbb{K}\langle p \rangle \otimes \hom(\xi_1, \xi_2)_p$.

Now we can state the simplest case of homological mirror symmetry:

Suppose Y is a compact symplectic manifold, with a compatible complex structure, and holomorphic volume form Ω. A Lagrangian L in Y is called *special* if $\Omega|_L = e^{i\pi\theta}\Omega_{\mathbb{R}}$ for some constant θ and some real-valued form $\Omega_{\mathbb{R}}$. Now suppose further that Y contains a special Langrangian torus L_1. It follows that Y contains a $\dim H^1(L_1, \mathbb{R})$-dimensional family of special Lagrangian tori. In the best case scenario Y is actually a special Lagrangian bundle over some base B. Let X be the fiberwise dual bundle, i.e., X is the bundle over the same B whose fiber over $b \in B$ is $L_b^\vee := \hom(\pi_1(L_b), U(1))$.

Claim. X has a natural complex structure, is a Calabi-Yau manifold and $D^{\mathrm{b}}\mathrm{Coh}(X) \cong D^{\Pi}\mathrm{Fuk}(Y)$. The equivalence is roughly as follows: given a Lagrangian L in Y assume it intersects every fiber L_b in a single point, that point corresponds to a $U(1)$-local system on $L_b^\vee$, so L gives rises to a family of local systems on X which can be completed to get a coherent sheaf.

A complex structure on Y gives rise to a Bridgeland stability condition on $D^{\Pi}\mathrm{Fuk}(Y)$ and an object (L, ξ) is stable if it has special Lagrangian representative and in that case the θ appearing in the definition of special is the phase of the object according to the stability condition.

Further reading

From here, the next step is to do an example! The case of an elliptic curve (whose mirror also turns out to be an elliptic curve) was worked out by Polishuk and Zaslow. For an introduction to the Fukaya category, see Dennis Auroux's notes. Homological mirror symmetry is described in Kontsevich's famous ICM address. For the case of Calabi-Yau hypersurfaces in $\mathbb{P}^n$, see Nick Sheridan's paper. He also has videos of lectures on the IAS webpage.