NEARBY CYCLES

1. Motivation

The start of this story begins with Milnor fibers. Consider the setting of an analytic map $f : \mathrm{X} \to \mathbf{C}$, where X is a complex manifold. We suppose that f has some sort of singularity at 0 (and only there), and would like to analyze it.

The idea Milnor had was to study the nature of the the singularity at $x \in f^{-1}(0) = \mathrm{X}_0$ by making precise what happens when we look at nearby fibers X_s for $|s| \approx 0$. The Milnor fiber F_x is given by $\mathrm{B}_{\varepsilon,x} \cap \mathrm{X}_s$ for s and ε sufficiently small. This is a fiber of the Milnor fibration $f^{-1}(\mathrm{D}^*) \cap \mathrm{B}_{\varepsilon,x}$.

The homotopy type of the Milnor fiber reveals important information about the singularity. For example, with an isolated hypersurface singularity it has the homotopy type of a wedge of spheres, and the number of spheres matches the Milnor number of the singularity.

The cohomology of the Milnor fiber sees a lot of information as well. In particular, by circling the base of the Milnor fibration we induce a monodromy operator T on the Milnor fibers. Computing the Lefschetz number gives 0 if $x \in \mathrm{X}_0$ is singular, so in particular for singular points the homotopy type is nontrivial and detected by cohomology.

The idea of the nearby cycle functor is to provide a map

$$\Psi_f : \mathbb{D}^b_c(\mathrm{X}, \mathbf{C}) \to \mathbb{D}^b_c(\mathrm{X}_0, \mathbf{C})$$

such that for $x \in \mathrm{X}_0$ we have

$$\mathcal{H}^k(\Psi_f \mathscr{F})_x \simeq \mathrm{H}^k(F_x, \mathscr{F}).$$

Thus, we can see that this quite literally gives information about nearby cycles. The idea is to globalize the previous theory. Note that here constructibility means constructible with respect to a Whitney stratification, so we are using the algebraic nature of the situation here (which is why we can take our sheaves on topological spaces and not worry about it). We will see later this has good properties with respect to perversity if shifted appropriately.

2. Construction and properties

To construct this map, we consider the following diagram.

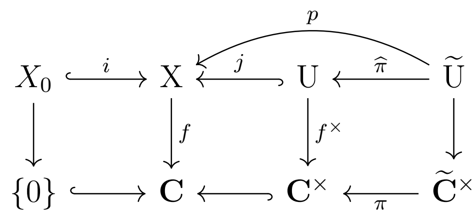

The map π is the infinite cyclic cover of $\mathbf{C}^\times$, and $\widehat{\pi}$ is the lift to U. It is very important later that in this construction that j is an affine morphism: this gives us t-exactness of j_* for the perverse t-structure. This is basically because the zero locus is cut out by a single equation $f = 0$.

A really silly (and wrong) way to define the functor we want is to just take $i^*\mathscr{F}$. But then we have

$$\mathcal{H}^k(i^*\mathscr{F})_x = \mathcal{H}^k(\mathscr{F})_x \simeq \mathrm{H}^k(\mathrm{B}_{\varepsilon,x}, \mathscr{F}).$$

This is of no use if we want to study the singularity.

What we really want is the fiber X_s for small s, so that when we work with a similar calculation as above we get $\mathrm{H}^k(\mathrm{B}_{\varepsilon,x} \cap \mathrm{X}_s, \mathscr{F})$ instead. This is the point of $p = j \circ \widehat{\pi}$.

> Lemma 2.1. The space $\widetilde{\mathrm{U}}$ is homotopy equivalent to the fiber X_s for s small.

Thus, this gives us a model for the generic fiber independent of s. We therefore make the following definition.

> Definition 2.2. We define
> $$\Psi_f\mathscr{F} := i^*\mathrm{R}(p_*p^*)\mathscr{F}.$$
> Here, $p = \widehat{\pi} \circ j$ is the map $\tilde{\mathrm{U}} \to \mathrm{X}$.

The claimed theorem about Milnor fibers is immediate from this construction.

> Proposition 2.3. For $\mathscr{F} \in \mathbb{D}^b_c(\mathrm{X})$ we have
> $$\mathcal{H}^k(\Psi_f\mathscr{F})_x \simeq \mathrm{H}^k(F_x, \mathscr{F}).$$
> Here, F_x is the Milnor fiber.

Proof. This is basically true by construction, assuming Lemma 2.1. ☐

It is not immediately clear that nearby cycles preserve constructibility (i.e. that we actually get a functor on derived category of constructible sheaves). This is because we used a non-algebraic map in the construction. It is a theorem of Deligne that this is fine, and we do actually get a constructible sheaf.

A related construction is the vanishing cycles functor. This wants to see the relative hypercohomology $\mathrm{H}^k(\mathrm{B}_{\varepsilon,x}, \mathrm{B}_{\varepsilon,x} \cap \mathrm{X}_s; \mathscr{F})$. In this sense it is literally seeing vanishing cycles.

The construction of this map is not too terrible either. We have a specialization morphism, given by the adjunction map $\mathscr{F} \to \mathrm{R}p_*p^*\mathscr{F}$ and then applying i^*:

$$i^*\mathscr{F} \to \Psi_f\mathscr{F}.$$

This is called this because if f is *proper* and $\mathscr{F} = \mathbf{C}$ is constant, we get a literal specialization map after applying hypercohomology $\mathrm{H}^k(\mathrm{X}_0, \mathbf{C}) \to \mathrm{H}^k(\mathrm{X}_s, \mathbf{C})$ which agrees with the morphism induced by the specialization map $\mathrm{X}_s \to \mathrm{X}_0$.

The mapping cone of this morphism is $\Phi_f\mathscr{F} : \mathbb{D}^b_c(\mathrm{X}) \to \mathbb{D}^b_c(\mathrm{X}_0)$. This makes it fit into the distinguished triangle

$$i^*\mathscr{F} \longrightarrow \Psi_f\mathscr{F} \xrightarrow{\text{can}} \Phi_f\mathscr{F} \longrightarrow i^*\mathscr{F}[1].$$

Note that this only determines it up to non-unique isomorphism, so more care is needed to see that this is actually functorial. Taking the long exact sequence in hypercohomology, we get the previously mentioned comparison in terms of Milnor fibers for vanishing cycles Φ_f. Namely,

$$\mathcal{H}^k(\Phi_f\mathscr{F})_x = \mathrm{H}^{k+1}(\mathrm{B}_{\varepsilon,x}, \mathrm{B}_{\varepsilon,x} \cap \mathrm{X}_s; \mathscr{F}).$$

This functor also preserves constructibility, as we already asserted implicitly.

Due to the definition in terms of an infinite cyclic covering, there exist monodromy actions on both functors just like we'd expect from a globalization of Milnor's classical theory. In particular, there is an action of $\mathrm{T} \in \pi_1(\mathbf{C}^\times, s)$ on the cohomology of nearby fibers X_s. Giving $i^*\mathscr{F}$ the trivial action, we get a diagram in the derived category:

$$\begin{array}{ccccccc}
i^*\mathscr{F} & \longrightarrow & \Psi_f\mathscr{F} & \xrightarrow{\text{can}} & \Phi_f\mathscr{F} & \longrightarrow & i^*\mathscr{F}[1] \\
\downarrow & & \downarrow{\scriptstyle \mathrm{T}-1} & & \downarrow{\scriptstyle \text{var}} & & \downarrow \\
0 & \longrightarrow & \Psi_f\mathscr{F} & = & \Psi_f\mathscr{F} & \longrightarrow & 0
\end{array}$$

Without the morphism var, the diagram commutes. We can use this to deduce its existence and that $\mathrm{T} - 1 = \text{var} \circ \text{can}$ (in either order).

The nearby and vanishing cycles functors respect perversity, which is important for applications to representation theory. Here, we now let $f : \mathrm{X} \to \mathbf{A}^1_{\mathbf{C}}$ be a regular function where $\mathrm{X}/\mathbf{C}$ is some algebraic variety.

THEOREM 2.4. Suppose that $\mathscr{F}$ is perverse. Then so are $\Psi_f\mathscr{F}[-1]$ and $\Phi_f\mathscr{F}[-1]$. Generally, we have t-exact functors

$$\Psi_f[-1], \Phi_f[-1] : \mathbb{D}_c^b(\mathrm{X}, \mathbf{C}) \to \mathbb{D}_c^b(\mathrm{X}_0, \mathbf{C})$$

for the perverse t-structure. They commute with Verdier duality up to natural isomorphisms.

For this reason, whenever we discuss these from now onward we will shift by $[-1]$. These are the "perverse" nearby and vanishing cycle functors. This of course doesn't affect the aforementioned properties.

There are good properties with respect to the derived tensor product.

THEOREM 2.5. There is a natural map

$$\Psi_f(\mathscr{F}) \otimes^{\mathbb{L}} \Psi_f(\mathscr{G}) \to \Psi_f(\mathscr{F} \otimes^{\mathbb{L}} \mathscr{G})[-1]$$

commuting with the monodromy action. Here, we are using the perverse nearby cycle functor, so everything is $[-1]$-shifted to respect perversity.

The same properties still hold when we define

$$\Psi_{f,\mathrm{U}} := \mathrm{R}(i^* j_* \widehat{\pi}_* \widehat{\pi}^*)[-1]$$

as a functor $\mathbb{D}_c^b(\mathrm{U}) \to \mathbb{D}_c^b(\mathrm{X}_0)$ (again not obvious that this holds). Note that the position where we put R doesn't much matter here, we can actually use $i^*\mathrm{R}j_*\mathrm{R}(\widehat{\pi}_*\widehat{\pi}^*)$ and get the same result. The old functor Ψ_f only depended on $\mathscr{F}|_{\mathrm{U}}$, so this functor is nearly identical and the only real distinction is the domain. All of the previous results stated will hold.

For Beilinson gluing, we will want to use this alternate version, and keep the shift by $[-1]$. This allows us to work things out more easily with perversity, and we will want to work on U so that we can glue from perverse sheaves on U and X_0 (so in particular, we will want to apply nearby cycles on U). This will be the convention from now on: $\Psi_f := \mathrm{R}(i^* j_* \widehat{\pi}_* \widehat{\pi}^*)[-1]$, which will send $\mathrm{Perv}(\mathrm{U}) \to \mathrm{Perv}(\mathrm{X}_0)$.

3. Beilinson gluing

We are initially interested in trying to figure out $\operatorname{Perv}(X)$ from a covering. One reason perverse sheaves are called sheaves when they are obviously complexes of sheaves is that they form a stack of abelian categories. In particular, they behave like sheaves in the sense that they can glue on an open cover. In particular, we can understand them by working locally.

However, a situation where we might not be so lucky is where we are attempting to glue with an open and a closed subscheme. Say $i : X_0 \to X$ is a closed immersion, and $j : U \to X$ is an open immersion as before. In our situation, remember j is affine. We get a diagram

$$\operatorname{Perv}(X_0) \xrightarrow{i_* = i_!} \operatorname{Perv}(X) \underset{j_*}{\overset{j_!}{\rightleftarrows}} \xrightarrow{j^* = j^!} \operatorname{Perv}(U).$$

The construction of $\operatorname{Perv}(X)$ from Z and U requires us to introduce the unipotent nearby cycles and vanishing cycles functors. There will be a third functor which won't appear in the statement, but does appear in the proof called the maximal extension functor.

The unipotent nearby cycles functor is not too complicated. The monodromy operator T induces a decomposition

$$\Psi_f \mathscr{F} = \bigoplus_\lambda \Psi_f^\lambda \mathscr{F}$$

into generalized eigenspaces. The most important is when $\lambda = 1$ which is precisely the unipotent nearby cycles functor (why is this? This eigenspace is when $T - 1$ acts nilpotently).

From now on, write Ψ_f^{un} for unipotent nearby cycles. Similarly, one can write Φ_f^{un}, by again decomposing with the monodromy operator. A large chunck of the content of Beilinson gluing is to be able to reinterpret vanishing cycles as cohomology of a diad, with a new function called the maximal extension functor (which only exists in the unipotent setting).

To omit technical details, we'll skip Beilinson's construction of the maximal extension functor $\Xi_f : \operatorname{Perv}(U) \to \operatorname{Perv}(X)$. What will matter here are its formal properties, as it only appears in the proof of Beilinson gluing.

On unipotent nearby cycles, we have the following distinguished triangle for $\mathscr{F}$ perverse on U:

$$i^* j_* \mathscr{F} \xrightarrow{[1]} \Psi_f^{\mathrm{un}} \mathscr{F} \xrightarrow{\mathrm{T}-1} \Psi_f^{\mathrm{un}} \mathscr{F} \longrightarrow i^* j_* \mathscr{F}[1]$$

This is because the cone of $\mathrm{T}-1$ on $\Psi_f\mathscr{F}$ (not shifted for perversity) is $i^*j_*\mathscr{F}$. This still holds when we consider unipotent nearby cycles.

The natural map $i^*j_*\mathscr{F} \to \Psi_f^{\mathrm{un}}\mathscr{F}[1]$ defines, by adjunction,

$$\mathrm{Hom}^1_{\mathrm{X}_0}(i^*j_*\mathscr{F}, \Psi_f^{\mathrm{un}}\mathscr{F}) = \mathrm{Hom}^1_{\mathrm{X}}(j_*\mathscr{F}, i_*\Psi_f^{\mathrm{un}}\mathscr{F}).$$

This, in turn, defines an object $\Xi_f\mathscr{F}$ sitting in

$$i_*\Psi_f^{\mathrm{un}}\mathscr{F} \longrightarrow \Xi_f\mathscr{F} \longrightarrow j_*\mathscr{F} \longrightarrow i_*\Psi_f^{\mathrm{un}}\mathscr{F}[1]$$

This can serve as an alternate characterization, but does not even show it is functorial (cones are not functorial!) and far less the important properties we summarize below, which are automatic in Beilinson's construction.

The functor Ξ_f is exact and commutes with Verdier duality, preserving perverse sheaves. There are canonical short exact sequences

$$0 \to j_! \to \Xi_f \to i_*\Psi_f^{\mathrm{un}} \to 0$$

and

$$0 \to i_*\Psi_f^{\mathrm{un}} \to \Xi_f \to j_* \to 0,$$

and $i_*\Psi_f^{\mathrm{un}} \to \Xi_f \to i_*\Psi_f^{\mathrm{un}}$ agrees with the nilpotent action of $T-1$.

We note that $j : \mathrm{U} \to \mathrm{X}$ introduces adjunctions $(j_!, j^*)$ and (j_*, j^*). Now we observe the following lemma, which is one of the key ideas.

LEMMA 3.1. We have

$$\Phi_f^{\mathrm{un}}\mathscr{F} := \mathrm{H}^0\left(\begin{array}{ccccc} & & \Xi_f j^*\mathscr{F} & & \\ & \nearrow & & \searrow & \\ j_!j^*\mathscr{F} & & \oplus & & j_*j^*\mathscr{F} \\ & \searrow & & \nearrow & \\ & & \mathscr{F} & & \end{array}\right)$$

where all maps are the canonical ones. On the bottom, they come from adjunctions, and that j^* sends them to the identity.

By this we mean the middle homology of the associated complex

$$0 \to j_!\mathscr{F} \to \Xi_f\mathscr{F} \oplus \mathscr{F} \to j_*\mathscr{F} \to 0.$$

This is because we start the complex in degree -1. In general, we call a diagram of this sort a *diad*.

As given, this shows unipotent vanishing cycles is an exact functor: there is only H^0. Note that the fact that j^* sends the bottom maps to the identity means we can compute $j^*\Phi_f^{un}\mathscr{F} = 0$ from this result as a sanity check. That is, it is supported on X_0, and indeed gives vanishing cycles.

With the technical details out of the way, we are ready to state and sketch a proof of Beilinson gluing.

> Theorem 3.2 (Gluing). The category $\mathrm{Perv}(X)$ is equivalent to the category $\mathrm{Glue}(U, X_0)$ given by
>
> $$\left\{ \mathscr{F}_U \in \mathrm{Perv}(U), \mathscr{F}_0 \in \mathrm{Perv}(X_0) : \begin{array}{ccc} \Psi_f^{un}\mathscr{F}_U & \xrightarrow{T-1} & \Psi_f^{un}\mathscr{F}_U \\ & \searrow \qquad \nearrow & \\ & \mathscr{F}_0 & \end{array} \right\}.$$
>
> We send $\mathscr{F} \mapsto (\mathscr{F}|_U, \Phi_f^{un}\mathscr{F})$ in other direction. In the other direction, associate to the diagram and our pair the H^0 of
>
> $$0 \longrightarrow i_*\Psi_f^{un}\mathscr{F}_U \longrightarrow \Xi_f(\mathscr{F}_U) \oplus i_*\mathscr{F}_0 \longrightarrow i_*\Psi_f^{un}\mathscr{F}_U \longrightarrow 0$$
>
> using that $i_*\Psi_f^{un} \to \Xi_f \to i_*\Psi_f^{un}$ is $T-1$ to complete the triangle into a diad.

Proof. The idea is the following: the gluing category clearly consists of diads, and $\mathrm{Perv}(X)$ can be made into a category of diads in a trivial way. There is an operation r on a diad D given by

$$D = \begin{array}{ccccc} & & A & & \\ & \overset{c_-}{\nearrow} & & \overset{c_+}{\searrow} & \\ C_- & & \oplus & & C_+ \\ & \searrow & & \nearrow & \\ & & B & & \end{array}$$

and replaces it with

$$r(D) = \begin{array}{ccccc} & & A & & \\ & \nearrow & & \searrow & \\ \ker c_+ & & \oplus & & \operatorname{coker} c_- \\ & \searrow & & \nearrow & \\ & & H^0(D) & & \end{array}.$$

It turns out $r^2 = 1$ for formal reasons. This will give our equivalence once we identify our functors in either direction with reflection functors, and then of course they give an equivalence.

As a diad category,

$$\mathrm{Perv}(\mathrm{X}) = \begin{array}{ccccc} & & \Xi_f(j^*\mathscr{F}) & & \\ & \nearrow & & \searrow & \\ j_!j^*\mathscr{F} & & \oplus & & j_*j^*\mathscr{F} \\ & \searrow & & \nearrow & \\ & & \mathscr{F} & & \end{array}$$

where we ask for diads with $\mathscr{F}$ perverse on X and the bottom two maps are isomorphisms on U.

Now, apply r. We get diads of the form

$$\begin{array}{ccccc} & & \Xi_f(j^*\mathscr{F}) & & \\ & \nearrow & & \searrow & \\ i_*\Psi_f^{\mathrm{un}}(j^*\mathscr{F}) & & \oplus & & i_*\Psi_f^{\mathrm{un}}(j^*\mathscr{F}) \\ & \searrow & & \nearrow & \\ & & i_*\Phi_f^{\mathrm{un}}(\mathscr{F}) & & \end{array}$$

where again $\mathscr{F}$ is a perverse sheaf on X. Here, we've used the lemma about unipotent vanishing cycles, as well as the distinguished triangle the maximal extension functor sits in.

Recalling the composition on top comes out to $\mathrm{T} - 1$ canonically, we have obtained the result of Perv(X) to the gluing category $\mathrm{Glue}(\mathrm{U}, \mathrm{X}_0)$: it is simply a reflection functor. We interpret this as a diad category using diagrams of the above form with $j^*\mathscr{F}$ being replaced by $\mathscr{F}_{\mathrm{U}}$ and the result of vanishing cycles being $i_*\mathscr{F}_0$. Because we've passed to perverse sheaves on X, we ask that the support of the sheaf on the bottom is in X_0 in our description.

Applying the reflection functor again, by construction we obtain our proposed inverse. This sends us back to Perv(X) interpreted as a diad category. In particular, since r is an involution and our functors are interpreted as exactly this involution, we get an equivalence. □

Consider the following example of gluing with $\mathbf{P}^1_{\mathbf{C}}$.

LEMMA 3.3. Let Λ be the stratification $0 \cup \mathbf{A}^1_{\mathbf{C}}$. Then we have

$$\mathrm{Perv}_\Lambda(\mathbf{P}^1_{\mathbf{C}}) \simeq \{a : \mathrm{V}_1 \to \mathrm{V}_0, b : \mathrm{V}_0 \to \mathrm{V}_1 | ab = 0\}.$$

Here, the V_i are just arbitrary vector spaces and a, b are linear maps. The vector space V_1 stands for nearby cycles at 0, and V_0 stands for vanishing cycles at 0.

To give an example, the constant sheaf is $\mathbf{C} \to 0, 0 \to \mathbf{C}$. The sheaf $j_* \underline{\mathbf{C}}_{\mathbf{A}^1}[1]$ is $0 : \mathbf{C} \to \mathbf{C}, 1 : \mathbf{C} \to \mathbf{C}$.

A really important property of perverse sheaves applying this is the following theorem of Beilinson.

THEOREM 3.4. There is an equivalence of categories $\mathbb{D}^b(\mathrm{Perv}(\mathrm{X})) \simeq \mathbb{D}^b_c(\mathrm{X})$ via the canonical realization functor.

This is very false if you pick a stratification! For example, try the trivial stratification on $\mathbf{P}^1_{\mathbf{C}}$. Certainly $\mathbb{D}^b(\mathbf{C}{-}\mathsf{Vect})$ is not the same as $\mathbb{D}^b_c(\mathbf{P}^1_{\mathbf{C}})$, since the latter has $\mathrm{Ext}^2(\mathbf{C}, \mathbf{C}) \neq 0$ as we can identify it with H^2 and there's a nontrivial cohomology class (e.g. $c_1(\mathcal{O}(1))$). The fault comes down to a gap in extensions, which goes away once we are allowed to refine stratifications.

In fact, to check this equivalence, it suffices to check extensions are the same on both sides. In particular, the fact we use is that if our functor induces an isomorphism on the hearts and induces isomorphisms between all higher extensions, then it is an equivalence.

This is where gluing comes in. Once we check that the equivalence of Ext groups holds at the generic point, the gluing construction allows us to use induction. In particular, when supports lie in $f^{-1}(0) = \mathrm{X}_0 \subseteq \mathrm{X}$, the idea is to use vanishing cycles to get an isomorphism on extension groups inductively. One shows this by identifying the Yoneda extension group $\mathrm{Ext}^i(M, N)$ with connected components of a category $\mathrm{E}^i(M, N)$, and using the facts about unipotent nearby cycles to produce a path to deduce the connected component.