Perverse Sheaves

Jackson Hopper

27 July, 2020

Abstract

Perverse sheaves are a powerful tool for understanding topology of algebraic maps, and have far-reaching applications in other areas including representation theory. I hope to introduce the concept of pure perverse sheaves with a thorough definition and some standard results, including their classification and the decomposition theorem of [2]. Then, time permitting, I will explain how the concept can be expanded to the statement of the geometric Satake correspondence, which for a reductive group G isomorphically relates the category of L^+G-equivariant perverse sheaves on the affine grassmannian of G to the representation semiring of the dual group $\hat{G}$.

1 Introduction

We will be talking about perverse sheaves, particularly in the setting of étale topology with $\overline{\mathbb{Q}}_\ell$-adic coefficients on separated schemes of finite type (hence, wlog, varieties) over a finite field. In this setting we have the additional advantage of the theory of weights. The goal is to summarize the basic theory, especially the classification and decomposition theorems, and to discuss how this theory can be applied to a Satake correspondence for reductive groups. Hopefully along the way I've gotten a couple of people interested.

A couple of strictly observed conventions for the first few sections:

- Sheaves and cohomology are always étale.
- Schemes of interest are always separated and finite type over a field.
- Given above, unless otherwise specified, schemes of interest are assumed reduced, and I will typically call them varieties. In particular, closed subschemes are always given reduced closed structure.
- Once defined, coefficients are ℓ-adic. I use the phrase ℓ-adic somewhat broadly, where other authors might use "adic" or "π-adic." (In particular, there is no residue field of order ℓ or uniformizer called ℓ for the integers of $\overline{\mathbb{Q}}_\ell$, but I will refer to $\overline{\mathbb{Q}}_\ell$-coefficients as ℓ-adic.) Someone correct me if that is an embarrassing faux pas, but I don't like making so many distinctions. $\ell \neq p$ will hold with all adic notions in use in this talk.
- Correspondingly, we have $q = p^r$ a prime power, and ℓ a distinct prime number. p and ℓ may be assumed fixed.
- I will only use an underline on constant sheaves when they are being literally used as a structure sheaf of rings. This is mainly to distinguish from the case where they are being figuratively used as a structure sheaf, but also to save on writing underlines in other cases. It's not mathematically justifiable, but I like the result.

- If not otherwise specified, a stratification $\{S_i\}$ of a space X is a partition of X into finitely many locally closed smooth subspaces S_i of dimension i such that $\overline{S}_i = \bigcup_{j \geq i} S_j$
- The only perversity in consideration is the middle perversity
- In discussions of purity, when working specifically over a finite field I will use a subscript 0 on the left. When passing to the algebraic closure (usually via base change), the subscript will be dropped. Passing to a larger finite field, I will use subscript of a larger integer. Additionally, I may use a subscript 0 to track the base field of a Galois action. (e.g. "Let ${}_0F$ generate the Weil group $W(\overline{\mathbb{F}}_q/\mathbb{F}_q)$, then let ${}_1F = {}_0F^n$ generate the Weil group $W(\overline{\mathbb{F}}_q/\mathbb{F}_{q^n})$")

My approach essentially is the one followed by [2], but it is strongly informed by [10], and the last section is based on [14]. The best overview—the thing I'm directly competing with here, and can't compare to—is [4].

1.1 Intersection complexes: Elementary definition

Perverse sheaves form an abelian category inside the "derived" category $D_c^b(X, \overline{\mathbb{Q}}_\ell)$. The category is both noetherian and artinian. As far as I know, there is no way to define this category without defining its t-structure. That said, if we are willing to accept the existence of a category called $D_c^b(X, \overline{\mathbb{Q}}_\ell)$, and not too worried to think of its objects as quasi-isomorphism classes of complexes of sheaves on X, then we could take the classification theorem of [2] as a definition of the *simple* perverse sheaves—the intersection complexes. We may make this definition elementarily with a formula using only derived direct images and truncations.

Let X be an integral variety of dimension d, and let $\{S_i\}_{i=0}^d$ be a stratification such that

- All S_i are smooth, locally closed subvarieties of X
- $\overline{S}_{i-1} \supset S_i$ for all $1 \leq i \leq d$
- $\dim S_i = i$

Given such a stratification, let $U_i = \bigcup_{k=i}^d S_k$. Let $j : U_d \hookrightarrow X$ and $j_i : U_i \hookrightarrow U_{i-1}$ be open immersions. Note that all such immersions are dense.

Now let $\mathcal{L}$ be a lisse ℓ-adic sheaf on U_d. If we like, we may consider $\mathcal{L}$ to be a continuous representation of $\pi_1(U_d, \bar{x})$ over $\overline{\mathbb{Q}}_\ell$ for some geometric point $\bar{x}$ lying over a closed point $x \in |U_d|$. Then we have the

Definition. The intersection complex $IC_X(\mathcal{L})$ is the complex in $D_c^b(X, \overline{\mathbb{Q}}_\ell)$ given by the formula

$$IC_X(\mathcal{L}) := \big(\tau_{\leq d-1} R(j_1)_* \cdots \tau_{\leq 1} R(j_{d-1})_* \tau_{\leq 0} R(j_d)_*(\mathcal{L})\big)[d]. \tag{1}$$

If $i : X \hookrightarrow Y$ is a closed immersion, then we have $IC_X(\mathcal{L}) \in D_c^b(Y, \overline{\mathbb{Q}}_\ell)$ via the direct image

$$Ri_* : D_c^b(X, \overline{\mathbb{Q}}_\ell) \to D_c^b(Y, \overline{\mathbb{Q}}_\ell).$$

A formula like (1) holds if U_d is any locally closed integral subvariety of X, but it reduces to simply pushing forward in the case U_d is closed. Thus the simplest way to define an intersection complex on U_d is always to apply formula (1) to define the complex to $\overline{U}_d$ and then push forward to X.

All other perverse sheaves are iterated extensions of objects like the above. Unfortunately, without defining the category in which they live, there is no way to intrinsically define what such an extension might look like.

To put it another way: being artinian, objects in the category of perverse sheaves are completely decomposable. Unfortunately, I am not aware of a nice and general characterization of indecomposable objects. See the decomposition theorem (§4) for amplification.

2 Definition and examples

Here we will define the derived category in which perverse sheaves live (sorry, that is a pain), and the perverse t-structure on that category (a little bit less of a pain). I will close with 3 essential examples.

2.1 The constructible category $\mathrm{Sh}_c(X, \underline{R})$

Before we get to the really hard definitions, I would like to establish the category $\mathrm{Sh}_c(X, \underline{R})$ for a noetherian ring R and finite-type scheme X. While this is a fairly easily-definable full subcategory of $\underline{R}$-modules on X, it can also be formed as an inductive limit over stratifications, and we will follow that induction to define a t-structure on it.

This category is an essential component of the framework of the Weil conjectures, as it contains the locally constant sheaves, and has some stability properties that category doesn't. In particular, we have the motivating

Theorem 2.1.1. *If $\mathcal{F}$ is a constructible sheaf over X and $f : X \to Y$ is a morphism of separable schemes of finite type over a field k, then $R^i f_! \mathcal{F}$ and $R^i f_* \mathcal{F}$ are constructible sheaves over Y for all i.*

If U is a dense open subvariety of a smooth variety X with inclusion j, then it is not even the case in general that $R^i j_* \mathbb{Q}_U$ is locally constant—and $j_! \mathbb{Q}_U$ is only locally constant if $U = X$.

Definition. A constructible sheaf $\mathcal{F}$ over X is an étale sheaf such that for some finite, locally closed, smooth stratification $X = \coprod S_i$, each restriction $i^*_{S_i} \mathcal{F}$ is a locally constant, finite $\underline{R}$-module. In the case $R = \mathbb{Z}$, this means finite as a set.

This condition can easily be verified locally, or by supports of sections.

Its objects can also be formed as the inductive limits of $(\{S_i\}, \mathcal{J})$-constructible sheaves, where $\{S_i\}$ is a finite smooth, locally closed stratification of X and $\mathcal{J}$ is a function assigning to each stratum S_i a finite set of finite irreducible locally constant sheaves $\mathcal{J}(S_i) \subset \mathrm{Sh}_{lc}(S_i, \underline{R})$. The introduction of $\mathcal{J}$ aids in inductive proofs—but does not alter the intuitive meaning of "constructible" (and is, strictly speaking, not necessary).

2.2 ℓ-adic constructible sheaves and their "derived" category $D^b_c(X, \overline{\mathbb{Q}}_\ell)$

Here we introduce the limiting processes used to form the category known as $D^b_c(X, \overline{\mathbb{Q}}_\ell)$.

Attempting to construct a triangulated category with appropriate natural t-structure, finiteness, and reasonable definitions of the 6 functors of Verdier duality, but also with coefficients in an algebraically closed field of characteristic 0, is a minefield for naive expectations. The one break we will catch has to do with the assumptions on the base field: because we are primarily working over algebraically closed or finite fields, we have finite Galois cohomology groups where appropriate. Unfortunately, to take advantage of this fact, we will have to introduce our most unfriendly tactic, the subcategory bounded tor-finite complexes.

Considering the manifold technical difficulties and workarounds, and my limited time to share them, I will only summarize a process known to be successful in producing $D^b_c(X, \overline{\mathbb{Q}}_\ell)$, providing neither full justification for nor sufficiency of sidesteps. Be comforted also by the words of [3] (whose approach I will not be following here):

"Nonetheless, in daily life, one pretends (without getting into much trouble) that $D^b_c(X, \overline{\mathbb{Q}}_\ell)$ is simply the full subcategory of some hypothetical derived category $D(X, \overline{\mathbb{Q}}_\ell)$ of all $\overline{\mathbb{Q}}_\ell$-sheaves spanned by those bounded complexes whose cohomology sheaves are locally constant along a stratification."

Outline:

- define ℓ-adic sheaves
- define the tor-finite category
- only now expand discussion to all characteristic $(0, \ell)$ local fields/integers
- discuss the projective systems of derived categories $\mathcal{O}_{\pi,r} \rightsquigarrow \mathcal{O}_\pi$
- define the 2-projective limit
- localize to extend scalars to fields $\mathcal{O}_\pi \rightsquigarrow E_\pi$
- discuss the inductive system, and define the inductive limit $E_\pi \rightsquigarrow \overline{\mathbb{Q}}_\ell$

First, we have to try and make sense of the cohomology sheaves of the complexes we'll be making. An ℓ-adic constructible sheaf is the projective limit of an ℓ-adic system of constructible sheaves $\mathcal{F}_r \in \mathrm{Sh}_c(X, \underline{\mathbb{Z}/\ell^r\mathbb{Z}})$. For us, a projective system $(\mathcal{F}_r)$ is *ℓ-adic* if for all $r \geq s$, we have

$$\mathcal{F}_r \otimes_{\mathbb{Z}/\ell^r\mathbb{Z}} \mathbb{Z}/\ell^s\mathbb{Z} = \mathcal{F}_r/\ell^{r-s}\mathcal{F}_r = \mathcal{F}_s.$$

Given that a system $(\mathcal{F}_r)$ is ℓ-adic, it is *lisse* if for each r the sheaf $\mathcal{F}_r$ is locally constant. This concept is critical to the classification of perverse sheaves. It is also geometrically meaningful, as a lisse ℓ-adic sheaf corresponds to a continuous representation of the étale fundamental group (see [12] for details).

Next, we will discuss our building block categories which are actually derived. First, $D_c(X, \underline{R})$ is the full subcategory of all complexes K of sheaves of $\underline{R}$-modules on X whose *cohomology sheaves* H^iK are constructible sheaves on X. (It may be the case that K has no quasi-isomorphic representative complex $K^\bullet$ such that all K^i are constructible.)

In order to pass from $D_c(X, \underline{\mathbb{Z}/\ell^r\mathbb{Z}})$ to $D_c(X, \mathbb{Z}_\ell)$, we need to further restrict to a nice subcategory of $D_c(X, \underline{\mathbb{Z}/\ell^r\mathbb{Z}})$ in order to take advantage of the Mittag-Leffler condition on terms of the complexes in the projective system, so that the limit functor is exact.

The bounded constructible tor-finite category $D^b_{ctf}(X, \underline{\mathbb{Z}/\ell^r\mathbb{Z}})$ is defined as the strictly full subcategory consisting of complexes represented by $K^\bullet$ where $K^\bullet$ is bounded (not just cohomologically bounded), each nonzero K^i is a flat $\underline{\mathbb{Z}/\ell^r\mathbb{Z}}$-module over X, and the cohomology sheaves $H^i(K^\bullet)$ are constructible.

Unfortunately, this subcategory does not inherit the natural t-structure from $D^b_c(X, \underline{\mathbb{Z}/\ell^r\mathbb{Z}})$, as the truncation functors $\tau_{\leq 0}$ and $\tau_{\geq 0}$ do not preserve flatness. So once we define the derived category $D^b_c(X, \mathbb{Z}_\ell)$, we still have to manually define cohomology sheaves and from there a "natural" t-structure.

But before defining $D^b_c(X, \mathbb{Z}_\ell)$, I want to expand the discussion somewhat to all characteristic-ℓ local fields, rather than just $\mathbb{Q}_\ell$, with the hope of efficiently reaching coefficients in $\overline{\mathbb{Q}}_\ell$. So for a finite extension

$E_\pi/\mathbb{Q}_\ell$, we have uniformizer π and ring of integers $\mathcal{O}_\pi$. We will refer to the quotient $\mathcal{O}_\pi/\pi^r\mathcal{O}_\pi$ as $\mathcal{O}_{\pi,r}$ for convenience, i.e.

$$\lim_{\substack{\leftarrow \\ r}} \mathcal{O}_{\pi,r} = \mathcal{O}_\pi.$$

All notions of ℓ-adic sheaves so far go through with cosmetic changes.

An object of $D_c^b(X, \mathcal{O}_\pi)$ is a projective system of complexes

$$K := \lim_{\substack{\leftarrow \\ r}} K_r,$$

where each $K_r \in D_{ctf}^b(X, \underline{\mathcal{O}}_{\pi,r})$, along with quasi-isomorphisms

$$\phi_{r+1} : K_{r+1} \otimes^L_{\mathcal{O}_{\pi,r+1}} \mathcal{O}_{\pi,r} \xrightarrow{\sim} K_r.$$

The existence of (ϕ_{r+1}) ensures that the system is ℓ-adic.

To define the t-structrure on $D_c^b(X, \mathcal{O}_\pi)$, first define the cohomology sheaves of a complex K. For all r, we have a cohomology sheaves $H^i(K_r) \in \mathrm{Sh}_c(X, \underline{\mathcal{O}}_{\pi,r})$. Then they form a projective system with limit

$$H^i(K) := \lim_{\substack{\leftarrow \\ r}} H^i(K_r).$$

Proposition 2.2.1. *The limit H^iK is representable by an ℓ-adic sheaf. That is, in the category of projective systems of sheaves modulo null systems (the Artin–Reese category), there is an ℓ-adic system isomorphic to H^iK.*

The above proposition, the fact that homomorphisms are well-defined in $D_c^b(X, \mathcal{O}_\pi)$ as ℓ-adic systems of homomorphisms, and the fact that this chomology sheaf definition provides a standard t-structure, all depend essentially on finiteness of Galois cohomology groups, and of the fibers of the complexes themselves. These theorems are overwhelmingly technical, and ultimately due to Deligne. See [10] and [8] for some details.

Now, let us localize to field coefficients on our derived category for the penultimate step. This is actually fairly simple. Given a derived category with integral coefficients $D_c^b(X, \mathcal{O}_\pi)$, define $D_c^b(X, E_\pi)$ to have objects $K \otimes^L_{\mathcal{O}_\pi} E_\pi$, where $K \in D_c^b(X, \mathcal{O}_\pi)$, and

$$\mathrm{Hom}\,(K \otimes^L_{\mathcal{O}_\pi} E_\pi, L \otimes^L_{\mathcal{O}_\pi} E_\pi) := \mathrm{Hom}\,(K, L) \otimes_{\mathcal{O}_\pi} E_\pi.$$

Finally, we pass from coefficients in a local field E_π of characteristic $(0, \ell)$ to the algebraic closure $\overline{\mathbb{Q}}_\ell$. Because the system here is inductive, the closest we have to a complication is keeping track of ramification. Suppose E_ρ/E_π is a finite extension with uniformizer ρ and ring of integers $\mathcal{O}_\rho$. Let e be the ramification index, so $\pi\mathcal{O}_\rho = \rho^e\mathcal{O}_\rho$. For a positive integer r, we have

$$\mathcal{O}_{\rho,re} = \mathcal{O}_\rho/\rho^{re}\mathcal{O}_\rho = \mathcal{O}_\rho/\pi^r\mathcal{O}_\rho,$$

and so since $\mathcal{O}_\rho$ is free over $\mathcal{O}_\pi$, we have $\mathcal{O}_{\rho,re}$ is free over $\mathcal{O}_{\pi,r}$. And so we have naturally defined functors

$$\begin{aligned} D_{ctf}^b(X, \underline{\mathcal{O}}_{\pi,r}) &\longrightarrow D_{ctf}^b(X, \underline{\mathcal{O}}_{\rho,re}) \\ K &\longmapsto K \otimes^L_{\mathcal{O}_{\pi,r}} \mathcal{O}_{\rho,re}, \end{aligned}$$

which easily defines functors $D_c^b(X, \mathcal{O}_\pi) \to D_c^b(X, \mathcal{O}_\rho)$ and $D_c^b(X, E_\pi) \to D_c^b(X, E_\rho)$. Therefore we can take $D_c^b(X, \overline{\mathbb{Q}}_\ell)$ to have objects of the form $K \otimes^L_{E_\pi} \overline{\mathbb{Q}}_\ell$ for $K \in D_c^b(X, E_\pi)$, where E_π is some finite extension of $\mathbb{Q}_\ell$. And similarly, given $K \in D_c^b(X, E_\pi)$ and $L \in D_c^b(X, E_\rho)$, we have them both defined (after base change) over the composite E_σ, and

$$\mathrm{Hom}\,(K_\sigma \otimes^L_{E_\sigma} \overline{\mathbb{Q}}_\ell, L_\sigma \otimes^L_{E_\sigma} \overline{\mathbb{Q}}_\ell) := \mathrm{Hom}\,(K, L) \otimes_{E_\sigma} \overline{\mathbb{Q}}_\ell$$

2.3 t-structures interlude

I've mentioned t-structures a lot by this point. So far, I've been using the natural t-structure, which I will actually define soon. But I want to first talk a bit about working in a triangulated category, with the assumption the reader is somewhat familiar with derived homological algebra. (But not trinagulated categories…Sorry.)

What is essential for us is that a triangulated category $\mathcal{D}$ is additive and not abelian, it has an autoequivalence $(\cdot)[1] : \mathcal{D} \to \mathcal{D}$ called the *shift functor*, and it has a distinguished class of "triangles," or sextuples (A, B, C, f, g, h) (typically referred to as triples (A, B, C)) forming diagrams

$$A \xrightarrow{f} B \xrightarrow{g} C \xrightarrow{h} A[1].$$

This distinguished class of triangles is called either *distinguished triangles* or *exact triangles*, and is strictly smaller than the class of all triangles. There are 4 axioms about them, which I will mention only if I use one. It might help in visualizing triangles view them like so:

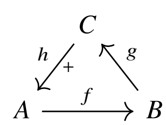

It may also help to know that if the triangle is exact, then C is called a *cone* of f and f is the *base* of the triangle. Additionally, (A, B, C, f, g, h) is exact if and only if $(B, C, A[1], g, h, -f)$ is.

A *cohomological functor* is a functor from a triangulated category to an abelian category sending exact triangles to long exact sequences. If the triangulated category is the derived category of an abelian category, say $\mathcal{D} = D(\mathcal{A})$, then the cohomology functor $H^0 : D(\mathcal{A}) \to \mathcal{A}$ is cohomological.

Lemma 2.3.1. *If (K, L, M) is an exact triangle in $\mathcal{D}$ such that* $\mathrm{Hom}\,(K, M) = 0$*, then (K, L, M) is the unique exact triangle in $\mathcal{D}$ with $K \to L$ as a base.*

Now we are ready to define t-structures.

Definition. A t-structure on a triangulated category $\mathcal{D}$ is a pair $(\mathcal{D}^{\leq 0}, \mathcal{D}^{\geq 0})$ of strictly full subcategories meeting the following criteria (where $\mathcal{D}^{\leq n}$ is defined to be $\mathcal{D}^{\leq 0}[-n]$, and $\mathcal{D}^{\geq n} := \mathcal{D}^{\geq 0}[-n]$.):

1. For all $K \in \mathcal{D}^{\leq 0}$ and $L \in \mathcal{D}^{\geq 1}$, $\mathrm{Hom}\,(K, L) = 0$
2. $\mathcal{D}^{\leq 0} \subset \mathcal{D}^{\leq 1}$ and $\mathcal{D}^{\geq 1} \subset \mathcal{D}^{\geq 0}$
3. For every object $K \in \mathcal{D}$, there is a distinguished triangle $(\tau_{\leq 0}K, K, \tau_{\geq 1}K)$ such that $\tau_{\leq 0}K \in \mathcal{D}^{\leq 0}$ and $\tau_{\geq 1}K \in \mathcal{D}^{\geq 1}$.

Note that although the objects $\tau_{\leq 0}K$ and $\tau_{\geq 1}K$ of axiom 3 are not defined to be unique, it is a consequence of Lemma 2.3.1 and axiom 1 that they are. In fact, $\tau_{\leq 0}$ is a left adjoint to the inclusion functor $\mathcal{D}^{\leq 0} \to \mathcal{D}$, and $\tau_{\geq 1}$ is a right adjoint to $\mathcal{D}^{\geq 1} \to \mathcal{D}$.

We can almost get away with only considering the

Example 2.3.1. Let $\mathcal{A}$ be an abelian category, and let $\delta : \mathcal{A} \to D(\mathcal{A})$ be the inclusion of $\mathcal{A}$ in its derived category. Then we have the *natural t-structure* on $D(\mathcal{A})$, given by

$$D(\mathcal{A})^{\leq 0} := \{K \in D(\mathcal{A}) \mid H^i K = 0 \text{ for } i > 0\}$$
$$D(\mathcal{A})^{\geq 0} := \{K \in D(\mathcal{A}) \mid H^i K = 0 \text{ for } i < 0\}$$

Then the truncation functors $\tau_{\leq 0}$ and $\tau_{\geq 0}$ are those defined by ordinary (good) truncation: If $K^\bullet$ is a complex representing $K \in D(\mathcal{A})$, then $\tau_{\leq 0}K$ and $\tau_{\geq 0}K$ are represented by complexes with the following terms:

$$(\tau_{\leq 0}K)^i = \begin{cases} K^i & i < 0 \\ Z^0(K) & i = 0 \\ 0 & i > 0 \end{cases} \qquad \text{and} \qquad (\tau_{\geq 0}K)^i = \begin{cases} 0 & i < 0 \\ K^0/B^0(K) & i = 0 \\ K^i & i > 0 \end{cases}$$

Given a t-structure on a triangulated category $\mathcal{D}$, there is an abelian subcategory ${}^pC(\mathcal{D})$ called the *heart*, or *core*, of $\mathcal{D}$, defined to be $\mathcal{D}^{\leq 0} \cap \mathcal{D}^{\geq 0}$. We then have a functor ${}^pH^0 : \mathcal{D} \to {}^pC(\mathcal{D})$ given by

$${}^pH^0 = \tau_{\leq 0}\tau_{\geq 0} = \tau_{\geq 0}\tau_{\leq 0}.$$

It is a fact that truncations commute, implying the second equality. ${}^pH^0$ is a cohomological functor.

Proposition 2.3.1. *In the case of the natural t-structure on $D(\mathcal{A})$, we have ${}^pC(D(\mathcal{A})) = \mathcal{A}$, and the heart functor ${}^pH^0$ coincides with cohomology H^0.*

We have a strictly fully faithful inclusion ${}^pC(\mathcal{D}) \to \mathcal{D}$, though it is not the case in general that $D({}^pC(\mathcal{D})) = \mathcal{D}$. Given an exact functor between triangulated categories $F : \mathcal{D} \to \mathcal{D}'$ with t-structures, we have an additive functor ${}^pF$ between abelian categories defined by composition:

$$\begin{array}{ccc} \mathcal{D} & \xrightarrow{F} & \mathcal{D}' \\ \uparrow & & \downarrow {}^pH^0 \\ {}^pC(\mathcal{D}) & \xrightarrow{{}^pF} & {}^pC(\mathcal{D}') \end{array}$$

If F is left (or right) t-exact, meaning it preserves $\mathcal{D}^{\geq 0}$ (or $\mathcal{D}^{\leq 0}$), then ${}^pF$ is left (or right) exact.

Recall that we have defined a cohomology functor H^0 for $D^b_c(X, \overline{\mathbb{Q}}_\ell)$. It is a theorem of Deligne that we can put a "natural" t-structure on this category using this functor, and we have so-called Deligne truncation functors acting somewhat like natural truncation functors. Specifically, we have the

Definition.

$$D^{b,\leq 0}_c(X, \overline{\mathbb{Q}}_\ell) := \{K \in D^b_c(X, \overline{\mathbb{Q}}_\ell) \mid H^i K = 0 \text{ for } i > 0\}$$
$$D^{b,\geq 0}_c(X, \overline{\mathbb{Q}}_\ell) := \{K \in D^b_c(X, \overline{\mathbb{Q}}_\ell) \mid H^i K = 0 \text{ for } i < 0\}$$

In the sequel this t-structure will be referred to as natural.

2.4 Gluing t-structures and perverse sheaves

We didn't define the whole framework of t-structures just to stick to the natural one. What we will produce will be built inductively out of almost-natural t-structures, again using two limiting processes. The basic step of the first induction process is the gluing t-structure: Given a variety X, we form a t-structure on the derived category of constructible sheaves $D^b_c(X, \overline{\mathbb{Q}}_\ell)$ from t-structures on a complementary pair of subschemes of X. I will discuss the primary framework for it, then (doubly) inductively apply it to $D^b_c(X, \overline{\mathbb{Q}}_\ell)$ to find $P(X, \overline{\mathbb{Q}}_\ell)$ in the heart.

Let $j : U \hookrightarrow X$ be a smooth dense open subscheme and let $i : Z \hookrightarrow X$ be a reduced closed immersion onto the complement of U. Then we have 6 functors between categories of schemes in two adjoint sequences:

$$(j_!, j^*, j_*) \qquad \text{and} \qquad (i^*, i_*, i^!)$$

$$\mathrm{Sh}(Z) \underset{i^*,\ i^!}{\overset{i_*}{\rightleftarrows}} \mathrm{Sh}(X) \underset{j_!,\ j_*}{\overset{j^*}{\rightleftarrows}} \mathrm{Sh}(U)$$

In this situation, we can call $j_!$ extension by 0 and $i^!$ inverse image with support in Z. We can also use the adjunctions to declare equalities $j^! := j^*$ and $i_! := i_*$.

We then have the following consequences:

1. $j^* i_* = 0$. By adjunction, we have the equivalent statements $i^* j_! = i^! j_* = 0$, and for sheaves $\mathcal{F}$ over Z and $\mathcal{G}$ over U,
$$\mathrm{Hom}\,(j_!\mathcal{G}, i_*\mathcal{F}) = \mathrm{Hom}\,(i_*\mathcal{F}, j_*\mathcal{G}) = 0$$

2. All inverse images are essentially surjective, and all direct images are fully faithful. Equivalently, the following transformations derived from adjunctions are isomorphisms of functors:
$$i^* i_* \xrightarrow{\sim} \mathrm{id}_{\mathrm{Sh}(Z)} \xrightarrow{\sim} i^! i_*$$
$$j^* j_* \xrightarrow{\sim} \mathrm{id}_{\mathrm{Sh}(U)} \xrightarrow{\sim} j^* j_!$$

3. Given a sheaf $\mathcal{F} \in \mathrm{Sh}(X)$, we have the following exact sequences derived from adjunctions:
$$0 \longrightarrow j_! j^* \mathcal{F} \longrightarrow \mathcal{F} \longrightarrow i_* i^* \mathcal{F} \longrightarrow 0$$
$$0 \longrightarrow i_* i^! \mathcal{F} \longrightarrow \mathcal{F} \longrightarrow j_* j^* \mathcal{F}$$

Passing to the derived categories $D(Z)$, $D(X)$, and $D(U)$ we can rewrite all functors as derived functors ($Rj_!$ etc.), and the above adjunctions and enumerated consequences still hold (reformulating 3. with exact triangles). In fact, we can take them as axioms for a framework in which we can glue t-structures on any two triangulated categories $\mathcal{D}_Z$ and $\mathcal{D}_U$ to form one on a third triangulated category $\mathcal{D}_X$.

Theorem 2.4.1. *Let $\mathcal{D}_Z \xrightarrow{Ri_*} \mathcal{D}_X \xrightarrow{Rj^*} \mathcal{D}_U$ be as above, and let $(\mathcal{D}_Z^{\leq 0}, \mathcal{D}_Z^{\geq 0})$ and $(\mathcal{D}_U^{\leq 0}, \mathcal{D}_U^{\geq 0})$ be t-structures on $\mathcal{D}_Z$, $\mathcal{D}_U$, respectively. Then we have subcategories of $\mathcal{D}_X$ defined as follows:*

$$\mathcal{D}_X^{\leq 0} = \{K \in \mathcal{D}_X \mid Rj^*K \in \mathcal{D}_U^{\leq 0} \text{ and } Ri^*K \in \mathcal{D}_Z^{\leq 0}\}$$
$$\mathcal{D}_X^{\geq 0} = \{K \in \mathcal{D}_X \mid Rj^*K \in \mathcal{D}_U^{\geq 0} \text{ and } Ri^!K \in \mathcal{D}_Z^{\geq 0}\}$$

These subcategories form a t-structure on $\mathcal{D}_X$, called the gluing t-structure.

Now we can inductively define the perverse t-structure relative to a smooth stratification by dimensions $X = \{S_i\}_{i=0}^{d}$ and function $\mathcal{J}$ assigning to each stratum a finite set of irreducible objects in $\mathrm{Sh}_c(S_i, \overline{\mathbb{Q}}_\ell)$. We have for all i, $\overline{S_i} = \bigcup_{k\leq i} S_i$. Then a complex $K \in D_c^b(X, \overline{\mathbb{Q}}_\ell)$ is in $D^b_{\{S_i\},\mathcal{J}}(X, \overline{\mathbb{Q}}_\ell)$ if and only if all of the cohomology sheaves of restrictions to strata $H^j i^*_{S_i} K$ are finite successive extensions by simple objects in $\mathcal{J}(S_i)$. This category is known as the $(\{S_i\}, \mathcal{J})$-constructible complexes over X.

Definition. On each stratum S_i, let ${}^pD^{b,\leq 0}_{\{S_i\},\mathcal{J}}(S_i, \overline{\mathbb{Q}}_\ell)$ be a shift to the left of the natural t-structure on $D^b_{\{S_i\},\mathcal{J}}(S_i, \overline{\mathbb{Q}}_\ell)$ by $i = \dim S_i$. We have $K \in {}^pD^{b,\leq 0}_{\{S_i\},\mathcal{J}}(S_i, \overline{\mathbb{Q}}_\ell)$ if and only if $H^j K = 0$ for all $j > -i$. Similarly, we have $K \in {}^pD^{b,\geq 0}_{\{S_i\},\mathcal{J}}(X, \overline{\mathbb{Q}}_\ell)$ if and only if $H^j K = 0$ for all $j < -i$.

Definition. We form the perverse t-structure relative to the stratification $\{S_i\}$ by gluing these t-structures inductively using the formalism above. We have the perverse and natural t-structures coincide on $D^b_{\{S_i\},\mathcal{J}}(S_0, \overline{\mathbb{Q}}_\ell)$. For $i \geq 1$ we have $\overline{S}_i = S_i \amalg \overline{S}_{i-1}$, with S_i open and $\overline{S}_{i-1}$ closed in $\overline{S}_i$. Also let $j_i : S_i \to \overline{S}_i$ and let $k_i : \overline{S}_{i-1} \to \overline{S}_i$. Then let

$$ {}^pD^{b,\leq 0}_{\{S_i\},\mathcal{J}}(\overline{S}_i, \overline{\mathbb{Q}}_\ell) := \{K \in D^b_{\{S_i\},\mathcal{J}}(S_i, \overline{\mathbb{Q}}_\ell) \mid Rj_i^* K \in D^{b,\leq -i}_{\{S_i\},\mathcal{J}}(S_i, \overline{\mathbb{Q}}_\ell) \text{ and } Rk_i^* K \in {}^pD^{b,\leq 0}_{\{S_i\},\mathcal{J}}(\overline{S}_{i-1}, \overline{\mathbb{Q}}_\ell)\} $$

$$ {}^pD^{b,\geq 0}_{\{S_i\},\mathcal{J}}(\overline{S}_i, \overline{\mathbb{Q}}_\ell) := \{K \in D^b_{\{S_i\},\mathcal{J}}(S_i, \overline{\mathbb{Q}}_\ell) \mid Rj_i^* K \in D^{b,\geq -i}_{\{S_i\},\mathcal{J}}(S_i, \overline{\mathbb{Q}}_\ell) \text{ and } Rk_i^! K \in {}^pD^{b,\geq 0}_{\{S_i\},\mathcal{J}}(\overline{S}_{i-1}, \overline{\mathbb{Q}}_\ell)\}. $$

Then a complex $\mathcal{F}$ over $X = \overline{S}_d$ is called a *perverse sheaf* relative to the stratification $\{S_i\}$ if it is in the heart of the perverse t-structure ${}^pD^{b,\leq 0}_{\{S_i\},\mathcal{J}}(X, \overline{\mathbb{Q}}_\ell) \cap {}^pD^{b,\geq 0}_{\{S_i\},\mathcal{J}}(X, \overline{\mathbb{Q}}_\ell)$. This category will be denoted $P_{\{S_i\},\mathcal{J}}(X, \overline{\mathbb{Q}}_\ell)$. Truncation for the perverse t-structure is denoted ${}^p\tau_{\leq 0}$ and perverse cohomology is denoted ${}^pH^0$.

Thus concludes the first induction of gluing. In order to define the perverse category $P(X, \overline{\mathbb{Q}}_\ell)$, we need to pass the definition above along successive refinement. It is clear what it means for one stratification to refine another. We say a pair $(\{T_k\}, \mathcal{K})$ refines $(\{S_i\}, \mathcal{J})$ if $\{T_k\}$ refines $\{S_i\}$ and all irreducible sheaves in $\mathcal{J}(S_i)$ are $(\{T_k\}, \mathcal{K})$-constructible, i.e. objects in the category $\mathrm{Sh}_{\{T_k\},\mathcal{K}}(S_i, \overline{\mathbb{Q}}_\ell)$.

Lemma 2.4.1. *If $(\{T_k\}, \mathcal{K})$ refines $(\{S_i\}, \mathcal{J})$, then the perverse t-structure relative to $(\{S_i\}, \mathcal{J})$ is induced by that relative to $(\{T_k\}, \mathcal{K})$. That is, a $(\{S_i\}, \mathcal{J})$-constructible complex K is in ${}^pD^{b,\leq 0}_{\{S_i\},\mathcal{J}}(X, \overline{\mathbb{Q}}_\ell)$ (or $^{\geq 0}$) if and only if K is in ${}^pD^{b,\leq 0}_{\{T_k\},\mathcal{K}}(X, \overline{\mathbb{Q}}_\ell)$ (or $^{\geq 0}$).*

This depends on the finiteness theorems of [5]. It is thus apparent that we can take a colimit of the filtered system of refinements of stratifications $\{S_i\}$ and finite functions $\mathcal{J}$ to finally form the perverse t-structure on $D^b_c(X, \overline{\mathbb{Q}}_\ell)$. Doing so, we have the

Proposition 2.4.1. *A complex $K \in D^b_c(X, \overline{\mathbb{Q}}_\ell)$ is perverse if and only if, for all points (closed or not) $x \in X$, with $i_x : \{x\}^- \to X$, we have*

$$ H^i Ri_x^* K = 0 \text{ for } i > -\dim\{x\}^-, \text{ and} $$
$$ H^i Ri_x^! K = 0 \text{ for } i < -\dim\{x\}^-. $$

We also have that this category is, as defined, self-dual. That is, taking the Verdier dual

$$\mathbb{D}K := R\mathcal{H}om(K, \omega_X),$$

we have $K \in {}^pD_c^{b,\geq n}(X, \overline{\mathbb{Q}}_\ell)$ if and only if $\mathbb{D}K \in {}^pD_c^{b,\leq n}(X, \overline{\mathbb{Q}}_\ell)$. Note also that for a morphism f of schemes, $\mathbb{D}$ interchanges ! and $*$. That is, $\mathbb{D}Rf_*K = Rf_!\mathbb{D}K$, etc.

To amplify this idea, we have the

Theorem 2.4.2. *Let X be a finite-type scheme over an algebraically closed field. A complex $K \in D_c^b(X, \overline{\mathbb{Q}}_\ell)$ is in ${}^pD_c^{b,\leq 0}(X, \overline{\mathbb{Q}}_\ell)$ if and only if for all affine schemes U with open immersion $u : U \to X$, we have global cohomology*

$$H^i(U, u^*K) = 0$$

for all $i > 0$.

This statement is a converse to the theorem of [7] that if X is affine over an algebraically closed field and $\mathcal{F}$ is an étale sheaf over X, then $H^i(X, \mathcal{F}) = 0$ for all $i > \dim X$.

2.5 The middle extension

Under the gluing formalism, the functors from adjunctions give us natural maps

$$Rj_!Rj^*K \to K \to Rj_*Rj^*.$$

For a complex $K \in D_c^b(U, \overline{\mathbb{Q}}_\ell)$, with $j : U \to X$ an open immersion, we can view K as $Rj^*Rj_!K = Rj^*Rj_*K$. So we have a naturally defined composite

$$Rj_!K \to Rj_*K.$$

If $\mathcal{F}$ is a perverse sheaf on U, then we have a natural map

$${}^pj_!\mathcal{F} \to {}^pj_*\mathcal{F}$$

in $P(X, \overline{\mathbb{Q}}_\ell)$. We have the

Definition. The *middle extension*, or *intermediate extension*, $j_{!*}\mathcal{F}$, is the image of ${}^pj_!\mathcal{F}$ in ${}^pj_*\mathcal{F}$. Note that it is a quotient of ${}^pj_!\mathcal{F}$ and subobject of ${}^pj_*$, and all three extensions of $\mathcal{F}$ in $P(X, \overline{\mathbb{Q}}_\ell)$ are perverse sheaves.

Let X be a d-dimensional variety with a smooth stratification by dimensions $\{S_i\}_{i=0}^d$. Then for $-d \leq k \leq 0$, let $Z_k = \bigcup_{k \leq i \leq 0} S_{-i}$, and $U_k = \bigcup_{-d \leq i \leq k} S_{-i}$. Let also $j_{k-1} : U_{k-1} \hookrightarrow U_k$. Note that $Z_{-d} = U_0 = X$. Then we have the following inductive formula for middle extensions of smooth dense open immersions:

Proposition 2.5.1. *Let $j : U_k \hookrightarrow X$, and let $\mathcal{F}$ be a perverse sheaf over U_k. Then*

$$j_{!*}\mathcal{F} = \tau_{\leq -1}(j_{-1})_*\tau_{\leq -2}(j_{-2})_* \cdots \tau_{\leq k}(j_k)_*\mathcal{F}.$$

This gives us the formula for intersection complexes in the introduction.

2.6 Examples

First I want to remark that there are two classes of very simple examples, which should form the basis of visualizing perverse sheaves.

Example 2.6.1. If $\mathcal{F}$ is a lisse ℓ-adic sheaf on a smooth finite type separated scheme X and $\dim X = n$, then $\mathcal{F}[n]$ is a perverse sheaf on X. Similarly, $\mathbb{D}(\mathcal{F}[n]) = \mathcal{F}^\vee[n]$ is a perverse sheaf. This includes constant sheaves and, over a finite field, the Tate twisting sheaf $\overline{\mathbb{Q}}_\ell(1)$.

Example 2.6.2. If $x \in |X|$ is a closed point and $\mathcal{F}$ is a skyscraper sheaf with support on x (or, equivalently, $\mathcal{F} = (i_x)_*M$ where M is a Galois module of $\kappa(x)$) then $\mathcal{F}$ is a perverse sheaf on X. Similarly, $\mathbb{D}\mathcal{F} = (i_x)_*M^\vee$ is perverse.

Example 2.6.3. If X is a n-dimensional variety with an isolated singularity at a closed point $x \in X$, and $\overline{\mathbb{Q}}_\ell$ is the constant sheaf on $X \setminus \{x\}$, then

$$j_{!*}\overline{\mathbb{Q}}_\ell[n] = (\tau_{\leq n-1}Rj_*\overline{\mathbb{Q}}_\ell)[n],$$

and the stalk of $R^i j_*\overline{\mathbb{Q}}_\ell$ on a geometric point $\bar{x}$ over x is the colimit of the ith cohomology taken over étale $U \to X$ lying over x,

$$(R^i j_*\overline{\mathbb{Q}}_\ell)_{\bar{x}} = \varinjlim_{U} H^i(U \times_X (X \setminus \{x\}), \overline{\mathbb{Q}}_\ell)$$

My étale cohomology is a little shaky, but I believe that is equal to $H^i(\operatorname{Spec}(\mathcal{O}^{sh}_{X,\bar{x}})\setminus\{\mathfrak{m}^{sh}_{\bar{x}}\}, \overline{\mathbb{Q}}_\ell)$. Anyway it's clear that this is typically not concentrated in a single degree.

3 Classification

3.1 Classification over an arbitrary field

The following powerful result is basically dependent on properties of the category of constructible sheaves, but the same doesn't hold in that category. One difference is that middle extensions of perverse sheaves are perverse—it's not the case that the middle extension, or any other (derived) extension, of a constructible sheaf is a constructible sheaf, though we do still have that higher direct images (resp. with compact support) of constructible sheaves are constructible.

Definition. Given a closed integral subvariety $i : Z \hookrightarrow X$ of dimension d and a lisse sheaf (ℓ-adic local system) $\mathcal{L} \in \mathrm{Sh}_c(Z^\circ, \overline{\mathbb{Q}}_\ell)$ on a smooth (dense) open subset $Z^\circ \subset Z$, let $j : Z^\circ \hookrightarrow X$ be a locally closed immersion. Then the intersection complex $IC_Z(\mathcal{L})$ is the complex ${}^p i_* j_{!*}(\mathcal{L}[d])$. It is a perverse sheaf on X.

Theorem 3.1.1 (Classification)**.** *The category $P(X, \overline{\mathbb{Q}}_\ell)$ is a finite-length category with simple objects $IC_Z(\mathcal{L})$, where $\mathcal{L}$ is an irreducible lisse sheaf on $Z^\circ \subset Z$.*

3.2 Working over a finite field: Frobenii and Weil sheaves

Here we make the transition to considering only varieties over a finite field $\mathbb{F}_q$ or its algebraic (=separable) closure $\overline{\mathbb{F}}_q$. A variety over a finite field is denoted with a subscript 0 on the left, for example ${}_0X$, and the base change ${}_0X_{\overline{\mathbb{F}}_q}$ is denoted simply X. On the other hand, if X is a variety over $\overline{\mathbb{F}}_p$ defined over $\mathbb{F}_q$, we may fix a variety ${}_0X$ such that ${}_0X_{\overline{\mathbb{F}}_q} = X$. Similarly, for a sheaf ${}_0\mathcal{F}$ or complex of sheaves ${}_0K$ over ${}_0X$, its pullback to X will be denoted $\mathcal{F}$ or K respectively.

This is the proper framework for discussing pure and mixed perverse sheaves, as well as Weil sheaves.

In order to carefully define the notion of a mixed sheaf, we will need a couple of Frobeniuses. However, we will soon be working with them on a formal enough level that you can ignore all but one or two of them. (I may not be perfectly careful, to be honest.)

Let F^{-1} be the (arithmetic) frobenius of schemes over $\mathbb{F}_q$ defined as identity on spaces and $x \mapsto x^q$ on structrure sheaves. Then the geometric Frobenius F is the inverse of F^{-1}, and generates the Weil group $W(\overline{\mathbb{F}}_q/\mathbb{F}_q) \subset \mathrm{Gal}(\overline{\mathbb{F}}_q/\mathbb{F}_q)$. If X is a scheme over $\overline{\mathbb{F}}_q$, then F acts as an automorphism of X as a scheme over $\mathbb{F}_q$, and as a finite endomorphism (called ${}_0Fr_X$) of schemes over $\overline{\mathbb{F}}_q$. Notably, for a $\overline{\mathbb{F}}_q$-point $\bar{x} \in X(\overline{\mathbb{F}}_q)$, $\bar{x}$ lies over a point in $X(\mathbb{F}_{q^n})$ if and only if $\bar{x}$ is fixed by F^n (which is the geometric Frobenius for $\mathbb{F}_{q^n}$). That is to say

$$X(\overline{\mathbb{F}}_q)^{F^n} = X(\mathbb{F}_{q^n}).$$

So the Frobenius is quite useful in counting points of X over finite fields. Furthermore, given an ℓ-adic sheaf ${}_0\mathcal{G}$ over ${}_0X$ and a closed point $x \in |{}_0X|$, $W(\overline{\mathbb{F}}_q/\mathbb{F}_q)$ acts linearly on the stalk ${}_0\mathcal{G}_{\bar{x}}$. First, if necessary, base change to a field $\mathbb{F}_{q^n}$ on which x is defined (for instance the residue field of x). Let ${}_1X$ and ${}_1\mathcal{G}$ be the base-changed variety and sheaf. Then ${}_1\mathcal{G}$ is represented by an étale algebraic space ${}_1\mathfrak{G} \to {}_1X$, and ${}_1Fr_{\mathfrak{G}}$ induces on the fiber a map

$$F_x := ({}_1Fr_{\mathfrak{G}})_{\bar{x}} : {}_1\mathcal{G}_{\bar{x}} \to {}_1\mathcal{G}_{\bar{x}}.$$

The eigenvalues of the stalk Frobenius F_x are the focus of the Riemann hypothesis and the definition of purity (and mixedness, etc.), so it is the most important Frobenius in consideration. However, I would like to make one more observation and definition before we can be done. Given a sheaf ${}_0\mathcal{G}$ over ${}_0X$, there is a canonical isomorphism ${}_0F^* : {}_0Fr_X^*(\mathcal{G}) \xrightarrow{\sim} \mathcal{G}$ arising from the algebraic space $\mathfrak{G} \to X$. This datum distinguishes sheaves over X arising from sheaves over ${}_0X$. In particular, we have the

Proposition 3.2.1. *The functor*

$$\begin{aligned} \mathrm{Sh}_c({}_0X) \to \{&\text{pairs } (\mathcal{G}, \phi) \text{ of sheaves } \mathcal{G} \in \mathrm{Sh}_c(X) \\ &\text{and isomorphisms } \phi : {}_0Fr_X^*(\mathcal{G}) \xrightarrow{\sim} \mathcal{G}\} \\ {}_0\mathcal{G} \mapsto &(\mathcal{G}, {}_0F^*) \end{aligned}$$

is fully faithful exact, with essential image closed under subquotients and extensions.

For this reason, some authors prefer to work with the category on the right hand side, called Weil sheaves. In this sense, the base change functor can be viewed as a forgetful functor, for which we have the

Lemma 3.2.1. *The base change functor* $\mathrm{Sh}_c({}_0X, \overline{\mathbb{Q}}_\ell) \to \mathrm{Sh}_c(X, \overline{\mathbb{Q}}_\ell)$ *kills* Ext^1.

3.3 Purity and mixed sheaves

Given the notion of Frobenius action on stalks of a sheaf, we can define the mixed category and pure perverse sheaves. In order to make sense of some of the definitions, it is worth stating at the outset the

Theorem 3.3.1 (Main theorem of [6]). *Let ${}_0\mathcal{F}$ be a pure constructible sheaf of weight w on a variety ${}_0X$, and let $f : {}_0X \to {}_0Y$ be a morphism of schemes. Then for all i in $\mathbb{Z}$, the sheaf $R^i f_!({}_0\mathcal{F})$ is mixed, with weights $\leq w + i$.*

In particular, this explains why we tend to prefer working with lower bounds on weights and why weights on the cohomology groups of a complex vary with i. Note that by Verdier duality, we have in the situation above that Rf_* and $Rf^!$ preserve upper bounds on weights, and Rf^*, like $Rf_!$, preserves lower bounds.

So let us first define what it means for a constructible sheaf to be pointwise pure.

Definition. A sheaf ${}_0\mathcal{F}$ is pointwise pure of weight $w \in \mathbb{Z}$ if, for all closed points $x \in |{}_0X|$, supposing x is defined over $\mathbb{F}_{q^n}$, the eigenvalues of the stalk Frobenius F_x are all pure algebraic numbers of absolute value $q^{wn/2}$.

This absolute value can be taken by isomorphism $\iota : \overline{\mathbb{Q}}_\ell \xrightarrow{\sim} \mathbb{C}$. The meaning of "pure" above refers to the requirement that all such isomorphisms—i.e. all automorphisms of $\overline{\mathbb{Q}}$—preserve this absolute value. Hence it is a very strong condition, met for example by $\zeta q^{wn/2}$, where ζ is a root of unity.

The correct categories to consider purity of perverse sheaves are the categories of mixed complexes, and thus of mixed perverse sheaves. We have

Definition. A constructible sheaf ${}_0\mathcal{F}$ is *mixed* if it is a finite successive extension of pointwise pure constructible sheaves of any weight. The set of integers for which irreducible subquotients of ${}_0\mathcal{F}$ are pure of weight w is called *the weights of* ${}_0\mathcal{F}$. If that set is bounded above by an integer w, ${}_0\mathcal{F}$ is called *mixed of weights* $\leq w$. Their respective categories are denoted $\mathrm{Sh}_m({}_0X, \overline{\mathbb{Q}}_\ell)$ and $\mathrm{Sh}_{\leq w}({}_0X, \overline{\mathbb{Q}}_\ell)$.

Definition. A complex ${}_0K$ of constructible sheaves is called *mixed* if every cohomology sheaf is a mixed constructible sheaf. ${}_0K$ is called a *mixed complex of weights* $\leq w$ if for all i the ith cohomology sheaf $H^i{}_0K$ is a mixed sheaf of weights $\leq w+i$. Their respective categories are denoted $D^b_m({}_0X, \overline{\mathbb{Q}}_\ell)$ and $D^b_{\leq w}({}_0X, \overline{\mathbb{Q}}_\ell)$. A complex is *mixed of weights* $\geq w$ if its Verdier dual $\mathbb{D}_0 K$ is mixed of weights $\leq -w$.

Definition. A complex ${}_0K$ is called *pure of weight* w if ${}_0K$ is mixed of weights $\leq w$ and $\geq w$. The same applies to perverse sheaves. Their respective categories are denoted $D^b_w({}_0X, \overline{\mathbb{Q}}_\ell)$ and $P_w({}_0X, \overline{\mathbb{Q}}_\ell)$.

Proposition 3.3.1. *If a mixed complex has lisse cohomology sheaves $H^i{}_0K$, then it is mixed of weights $\leq w$ (respectively $\geq w$) if and only if for all closed points $x \in |{}_0X|$ the eigenvalues of the stalk Frobenius F_x on $(H^i{}_0K)_x$ are all pure algebraic numbers of absolute value $\leq q^{(w+i)n/2}$ (respectively $\geq q^{(w+i)n/2}$). Thus a mixed constructible sheaf is pure of weight w if and only if it is pointwise pure of weight w, and a mixed complex is pure of weight w if and only if its cohomology sheaves $H^i{}_0K$ are pointwise pure of weight $w+i$.*

Although the following theorem is deeper into the theory of perverse sheaves than other theorems around it, it is essentially a strengthening of above.

Proposition 3.3.2. *A mixed complex ${}_0K$ has weights $\leq w$ (resp. $\geq w$) if and only if all of its perverse cohomology sheaves ${}^pH^i({}_0K)$ are mixed of weights $\leq w+i$ (resp. $\geq w+i$).*

Lemma 3.3.1. *Given mixed perverse sheaves ${}_0\mathcal{F}$ of weights $\leq w$ and ${}_0\mathcal{G}$ of weights $\geq w'$, with $w' > w$, we have*

$$\mathrm{Ext}^1({}_0\mathcal{F}, {}_0\mathcal{G}) = 0.$$

Lemma 3.3.2. *The categories $P_{\leq w}({}_0X, \overline{\mathbb{Q}}_\ell)$ and $P_{\geq w}({}_0X, \overline{\mathbb{Q}}_\ell)$ are closed under subquotients in $P({}_0X, \overline{\mathbb{Q}}_\ell)$.*

Theorem 3.3.2 (Weight filtration)**.** *Given a mixed perverse sheaf ${}_0\mathcal{F}$ over ${}_0X$, there is an increasing filtration $\mathcal{W}$, called the* weight filtration, *such that $\mathrm{gr}^{\mathcal{W}}_i({}_0\mathcal{F})$ is pure of weight i.*

Theorem 3.3.3 (Classification theorem, mixed case). *The category $P_m({}_0X, \overline{\mathbb{Q}}_\ell)$ is a full subcategory of $P({}_0X, \overline{\mathbb{Q}}_\ell)$ closed under subquotients and extensions. Its simple objects are of the form $IC_{{}_0Z}({}_0\mathcal{L})$, where ${}_0\mathcal{L}$ is an irreducible* pure *lisse sheaf.*

Again, the following is a bit deeper than the preceding, but it fits here thematically.

Proposition 3.3.3. *The indecomposable objects of $P_m({}_0X, \overline{\mathbb{Q}}_\ell)$ are of the form $\mathcal{S}_0 \otimes \mathcal{J}_n$, where $\mathcal{S}_0$ is a simple pure perverse sheaf (i.e. an intersection complex defined by a pure irreducible lisse sheaf), and $\mathcal{J}_n$ is the n-dimensional unipotent sheaf where the Frobenius acts on all stalks via a Jordan block with eigenvalue 1.*

Proposition 3.3.4. *Let ${}_0\mathcal{F}$ be a pure perverse sheaf of weight w over ${}_0X$. Then for all open subschemes $j : {}_0U \hookrightarrow {}_0X$ with closed complement $i : {}_0Z \hookrightarrow {}_0X$, we have the unique decomposition*

$$ {}_0\mathcal{F} \cong j_{!*}({}_0\mathcal{F}') \oplus i_*({}_0\mathcal{F}'') $$

for some pure perverse sheaves ${}_0\mathcal{F}' \in P_w({}_0U, \overline{\mathbb{Q}}_\ell)$ and ${}_0\mathcal{F}'' \in P_w({}_0Z, \overline{\mathbb{Q}}_\ell)$.

Applying noetherian induction, we have

$$ {}_0\mathcal{F} \cong \bigoplus_{{}_0Z \hookrightarrow {}_0X} IC_{{}_0Z}({}_0\mathcal{L}), $$

where the sum runs over finitely many closed integral subschemes with possible multiplicity, and ${}_0\mathcal{L}$ is a pure indecomposable lisse sheaf of weight $(w - \dim {}_0Z)$ on an open dense ${}_0Z^\circ \subset {}_0Z$.

4 Decomposition and Relative Hard Lefschetz

For some background, we have two powerful and impressive theorems of Deligne. The first is something I've seen called Deligne's decomposition theorem.

Theorem 4.0.1 (Decomposition Theorem for smooth varieties). *Given a proper smooth map $f : {}_0X \to {}_0Y$ between nonsingular varieties over a finite field and a pure constructible sheaf ${}_0\mathcal{F} \in \mathrm{Sh}_w({}_0X, \overline{\mathbb{Q}}_\ell)$, the direct image complex $Rf_*({}_0\mathcal{F})$ is pure of weight w and decomposes as follows in $D_c^b(X, \overline{\mathbb{Q}}_\ell)$:*

$$ Rf_*(\mathcal{F}) \cong \bigoplus_{i \geq 0} R^i f_*(\mathcal{F})[-i]. $$

Theorem 4.0.2 (Relative Hard Lefschetz Theorem for smooth varieties). *Given a projective smooth map $f : {}_0X \to {}_0Y$ of relative dimension d between nonsingular varieties, letting $\eta \in H^2(X, \overline{\mathbb{Q}}_\ell)$ be the first Chern class of a relatively ample line bundle on X, the iterated cup product with η produces isomorphisms*

$$ (- \cup \eta)^i : R^{d-i} f_* \overline{\mathbb{Q}}_\ell \xrightarrow{\sim} R^{d+i} f_* \overline{\mathbb{Q}}_\ell $$

of pure constructible sheaves on X.

Note in both cases the base change to the algebraic closures.

Significantly stronger versions of both theorems are known using the formalism of perverse sheaves. We have the restatement

Theorem 4.0.3 (Weil II, stated more maturely). *Let ${}_0K$ be a mixed complex of weights $\leq w$ over ${}_0X$, and let $f : {}_0X \to {}_0Y$ be a separated morphism of schemes. Then $Rf_!({}_0K)$ is a mixed complex of weights $\leq w$ over ${}_0Y$. By duality, we have Rf_* and $Rf^!$ preserve (mixedness and) lower bounds, and Rf^* also preserves upper bounds.*

Lemma 4.0.1 (Gabber purity). *If $f : {}_0X \to {}_0Y$ is a morphism of schemes and ${}_0\mathcal{F}$ is a mixed complex over ${}_0X$, then ${}^pf_!$ preserves lower bounds on its weights, and ${}^pf_*$ preserves upper bounds. Similarly, ${}^pf^*$ and ${}^pf^!$ preserve lower and upper bounds, respectively, of a sheaf ${}_0\mathcal{G}$ over ${}_0Y$.*

Example 4.0.1. In the case f is quasifinite, we can define a functor $f_{!*}$ as the image of ${}^pf_!$ in ${}^pf_*$, and it preserves perversity and purity.

If f is smooth, then $Rf^*[d]$ preserves perversity, and sends complexes of weights $\leq w$ to complexes of weights $\leq (w-d)$. Then if f is étale, we have $f^* = Rf^* = Rf^!$, and it preserves perversity and purity, as does $f_{!*}$.

On the other hand, if f is proper, then we have $f_* = f_!$, and so ${}^pf_*$ and even Rf_* preserve purity. If, furthermore, f is finite, then $f_* = Rf_* = f_{!*}$, and f_* preserves perversity and purity.

Lemma 4.0.2. *If K is a pure complex over X derived via base change from a pure complex ${}_0K$ over ${}_0X$, then we have the decomposition*

$$K \cong \bigoplus_{i\in\mathbb{Z}} {}^pH^iK[-i].$$

And if $\mathcal{F}$ is a pure perverse sheaf derived from ${}_0\mathcal{F}$, then we have the decomposition

$$\mathcal{F} \cong \bigoplus_{Z\hookrightarrow X} IC_Z(\mathcal{L}),$$

where the sum runs over closed integral subschemes of X, with possible multiplicity, and all $\mathcal{L}$ appearing in the sum are irreducible pure lisse sheaves of weight $(w - \dim Z)$ on some dense open $Z^\circ \subset Z$.

Combining 4.0.1 and 4.0.2, we have the

Theorem 4.0.4 (Decomposition Theorem for perverse sheaves). *Let $f : {}_0X \to {}_0Y$ be a proper morphism of schemes. Let ${}_0\mathcal{F}$ be a pure perverse sheaf of weight w over ${}_0X$. Then we have the decomposition*

$$Rf_*\mathcal{F} \cong \bigoplus_{i\geq 0} {}^pH^i(Rf_*\mathcal{F})[-i]$$

and for all i we have the further decomposition

$${}^pH^i(Rf_*\mathcal{F}) \cong \bigoplus_{Z\hookrightarrow X} IC_Z(\mathcal{L}),$$

with all $\mathcal{L}$ lisse, irreducible, and pure of weight $(w + i - \dim Z)$.

Theorem 4.0.5 (Relative Hard Lefschetz Theorem for perverse sheaves). *Let $f : {}_0X \to {}_0Y$ be a projective morphism of schemes, and let ${}_0\mathcal{F}$ be a perverse sheaf of weight w over ${}_0X$. Let η be the first Chern class of an f-ample line bundle in $H^2({}_0X, \overline{\mathbb{Q}}_\ell(1))$. Then the iterated cup product with η produces isomorphisms for all i*

$$(-\cup\eta)^i : {}^pH^{-i}Rf_*({}_0\mathcal{F}) \xrightarrow{\sim} {}^pH^iRf_*({}_0\mathcal{F})(i),$$

where (i) on the right hand side represents the i-fold tate twist $- \otimes \overline{\mathbb{Q}}_\ell(1)^{\otimes i}$.

5 Geometric Satake Correspondence

One amazing consequence of the theory of perverse sheaves is the geometric Satake correspondence of [13]. I will describe an approach due to [14]. This will require extending the kinds of spaces that perverse sheaves can live on, defining some notion of equivariant perverse sheaves, and giving the category of L^+G-equivariant perverse sheaves on $\mathcal{G}r_G$ a convolution product. Then we have a notion of the left hand side, and the statement is an isomorphism.

I want to state the necessary Schubert structure on $\mathcal{G}r_G$ with its L^+G-action, the two basic equivalent ideas on what an equivariant perverse sheaf should be, the multiplication on $P_{L^+G}(\mathcal{G}r_G)$, and a monoidal functor ω (defined simply as a direct sum of hypercohomology groups!) to $\mathrm{Vect}_{\overline{\mathbb{Q}}_\ell}$. There is a $\hat{G}$-action on fibers of the functor ω, respecting $\otimes$, producing an isomorphism of symmetric monoidal categories

$$\omega : (P_{L^+G}(\mathcal{G}r_G), \star) \xrightarrow{\sim} (\mathrm{Rep}_{\overline{\mathbb{Q}}_\ell}(\hat{G}), \otimes) \ .$$

So first recall a rather fundamental theorem that if $i : Z \hookrightarrow X$ is a closed immersion, then the functor ${}^p i_* = i_* : P(Z, \overline{\mathbb{Q}}_\ell) \to P(X, \overline{\mathbb{Q}}_\ell)$ is fully faithful, with essential image closed under subquotients and extensions.

Now recall that if G is a reductive group over a field k of positive characteristic (we may take $k = \overline{\mathbb{F}}_p$), we have $\mathcal{G}r_G := LG/L^+G$ is an ind-(projective scheme) over k, and thus ind-finite type. So given an ind-presentation (X_i) of $\mathcal{G}r_G$ such that each X_i is of finite type, we have well-defined $P(X_i, \overline{\mathbb{Q}}_\ell)$ and the maps $X_i \to X_j$ preserve that perverse structure such that we may take

$$P(\mathcal{G}r_G, \overline{\mathbb{Q}}_\ell) := \varinjlim_i P(X_i, \overline{\mathbb{Q}}_\ell).$$

It is a theorem of [11] that this definition of the perverse category does not depend on the ind-presentation.

As it happens, we will have a favorite presentation, given by L^+G-orbits. In fact, we have a kind of Schubert decomposition of $\mathcal{G}r_G$ into Schubert cells $\mathcal{O}_\mu := L^+G \cdot t^\mu$ for dominant cocharacters μ of G. Furthermore, these orbits actually produce an ind-presentation—their closures are actually projective schemes, and unions of cells indexed by smaller cocharacters. Thus we have $\mathcal{G}r_G = \varinjlim \overline{\mathcal{O}}$.

More usefully yet, the L^+G-action on each orbit actually decomposes through a quotient of L^+G, the ith jet group G_i, which is a finite-type affine group scheme over k. And so we can define on each orbit (bearing in mind the invariance of the étale site under homeomorphism) a category of L^+G-invariant perverse sheaves.

Roughly speaking, given a finite type reduced affine group H and a finite type scheme X with an H-action $m : H \times X \to X$, the category of H-invariant perverse sheaves on X is the category of perverse sheaves $\mathcal{F}$ with some isomorphism $\theta : m^*\mathcal{F} \xrightarrow{\sim} \mathrm{pr}^*\mathcal{F}$ over $H \times X$. This category is known as $P_H(X, \overline{\mathbb{Q}}_\ell)$. More precisely, it is shown in [11] that $P_H(X, \overline{\mathbb{Q}}_\ell)$ is naturally equivalent to the full category of perverse sheaves $P(H\backslash X, \overline{\mathbb{Q}}_\ell)$, where $H\backslash X$ is the stack quotient.

And so, since we have the action of L^+G factors through a finite-type group G_i on $\mathcal{O}$, we have an inductive system of L^+G-equivariant perverse sheaves, allowing us to define $P_{L^+G}(\mathcal{G}r_G, \overline{\mathbb{Q}}_\ell)$.

Convolution is a bit tougher. It helps to work with the Beauville–Laszlo (cf [1]) conception of an affine grassmannian satisfying a certain moduli problem. I'm really not in a place where I can define the convolution product.

Then cohomology, just the direct sum of global cohomology groups, is a symmetric monoidal functor with respect to $\star$.

To find the isomorphism with representations of $\hat{G}$, use the Tannakian formalism and consider the group of automorphisms of the functor ω. Using a result of [9], we can reconstruct the root datum of the group, and we find it is precisely dual to the root datum of G, implying the space on the right hand side is indeed the representations of $\hat{G}$. This can also be viewed as a non-combinatorial construction of the dual group.

References

[1] Arnaud Beauville and Yves Laszlo. Un lemme de descente. *Comptes Rendus de l'Academie des Sciences-Serie I-Mathematique*, 320(3):335–340, 1995.

[2] Alexander Beilinson, Joseph Bernstein, Pierre Deligne, and Ofer Gabber. *Faisceaux pervers.* Société Mathématique de France, 2018.

[3] Bhargav Bhatt and Peter Scholze. The pro-étale topology for schemes. *arXiv preprint arXiv:1309.1198*, 2013.

[4] Mark de Cataldo and Luca Migliorini. The decomposition theorem, perverse sheaves and the topology of algebraic maps. *Bulletin of the American Mathematical Society*, 46(4):535–633, 2009.

[5] Pierre Deligne. SGA 4 1/2–Cohomologie étale. *Lecture Notes in Mathematics*, 569, 1977.

[6] Pierre Deligne. La conjecture de Weil II. *Publications Mathématiques de l'Institut des Hautes Etudes Scientifiques*, 52(1):137–252, 1980.

[7] Pierre Deligne, Michael Artin, Bernard Saint-Donat, Alexander Grothendieck, and Jean-Louis Verdier. *Théorie des Topos et Cohomologie Etale des Schémas. Séminaire de Géométrie Algébrique du Bois-Marie 1963-1964 (SGA 4)*, volume 3. Springer, 2006.

[8] Eberhard Freitag and Reinhardt Kiehl. *Etale cohomology and the Weil conjecture*, volume 13. Springer, 2013.

[9] David Kazhdan, Michael Larsen, and Yakov Varshavsky. The Tannakian formalism and the Langlands conjectures. *Algebra & Number Theory*, 8(1):243–256, 2014.

[10] Reinhardt Kiehl and Rainer Weissauer. *Weil conjectures, perverse sheaves and l'adic Fourier transform*, volume 42. Springer Science & Business Media, 2013.

[11] Yves Laszlo and Martin Olsson. Perverse t-structure on Artin stacks. *Mathematische Zeitschrift*, 261(4):737–748, 2009.

[12] James S Milne. Lectures on étale cohomology. *Available on-line at http://www. jmilne. org/math/CourseNotes/LEC. pdf*, 1998.

[13] Ivan Mirković and Kari Vilonen. Geometric Langlands duality and representations of algebraic groups over commutative rings. *Annals of Mathematics*, pages 95–143, 2007.

[14] Timo Richarz. A new approach to the geometric Satake equivalence. *arXiv preprint arXiv:1207.5314*, 2012.