12.

We introduce the notion of categorification. Then we introduce constructible sheaves, and discuss some naive attempts to categorify the Iwahori–Hecke algebra using them.

Throughout this course, the main reference on sheaves and their technical details will be Achar's book, supplemented by SGA and the texts by Freitag–Kiehl, Kiehl–Weissauer, and Milne. As for categorification of the Hecke algebra, a possible reference is Lecture 24 in Romanov–Williamson's lecture notes.

12.1.

Categorification of an additive group A means constructing an additive category C such that A is the Grothendieck group of C in an appropriate sense.

There are several different kinds of additive category, each with its own notion of Grothendieck group. In each case, we assume that C admits a small skeleton; the Grothendieck group is generated by the isomorphism classes of objects in the skeleton modulo certain relations.

(1) For any C, the *split Grothendieck group* $[\mathsf{C}]_\oplus$ is given by the relations

$$[c] = [c'] + [c''] \quad \text{for any } c \simeq c' \oplus c''.$$

(2) For C abelian, the usual *Grothendieck group* $[\mathsf{C}]$ is given by the relations

$$[c] = [c'] + [c''] \quad \text{for any exact sequence } 0 \to c' \to c \to c'' \to 0.$$

(3) For C triangulated, the *triangulated Grothendieck group* $[\mathsf{C}]_\triangle$ is given by the relations

$$[c] = [c'] + [c''] \quad \text{for any exact triangle } c' \to c \to c'' \to c'[1].$$

Note that for any c, the triangle $c \to 0 \to c[1] \to c[1]$ is exact, giving $[c[1]] = -[c]$. That is, the shift $[1]$ must decategorify to scaling by -1.

It appears to be well-known that if C is abelian and $\mathsf{D}^b(\mathsf{C})$ is the bounded derived category of complexes of objects in C, then $[\mathsf{D}^b(\mathsf{C})]_\triangle = [\mathsf{C}]$. Seemingly less-known, but important for our goals, is a result recorded by David Rose in "A Note on the Grothendieck Group..." Below, for any additive C, let $\mathsf{K}^b(\mathsf{C})$ be the bounded homotopy category of complexes of objects in C.

Theorem 12.1 (Rose). *We have* $[\mathsf{K}^b(\mathsf{C})]_\triangle = [\mathsf{C}]_\oplus$.

Remark 12.2. For C abelian, $[\mathsf{C}]_\oplus$ is usually larger than $[\mathsf{C}]$. This corresponds to the fact that a short exact sequence of complexes in C will give rise to an exact triangle in $\mathsf{D}^b(\mathsf{C})$ but not necessarily in $\mathsf{K}^b(\mathsf{C})$.[1]

---

[1] Thank-you to David B. for spotting an error here during the lecture.

Categorification of a ring R begins with categorification of the underlying additive group to some category C. We then need to construct some monoidal product $*$ on C that distributes over the direct sum $\oplus$, such that the relations

$$[c] = [c'][c''] \quad \text{for any } c \simeq c' * c''$$

define the multiplication on R.

Our goal is to build a categorification of the Iwahori–Hecke algebra $H_W(\mathsf{x})$ involving geometric objects. Recall that for suitable G, F, B, we have

$$H_W(\mathsf{x})|_{\mathsf{x} \to q^{1/2}} \simeq \operatorname{End}_{AG^F}(\operatorname{Ind}_{AB^F}^{AG^F}(1)), \quad \text{where } A = \mathbf{Z}[q^{\pm 1/2}],$$

and that the right-hand side can be rewritten in terms of G^F-invariant, A-valued functions on $(G/B \times G/B)^F$. It turns out that such functions arise from G-equivariant sheaves on $G/B \times G/B$ in a precise sense. Moreover, the sheaves involved give rise to a cohomology theory that recovers the notion of étale cohomology discussed earlier. However, getting the "right" sheaves to appear, with the "right" cohomology, is much trickier than one might expect.

12.2.

Fix an algebraically closed field k and a scheme X of finite type over k. Recall that étale cohomology behaves like singular cohomology with *locally constant* coefficients. So we might consider étale sheaves that are locally constant on X: *i.e.*, sheaves that trivialize after pullback along a finite étale cover of X. It turns out that this sometimes gives more sheaves than expected, and even when it doesn't, it gives the wrong cohomology when the sheaves have non-torsion coefficients.

To explain further: In analogy with classical topology, our sheaf theory should provide a (functorial) equivalence between locally constant sheaves on X and representations of some kind of fundamental group $\pi_1(X)$. In algebraic geometry, the usual choice is the étale fundamental group $\pi_1^{\text{ét}}(X)$, essentially because when we take $X = \operatorname{Spec} k_1$ for a field k_1, it recovers the absolute Galois group of k_1. But the étale fundamental group carries an intrinsic topology: the profinite topology. Finite sets with a $\pi_1^{\text{ét}}(X)$-action correspond to finite étale covers of X, so *continuous* finite-rank representations of $\pi_1^{\text{ét}}(X)$ correspond to étale-locally constant sheaves on X.

Example 12.3. In the case where X is a rational curve with a node, $\pi_1^{\text{ét}}(X)$ is procyclic, hence compact. At the same time, any étale-locally constant sheaf of rank 1 over X is determined by an *arbitrary* nonzero scalar describing the gluing at the node. If the ring of coefficients is not compact, then there are more possible scalars than representations of $\pi_1^{\text{ét}}(X)$.

Example 12.4. In the case where $X = \operatorname{Spec} k_1$ case, the issue is not with the sheaves, but with their cohomology. It turns out that the cohomology of a locally constant sheaf should be the group cohomology of the corresponding $\operatorname{Gal}(k/k_1)$-module. But one can show that for any profinite group Γ, written as an inverse limit of finite quotients $\Gamma = \varprojlim_i \Gamma/\Gamma_i$, and any continuous Γ-module M, we have an isomorphism of cochain groups

(12.1) $$C^*(\Gamma, M) \simeq \varinjlim_i C^*(\Gamma/\Gamma_i, M^{\Gamma_i}).$$

Since the Γ/Γ_i are finite, the right-hand side will fail to detect the non-torsion part of M. See SGA 4, Exposé IX or §12 of Freitag–Kiehl's book for related discussions of this issue.

Remark 12.5. Why doesn't Achar encounter these issues in Chapter 2 of his book? There, he isn't working in the étale topology at all: He is taking $k = \mathbf{C}$ and working in the *analytic* topology on the underlying topological spaces.

12.3.

There are two ways to fix the issues above.

12.3.1. The old way

One avoids defining a constructible sheaf as an étale-locally constant sheaf, except when special conditions hold on the coefficient ring as well as the sheaf itself. Namely, the ring must be finite and self-injective, like $\mathbf{Z}/\ell^m\mathbf{Z}$, and the sheaf Tor-finite: See Achar, pp. 221–222.

One instead defines the *constructible derived category* in an indirect way, using an inverse limit of the categories arising from the special cases. Then one constructs a t-structure on the result,[2] and defines constructible sheaves to be the objects of the heart of this t-structure. The details are laid out in Definition 5.1.14 of Achar's book.

12.3.2. The new way

Due to Bhatt–Scholze. One replaces the (small) étale site of X with the pro-étale site of X, whose objects are so-called weakly-étale schemes over X and whose covers are the fpqc covers. The étale site embeds into the pro-étale site with a right adjoint. In general, the latter is much larger, in that its typical objects involve non-noetherian constructions even when X is noetherian.

When X is the nodal curve, the pro-étale fundamental group is just $\mathbf{Z}$, fixing the issue seen earlier. In general, the pro-étale fundamental group merely has the étale one as its profinite completion.

---

[2]To be discussed next week.

When $X = \operatorname{Spec} k_1$, the groups are the same. Profinite $\operatorname{Gal}(k/k_1)$-sets correspond to objects of the pro-étale site, so continuous representations of $\operatorname{Gal}(k/k_1)$, not necessarily of finite rank, correspond to pro-étale-locally constant sheaves. But due to the existence of more covers, the pro-étale analogue of (12.1) does not hold. Instead, the pro-étale cohomology of locally constant sheaves behaves like the naïve expectation, even with arbitrary coefficients, as long as the characteristic of k is invertible in the coefficient ring.

The current best reference for the pro-étale topology is its chapter in the Stacks Project, which fixes many minor errors in the original text by Bhatt–Scholze.

12.3.3.

We will adopt the old way, going forward, because the enhancements that we actually need—mixedness and equivariance—only seem to be written down carefully in the old way at present. Indeed, providing an accurate summary of these enhancements takes up Chapters 5 and 6 of Achar's book.

Suppose that $\mathfrak{k}$ is a finite self-injective ring. An étale sheaf of $\mathfrak{k}$-modules $\mathcal{F}$ over X is *lisse* if and only if it is locally constant of finite type. More generally, if $\mathcal{S}$ is a finite stratification of X by constructible subschemes,[3] then $\mathcal{F}$ is *constructible with respect to $\mathcal{S}$* if and only if its restriction to each stratum is lisse. We say that it is *constructible* if and only if it is constructible with respect to some $\mathcal{S}$. Let $\mathsf{Shv}(X, \mathfrak{k})$ be the (abelian, $\mathfrak{k}$-linear) category of constructible sheaves of $\mathfrak{k}$-modules.

Now suppose that $\mathfrak{o}$ is the ring of integers in a finite extension of $\mathbf{Q}_\ell$, where ℓ is a prime different from the characteristic of k. Let $\mathfrak{m} \subseteq \mathfrak{o}$ be the maximal ideal. Let $\mathsf{Shv}(X, \mathfrak{o})$ be the abelian, $\mathfrak{o}$-linear category of *constructible $\mathfrak{o}$-sheaves* on X in the sense of Achar Definition 5.1.16. For all $j > 0$, there is a functor $\mathsf{Shv}(X, \mathfrak{o}) \to \mathsf{D}^{\leq 0}(\mathsf{Shv}(X, \mathfrak{o}/\mathfrak{m}^j))$, where $\mathsf{D}^{\leq 0}$ means the full subcategory of the derived category of objects with no cohomology in positive degrees. These functors more-or-less characterize $\mathsf{Shv}(X, \mathfrak{o})$: See Achar pp. 222–223.

Finally, let $\mathsf{Shv}(X, \bar{\mathbf{Q}}_\ell)$ be the abelian, $\bar{\mathbf{Q}}_\ell$-linear category of *constructible $\bar{\mathbf{Q}}_\ell$-sheaves* on X in the sense of Achar Definition 5.1.16. We will just call them *sheaves* when convenient. For any ring $\mathfrak{o}$ as above, extension of scalars induces a functor $\mathsf{Shv}(X, \mathfrak{o}) \to \mathsf{Shv}(X, \bar{\mathbf{Q}}_\ell)$. Again, these functors more-or-less characterize $\mathsf{Shv}(X, \bar{\mathbf{Q}}_\ell)$.

The facts we need about $\mathsf{Shv}(X, \mathbf{k})$ when $\mathbf{k} \in \{\mathfrak{k}, \mathfrak{o}, \bar{\mathbf{Q}}_\ell\}$:

(1) A map $p : Y \to X$ of schemes over k that sends strata onto strata induces a *pullback* $p^* : \mathsf{Shv}(X, \mathbf{k}) \to \mathsf{Shv}(Y, \mathbf{k})$. In particular, if $j : U \to X$ is

---

[3]Technically, a stratification is stricter than a *partition*. See Achar Definition 2.3.1.

an étale open, then for any object $\mathcal{F}$ of $\mathsf{Shv}(X, \mathbf{k})$, we set

$$\mathcal{F}(U) = \mathrm{Hom}(\mathbf{k}_U, j^*\mathcal{F})$$

When $\mathbf{k} = \mathfrak{k}$, this is precisely the space of sections of $\mathcal{F}$ over U in the usual sense.

(2) Set $pt = \mathrm{Spec}\, k$. Then all stratifications of pt are the same, and $\mathsf{Shv}(pt, \mathbf{k})$ is equivalent to the category of $\mathbf{k}$-modules of finite rank. For any object $\mathcal{F}$ of $\mathsf{Shv}(X, \mathbf{k})$ and any $\bar{x} \in X(k)$ viewed as a map $\bar{x} : pt \to X$, we define the *stalk* of $\mathcal{F}$ at $\bar{x}$ to be

$$\mathcal{F}_{\bar{x}} = \bar{x}^*\mathcal{F} = \varprojlim_{\substack{\text{étale open } U \to X \\ \bar{x} \in U(k)}} \mathcal{F}(U),$$

viewed as a $\mathbf{k}$-module.

(3) The usual tensor product induces a monoidal product $\otimes$ on $\mathsf{Shv}(X, \mathbf{k})$. The monoidal unit is the *constant sheaf*

$$\mathbf{k}_X = a^*\mathbf{k},$$

where $a : X \to pt$ is the unique map.

Moreover, all of these constructions generalize to the setting where we replace k with a subfield k_1, and X with a k_1-structure/form X_1:

(4) Pullback induces a functor $\mathsf{Shv}(X_1, \mathbf{k}) \to \mathsf{Shv}(X, \mathbf{k})$. In particular, if we set $pt_1 = \mathrm{Spec}\, k_1$, then $\mathsf{Shv}(pt_1, \mathbf{k})$ is equivalent to the category of $\mathbf{k}\mathrm{Gal}(k/k_1)$-modules of finite rank over $\mathbf{k}$, and $\mathsf{Shv}(pt_1, \mathbf{k}) \to \mathsf{Shv}(pt, \mathbf{k})$ is the forgetful functor.

In the rest of this lecture, we work only with $\mathbf{k} = \bar{\mathbf{Q}}_\ell$. Hence we will abbreviate by writing $\mathsf{Shv}(X) = \mathsf{Shv}(X, \bar{\mathbf{Q}}_\ell)$.

12.4.

If H is an algebraic group over k, acting on X with finitely many orbits, then it is natural to study sheaves constructible with respect to the stratification of X by H-orbits. In particular, if G is a reductive algebraic group over k and $\mathcal{B}$ is its flag variety, then we might take $X = \mathcal{B} \times \mathcal{B}$ and $H = G$. Alternatively, if $B \subseteq G$ is a fixed Borel, then we might take $X = G$ and $H = B \times B$.

How can we recover the Iwahori–Hecke algebra $H_W = H_W(\mathsf{x})$ from sheaves in this setup? Recall that when we identify $H_W(q^{1/2})$ with G^F-invariant functions on $\mathcal{B}^F \times \mathcal{B}^F$, the rescaled standard basis elements $q^{\ell(w)/2}\sigma_w$ correspond to the indicator functions on G^F-orbits.

The sheafy analogue of an indicator function is the extension-by-zero of a constant sheaf.[4] This leads to a naive guess: Since $\mathsf{Shv}(X)$ is abelian, we can

---

[4]Another notion to be discussed next week.

form either of the Grothendieck groups $[\mathsf{Shv}(X)]_\oplus$ or $[\mathsf{Shv}(X)]$. For $X = \mathcal{B} \times \mathcal{B}$, we might hope that the extensions-by-zero to X of the constant sheaves on its G-orbits decategorify to a rescaled standard basis of $H_W(\mathsf{x})$.

Unfortunately, this can't work. Take $G = \mathrm{GL}_1$ (or any other torus). Then $\mathcal{B} = pt$, so $X = pt$. By item (3) above, $[\mathsf{Shv}(pt)]_\oplus = [\mathsf{Shv}(pt)] = \mathbf{Z}$, whereas we want $H_W(\mathsf{x}) = \mathbf{Z}[\mathsf{x}^{\pm 1}]$.

12.5.

In trying to fix this issue, we might set $k = \bar{\mathbf{F}}_q$ and try to specialize x to a square root of q, like we did to compare the Hecke algebra to functions on $\mathcal{B}^F \times \mathcal{B}^F$. However, we immediately realize that k itself does not see q, but only the underlying prime of which q is a power. This suggests working with sheaves equipped with extra structure coming from an $\mathbf{F}_q$-structure on X, and using this structure to enrich our Grothendieck groups.

Henceforth, $k = \bar{\mathbf{F}}_q$, so that $\ell \nmid q$. Suppose that $X = X_1 \otimes k$ and $\mathcal{F} = \mathcal{F}_1|_X$ for some scheme X_1 over $\mathbf{F}_q$ and sheaf $\mathcal{F}_1$ over X_1. We will describe the *fonctions-faisceaux* (or *function-sheaf*) *dictionary*, which sends $\mathcal{F}$ to a collection of *trace of Frobenius* functions

$$\mathbf{t}_{\mathcal{F}} = \mathbf{t}_{\mathcal{F},d} : X_1(\mathbf{F}_{q^d}) \to \bar{\mathbf{Q}}_\ell \quad \text{for } d \geq 1.$$

Theorem 1.1.2 in Laumon's "Transformation de Fourier..." states a precise sense in which (under certain conditions) the functions $\mathbf{t}_{\mathcal{F},d}$ determine the class of $\mathcal{F}$ in a certain Grothendieck group.

Recall that given Z over $\mathbf{F}_q$, the absolute Frobenius map $\sigma_Z : Z \to Z$ is the map over $\mathbf{F}_q$ that fixes the underlying topological space and sends $f \mapsto f^q$ on sections of the structure sheaf $\mathcal{O}_Z$. We also worked with the relative Frobenius maps $F = \sigma_{X_1} \otimes \mathrm{id} : X \to X$, which are maps over k. It turns out that if $\mathcal{F} = \mathcal{F}_1|_X$, then there is an isomorphism

$$(12.2) \qquad\qquad\qquad F^* \mathcal{F} \xrightarrow{\sim} \mathcal{F}$$

induced from $\mathcal{F}_1$ via étale descent. The rough idea is that $\mathcal{F}_1$ is built up from étale algebraic spaces $E_1 \to X_1$. The relative Frobenius of $E = E_1 \otimes k$, given by $\sigma_{E_1} \otimes \mathrm{id}$, factors through an isomorphism $E \xrightarrow{\sim} (\sigma_{X_1} \otimes \mathrm{id})^* E = F^* E$, and the inverse isomorphisms $F^* E \xrightarrow{\sim} E$ give rise to (12.2). An arbitrary sheaf $\mathcal{F}$ equipped with an isomorphism of the form (12.2) is called a *Weil sheaf* with respect to $F : X \to X$.

Let $\sigma = \sigma_{\mathrm{Spec}\, k}$ in what follows. We can rewrite (12.2) in terms of σ^{-1} through the following trick. First, observe that

$$\sigma_X = (\mathrm{id}_{X_1} \otimes \sigma) \circ F = F \circ (\mathrm{id}_{X_1} \otimes \sigma).$$

Next, it turns out that $\sigma_X^* \simeq \mathrm{id}_{\mathsf{Shv}(X)}$ as functors, roughly because σ_X pulls back to σ_U for any étale open $U \to X$: See Lemma Achar 5.3.6. So we can rewrite (12.2) as an isomorphism

$$(12.3) \qquad (\mathrm{id}_{X_1} \otimes \sigma^{-1})^* \mathcal{F} \xrightarrow{\sim} \mathcal{F}.$$

We can also view σ^{-1} as a pro-generator of $\mathrm{Gal}(k/\mathbf{F}_q)$, sometimes called the *geometric Frobenius*. The isomorphism (12.3) forms a descent datum for $\mathcal{F}$ from X to X_1 if and only if it extends to a system of compatible isomorphisms for each element of $\mathrm{Gal}(k/\mathbf{F}_q)$. We will focus on the Weil sheaves for which this occurs: *i.e.*, those that take the form $\mathcal{F}_1|_X$ for some $\mathcal{F}_1$ on X_1.

Example 12.6. Via property (4) from §12.3.3, a sheaf on $pt = \mathrm{Spec}\, k$ that descends to $pt_1 = \mathrm{Spec}\, \mathbf{F}_q$ is equivalent to a continuous, finite-dimensional representation of $\mathrm{Gal}(k/\mathbf{F}_q)$ over $\bar{\mathbf{Q}}_\ell$. In particular, a 1-dimensional representation is equivalent to a homomorphism $\mathrm{Gal}(k/\mathbf{F}_q) \to \bar{\mathbf{Q}}_\ell^\times$. Continuity forces the image of σ^{-1} under such a homomorphism to be an element of $\bar{\mathbf{Z}}_\ell^\times$, the maximal compact subgroup of $\bar{\mathbf{Q}}_\ell^\times$.

We want to use (12.2) to define a Frobenius action on stalks. It turns out that for any $\bar{x} \in X(k)$, we have an identification

$$\mathcal{F}_{F(\bar{x})} = \varinjlim_{U \ni F(\bar{x})} \mathcal{F}(U) = \varinjlim_{V \ni \bar{x}} \mathcal{F}(F(V)) = (F^* \mathcal{F})_{\bar{x}}.$$

One can check that if $\bar{x}$ factors through a point $x \in X_1(\mathbf{F}_{q^d})$, then F^d fixes $\bar{x}$. So for such $\bar{x}$ and x, we obtain a morphism of $\bar{\mathbf{Q}}_\ell$-vector spaces

$$F^d : \mathcal{F}_{\bar{x}} = F^* \mathcal{F}_{\bar{x}} \xrightarrow{(12.2)} \mathcal{F}_{\bar{x}}.$$

The trace of this morphism only depends on x, so we can set

$$\mathbf{t}_{\mathcal{F},d}(x) = \mathrm{tr}(F^d \mid \mathcal{F}_{\bar{x}}).$$

Remark 12.7. Above, the action of F on $X(k)$ can be rewritten in terms of σ^{-1}, as $\bar{x} \mapsto (\mathrm{id}_{X_1} \otimes \sigma) \circ \bar{x} \circ \sigma^{-1}$. Similarly, the action of F^d on $\mathcal{F}_{\bar{x}}$ can be rewritten in terms of σ^{-d}: hence, in terms of $\mathrm{Gal}(k/\mathbf{F}_{q^d})$.

In the setting where $X = \mathrm{Spec}\, \mathbf{F}_{q^d}$ and we replace $\mathcal{F}$ with a lisse sheaf of $\mathfrak{k}$-modules, for some finite self-injective ring $\mathfrak{k}$, the map $\mathcal{F} \mapsto \mathcal{F}_{\bar{x}}$ induces the equivalence between lisse sheaves of $\mathfrak{k}$-modules and continuous representations of $\pi_1^{\text{ét}}(X, \bar{x}) = \mathrm{Gal}(k/\mathbf{F}_{q^d})$ over $\mathfrak{k}$.

12.6.

The above discussion gives the impression that we should replace $\mathsf{Shv}(X)$ with $\mathsf{Shv}(X_1)$, viewed as a full subcategory of the category of Weil sheaves on X.

Unfortunately, we have now overshot the size of the Grothendieck groups. Again taking G to be a torus, so that $X = pt$, we find from Example 12.6 that $[\mathsf{Shv}(pt_1)]_\oplus = [\mathsf{Shv}(pt_1)] = \mathbf{Z}[\bar{\mathbf{Z}}_\ell^\times]$.

So we should only use a certain subcategory of $\mathsf{Shv}(X_1)$. Based on our example, we might try to restrict the possible eigenvalues that occur in the action of F^d on $\mathcal{F}_{\bar{x}}$ discussed above. For instance, we could restrict the eigenvalues to be powers of $q^{1/2}$.

But more issues arise when we try to define a monoidal product $*$ that corresponds to the multiplication in H_W. From previous lectures, we expect $*$ to arise from some kind of convolution on X that should preserve our subcategory. A priori, it is not clear how to ensure that restrictions on eigenvalues would be preserved by this convolution.

Finally, we also want our subcategory to contain objects that, under the function-sheaf dictionary, recover the two bases that we have studied in detail: the standard basis $(\sigma_w)_w$ and the Kazhdan–Lusztig basis $(c_w)_w$. In future lectures, we will find that the approach that works is:

(1) First, categorify H_W to a merely additive category C preserved by a convolution $*$, the element x to a *new* kind of shift functor $\langle 1 \rangle$, and the basis $(c_w)_w$ to a collection of objects that generate this category under $\oplus$ and $\langle 1 \rangle$.

(2) Next, form objects in the triangulated category $\mathsf{K}^b(\mathsf{C})$ that categorify the standard basis $(\sigma_w)_w$.

We can give a preview of how $\langle 1 \rangle$ arises.

Recall that the *Tate twist* $\bar{\mathbf{Q}}_\ell(1)$ is the 1-dimensional vector space on which F acts by $q^{-1} \in \bar{\mathbf{Z}}_\ell^\times$. (Here we use the hypothesis that $\ell \nmid q$.) Fixing a square root of q allows us to define a *half Tate twist* $\bar{\mathbf{Q}}_\ell(\frac{1}{2})$, on which F acts by $q^{-1/2}$. We will construct C inside a larger triangulated category called the constructible derived category, then set $\langle 1 \rangle = (-) \otimes \bar{\mathbf{Q}}_\ell(\frac{1}{2})[1]$.