POINCARÉ DUALITY IN ÉTALE COHOMOLOGY, PART 1

Contents

1. Introduction 1
2. Why derived categories? 2
2.1. Ordinary Poincaré duality 2
3. Derived categories 4
4. Derived categories of étale sheaves 5
5. Poincaré duality in étale cohomology: statement 5
5.1. Compactly supported étale cohomology 5
5.2. The statement 7

1. Introduction

Today, we're going to start proving Poincaré duality in étale cohomology. But most of the 'work' will be in setting up some interesting category theory.

Peter May loves to say "the purpose of category theory is to make that which is trivial, trivially trivial." I think this fits the category theory we're going to do very well. Let me remind you of the proof (and statement!) of Poincaré duality for smooth manifolds.

Theorem 1.1 (Poincaré duality)**.** *Let M be a smooth, oriented closed d-manifold. Then there exists an isomorphism*

$$H^k(M;\mathbb{Z}) \simeq H_{d-k}(M;\mathbb{Z}),$$

for all integers k. This isomorphism is functorial in the oriented manifold M.

There are many proofs of Poincaré duality. Poincaré's proof – which I still think is the most geometric – goes by triangulating M and then replacing the polytopes in the triangulation with their duals. Unfortunately, this kind of proof, as far as I know, just has no hope of working in étale cohomology; we don't really have triangulations of smooth schemes.

So, we're going to copy a different proof of Poincaré duality, one which you might have seen in Hatcher's book. I don't want to prove Poincaré duality for smooth manifolds, so let me just state the process Hatcher goes through.

First, realize Poincaré duality can actually be generalized.

Theorem 1.2 (Poincaré duality, revisited)**.** *Let M be a smooth, oriented d-manifold. Then there exists an isomorphism*

$$H^k_c(M;\mathbb{Z}) \simeq H_{d-k}(M;\mathbb{Z}),$$

for all integers k, which is functorial in the oriented manifold M.

Here, $H^k_c(M;\mathbb{Z})$ is compactly supported cohomology.

Date: October 26, 2026.

By the 'paradox of induction,' Theorem 1.2 is easier to prove than Poincaré duality stated for closed manifolds – but why? The trick is a very simple two-step argument:

(1) prove Theorem 1.2 for $\mathbb{R}^d$ by a direct calculation,
(2) prove that Theorem 1.2 can be checked 'locally' on M, and hence (using that M is locally diffeomorphic to $\mathbb{R}^d$) reduce to the case of $\mathbb{R}^d$, where it follows from the first calculation.

When you read this proof in Hatcher's book, it's really easy to think that step 1 is trivial and step 2 is hard. But actually I think the situation is reversed: step 1 is hard and step 2 is trivial!

The reason step 2 seems so hard in Hatcher's text is that he doesn't develop the correct category theory that makes it trivial. In today's talk, we'll attempt to motivate this bit of category theory, and hence finish step 2 of the argument. Next time, I'll speak about step 1.

2. Why derived categories?

Why should step 2 be true? That is, why should Poincaré duality be a local statement?

2.1. **Ordinary Poincaré duality.** To make things concrete, choose M a smooth oriented d-manifold.

First let me make a small maneuver. As an algebraic geometer, local statements scream "sheaf!" at me. But when I think of sheaves, cohomology seems good, and homology scares me a little. So, we're going to try and eliminate the homology. Normally turning homology into cohomology requires some universal coefficients theorem; this can be done easily if you use coefficients in $D(\mathbb{Z})$, but I don't want to do that for the purposes of this talk, and there's a simplification we can make anyways: in étale cohomology, we care about rings like $\mathbb{Z}/\ell^n$, and those rings are always injective as modules over themselves. The universal coefficient theorem tells us that

$$H_i(M;\Lambda)^\vee \simeq H^i(M;\Lambda)$$

for any ring Λ which is injective over itself, and so let us actually just try and prove

$$H^k(M;\Lambda) \simeq H_c^{d-k}(X;\Lambda)^\vee \tag{2.1}$$

for self-injective rings Λ. How can we make this isomorphism sheaf-y?

The left hand side $H^k(M;\Lambda)$ is just the cohomology of the constant sheaf $\underline{\Lambda}$ on M. So, it's natural to ask the following question.

Question 2.2. Is there some sheaf $\mathcal{F}$ whose cohomology in degree k is 'naturally' $H_c^{d-k}(X;\Lambda)^\vee$?

If we could find such an $\mathcal{F}$, we could hope to prove (2.1) by showing $\mathcal{F}$ was isomorphic to $\underline{\Lambda}$. Unfortunately, it seems quite hard to invent such a sheaf $\mathcal{F}$.

However, there *is* an 'obvious' guess if we expand our mind from sheaves of abelian groups to sheaves of *chain complexes*.

Let's recall how cohomology and compactly-supported cohomology are defined. There is some cochain complex $C^\bullet(M;\Lambda)$, and

$$H^k(M;\Lambda) = H^k(C^\bullet(M;\Lambda)).$$

This cochain complex can be refined to a natural presheaf: define

$$U \mapsto C^\bullet(U; \Lambda)$$

sending an open U to the usual cochain complex with coefficients in Λ. Here, the restriction maps come from the fact that this cochain complex is contravariant in U (you should think that a cochain is a linear functional on chains; since chains are covariant, functionals on them are contravariant).

Unfortunately, this presheaf *is not a sheaf*!

Example 2.3. Say $M = \mathbb{R}$. Remember that an element of $C^1(\mathbb{R}; \Lambda)$ is a Λ-linear functional on 1-simplices.

Consider the open cover $\mathbb{R} = (-\infty, +1) \cup (-1, +\infty)$. Let σ be the 1-simplex corresponding to $[-2, 2]$.

Then $\sigma^\vee$ (the cochain which sends σ to 1 and every other simplex to 0) is zero on $(-\infty, +1)$ and zero on $(-1, +\infty)$, but clearly nonzero on the whole of M. Thus we do not have a sheaf!

The trouble, basically, is that only the quasi-isomorphism types of these cochain complexes are meaningful; neither complex alone is all that useful.

However, the Mayer–Vietoris sequence actually solves this problem.

Remark 2.4. Say $M = U \cup V$. The usual sheaf axiom for a presheaf $\mathcal{F}$ demands

$$\mathcal{F}(M) \simeq \ker\left(\mathcal{F}(U) \oplus \mathcal{F}(V) \to \mathcal{F}(U \cap V)\right), \tag{2.5}$$

where the map we take the kernel of is the 'subtraction restrictions' map.

Mayer–Vietoris says that there's a long exact sequence of cohomologies

$$\cdots \to H^k(M) \to H^k(U) \oplus H^k(V) \to H^k(U \cap V) \to H^{k+1}(M) \to \cdots,$$

where the map in the middle is some version of "subtract restrictions."

How is Mayer–Vietoris *proven*? Usually by showing that there is an exact sequence of chain complexes

$$0 \to C^\bullet_{U,V}(M) \to C^\bullet(U) \oplus C^\bullet(V) \to C^\bullet(U \cap V) \to 0,$$

where $C^\bullet_{U,V}(M)$ is the dual to the chain complex $C_\bullet^{U,V}(M) \subseteq C_\bullet(M)$ spanned by chains supported on either U or V.

Thus the Mayer–Vietoris sequence actually tells us the following:

$$U \mapsto C^\bullet(U; \Lambda)$$

is morally a sheaf of *chain complexes up to quasi-isomorphism*, where 'up to' means we work in the category obtained by formally inverting quasi-isomorphisms.

Warning 2.6. I say 'morally' here because this statement is not literally true; the trouble is that in this "category of chain complexes up to quasi-isomorphism," limits are very poorly behaved. But this is really a defect of 1-category theory giving the wrong notion of limits, and ∞-categories fix this. Next semester I will run a learning seminar on ∞-categories if you want to learn how to understand this problem!

Remark 2.7. In fact, if you read Hatcher's proof of Poincaré duality, he uses Mayer–Vietoris in his proof of Poincaré duality, exactly in his proof that Poincaré duality can be checked locally. So the fact that we see it cropping up when doing sheaf theory suggests we're on the right track!

Similarly, compactly-supported cohomology is just the cohomology of the cochain

$$C_c^\bullet(U;\Lambda) := \mathrm{colim}_{K\subseteq U} C^\bullet(U, U\setminus K;\Lambda),$$

where on the right we're taking relative cohomology of U relative to $U\setminus K$. Equivalently, $C_c^\bullet(U;\Lambda)$ is the subcomplex of $C^\bullet(U;\Lambda)$ consisting of cochains whose supports are compact.

Unfortunately, if $U\subseteq V$, then the restriction map

$$C^\bullet(V;\Lambda)\to C^\bullet(U;\Lambda)$$

does not respect the compact support condition at all: the intersection of a compact set and an open set need not be compact! However, if you've studied cohomology with compact supports, you might know that it acts *covariantly* by extension by zero. That is, there are natural maps

$$j_!: H_c^q(U;\Lambda)\to H_c^q(V;\Lambda)$$

for all q. These maps unfortunately are *not* induced by a map of chain complexes! So, where do these maps on cohomology come from? There is a natural map

$$\mathrm{colim}_{K\subseteq U} C^\bullet(X, X\setminus K;\Lambda)\to C_c^\bullet(X;\Lambda),$$

and a natural quasi-isomorphism

$$\mathrm{colim}_{K\subseteq U} C^\bullet(X, X\setminus K;\Lambda)\to C_c^\bullet(U;\mathbb{Z}),$$

with this latter map coming as the filtered colimit of the maps

$$C^\bullet(X, X\setminus K;\Lambda)\to C^\bullet(U, U\setminus K;\Lambda).$$

Thus $C_c^\bullet(U;\Lambda)$ is actually a *co-presheaf* valued in the category of "chain complexes up to quasi-isomorphism". However, we're interested in the dual; so, we look at the presheaf

$$U\mapsto \mathrm{Hom}(C_c^\bullet(U;\Lambda),\Lambda)[-d].$$

Here, the $[-d]$ is a normalizing convention which moves our complex forward by d; this way, the cohomology in degree k will be $H_c^{d-k}(U;\Lambda)^\vee$ instead of $H_c^{-k}(U;\Lambda)^\vee$, making for a more convenient comparison.

Remark 2.8. Observe how we turned a cosheaf into a sheaf by dualizing. This is why sheaf duals are computed by doing things like saying

$$\mathcal{F}^\vee := \underline{\mathrm{Hom}}(\mathcal{F}, \underline{\Lambda}),$$

instead of section-wise constructions!

3. Derived categories

At the end of the last section, we saw the need to consider chain complexes *up to quasi-isomorphism*. This motivates *derived categories*.

Definition 3.1. Let $\mathcal{A}$ be an abelian category. The *derived category* of $\mathcal{A}$ is the category obtained by starting with $\mathrm{Ch}(\mathcal{A})$, and then formally inverting all quasi-isomorphisms. We denote the derived category by $D(\mathcal{A})$.

That is, $D(\mathcal{A})$ is the 1-category obeying the following universal property: functors

$$D(\mathcal{A})\to\mathcal{C}$$

are the same as functors $\mathrm{Ch}(\mathcal{A})\to\mathcal{C}$ which send quasi-isomorphisms to isomorphisms in $\mathcal{C}$.

Warning 3.2. There are some set-theoretic troubles with making this definition work. Really one should require $\mathcal{A}$ is Grothendieck abelian for this to be a reasonable object.

I won't dwell on what Grothendieck means; think of it as a condition saying our category is generated by a set-theoretically small amount of information, even if the category itself is large. If this is your first time learning what a Grothendieck abelian category is, then I guarantee that every abelian category you know is Grothendieck abelian so I would suggest not worrying. The one thing I will say is that you can represent the derived category very explicitly in this case, as a category whose objects are "K-injective complexes," whatever that means, and whose morphisms are homotopy classes of chain maps.

If $f : \mathcal{A} \to \mathcal{B}$ is a left exact functor of Grothendieck abelian categories, then applying it termwise to K-injective complexes we can define its right derived functor

$$Rf : D(\mathcal{A}) \to D(\mathcal{B}).$$

If f is *exact*, then actually it preserves quasi-isomorphisms on the nose, and so f extends automatically to $D(\mathcal{A}) \to D(\mathcal{B})$, without doing this trick of injectives.

4. Derived categories of étale sheaves

Let me make a few technical assumptions now. Fix a prime ℓ. Set $\Lambda = \mathbb{Z}/\ell^n$.

Notation 4.1. Throughout this note, *scheme* means **qcqs $\mathbb{Z}[\frac{1}{\ell}]$-scheme.** Here, qcqs means quasicompact quasiseparated.

Notation 4.2. Here, *étale sheaf* means **étale sheaf of Λ-modules**. That is, when I talk about étale sheaves $\mathcal{F}$ on a scheme X, I will always assume that you've fixed a $\underline{\Lambda}_X$-module structure on $\mathcal{F}$. This is because étale cohomology is a bit subtle for coefficients which are not torsion, and I don't want to deal with this.

We write $D(X_{\text{ét}}; \Lambda)$ for the derived category of $\mathrm{Sh}(X_{\text{ét}}; \Lambda)$.

5. Poincaré duality in étale cohomology: statement

5.1. **Compactly supported étale cohomology.** What is compactly supported étale cohomology? We will write $H^i_c(X; \mathcal{F})$ for the compactly supported cohomology of an étale sheaf of $\mathcal{F}$.

Well, if X is some proper scheme, probably compactly supported étale cohomology should just equal étale cohomology. Let's then define

$$H^i_c(X; \mathcal{F}) := H^i(X; \mathcal{F}) \text{ when } X \text{ is proper.}$$

When X is not proper, what should we do?

Well, in decent situations we can find an open immersion $j : X \hookrightarrow \overline{X}$ where $\overline{X}$ is proper. We can hope to define compactly supported cohomology on X using cohomology on $\overline{X}$.

When we take an étale sheaf $\mathcal{F}$ on X, you already know one way to get an étale sheaf on $\overline{X}$: write down $Rj_*\mathcal{F}$. Unfortunately, the cohomology $H^i(\overline{X}, Rj_*\mathcal{F})$ is always equal to $H^i(X, \mathcal{F})$, when we want the compactly supported cohomology to differ in this non-proper case, somehow remembering that X is not proper.

However, you might remember from your Hartshorne book a different way to do compactly supported cohomology: $j_!$. We can make this definition work in étale cohomology.

Definition 5.1. Let $j : X \hookrightarrow \overline{X}$ be an étale morphism of schemes. We define *extension by zero*

$$j_! : \mathrm{Sh}(X_{\text{ét}}; \Lambda) \to \mathrm{Sh}(\overline{X}_{\text{ét}}; \Lambda)$$

as follows: $j_!\mathcal{F}$ is the sheafification of the étale presheaf

$$V \mapsto \bigoplus_{\varphi \in \mathrm{Hom}_{\overline{X}}(V, X)} \mathcal{F}(V \xrightarrow{\varphi} X).$$

Remark 5.2. It is easy to show that, when $\bar{x}$ is a geometric point of $\overline{X}$,

$$(j_!\mathcal{F})_{\bar{x}} = \bigoplus_{\bar{x}' \in j^{-1}(\bar{x})} \mathcal{F}_{\bar{x}'}.$$

By this stalk computation, $j_!$ is exact; it therefore extends automatically to a functor

$$j_! : D(X_{\text{ét}}; \Lambda) \to D(\overline{X}_{\text{ét}}; \Lambda),$$

computed by applying $j_!$ to complexes termwise.

We then define

$$H^i_c(X; \mathcal{F}) := H^i(\overline{X}, j_!\mathcal{F}).$$

This is actually independent of the choice of compactification; moreover, this operation can be done in families.

Theorem 5.3. *Let $f : X \to S$ be a finite-type, separated morphism of schemes (remember: scheme means qcqs to us!). By Nagata compactification (see tag 0F41 in the stacks project), we may find a factorization of f as*

$$X \xrightarrow{j} \overline{X} \xrightarrow{f'} S$$

where f' is proper and j is an open immersion.

We define

$$Rf_! : D(X_{\text{ét}}; \Lambda) \to D(S_{\text{ét}}; \Lambda)$$

by

$$Rf_!(\mathcal{F}) := Rf'_*(j_!\mathcal{F}).$$

This functor $Rf_!$ is independent of the choice of Nagata compactification of f, up to canonical natural isomorphism of functors.

Proof. See tag 0F7H in the stacks project for a proof of independence of compactification; it is an obvious kind of diagram chase, using proper base change. That section of the stacks project also proves that all the 'expected' properties of $Rf_!$ hold. □

Warning 5.4 (Do not do this non-derivedly!)**.** Despite the notation $Rf_!$, it is *not true* that $Rf_!$ is the right derived functor of

$$f_! := R^0 f_! : \mathrm{Sh}(X_{\text{ét}}; \Lambda) \to \mathrm{Sh}(S_{\text{ét}}; \Lambda).$$

The trouble is that $j_!$ has no reason to send injectives to acyclics. This is even false when $X = \mathbb{A}^1_k$, which is easy to see just by unwinding definitions; we leave it as an exercise to the reader.

Remark 5.5. When f is étale, $f_! = Rf_!$.

In the sequel, we will need the following.

Theorem 5.6 (Projection formula)**.** *For $f : X \to S$ a finite-type separated morphism, there is a (functorial in $\mathcal{F}, \mathcal{G}$) isomorphism*

$$Rf_!(\mathcal{F} \otimes^{\mathbb{L}}_{\Lambda} f^*\mathcal{G}) \simeq Rf_!\mathcal{F} \otimes^{\mathbb{L}}_{\Lambda} \mathcal{G}.$$

Proof. See 0GL5 in the stacks project, but it is similar to the usual projection formula. □

5.2. **The statement.** With compactly supported étale cohomology in mind, we can try formulating a very general Poincaré duality theorem.

Take $f : X \to S$ a smooth, separated morphism of relative dimension d.

Poincaré duality should compare cohomology and compactly supported cohomology; so we should expect to be comparing $f_*\underline{\Lambda}_X$ and $f_!\underline{\Lambda}_X$. More precisely, we expect something like

$$Rf_*\underline{\Lambda}_X \approx R\underline{\mathrm{Hom}}(Rf_!\underline{\Lambda}_X, \underline{\Lambda}_S)[-2d]. \tag{5.7}$$

Warning 5.8. (5.7) is actually not quite right! We will fix it soon; for now just know we expect something like it to hold.

At this point, Verdier realized something: the easiest way to prove something like (5.7) would be if $Rf_!$ admitted a right adjoint. And fortunately, it does for 'silly' reasons: the adjoint functor theorem. More precisely,

(1) for usual derived categories, *Brown representability* tells us the existence;
(2) for derived ∞-categories, it turns out that $D(X_{\text{ét}}; \Lambda)$ is *presentable* and so one has the literal adjoint functor theorem.

We write

$$f^! : D(S_{\text{ét}}; \Lambda) \to D(X_{\text{ét}}; \Lambda)$$

for the right adjoint to $Rf_!$.

Theorem 5.9 (Poincaré duality in étale cohomology, part 1)**.** *Let $f : X \to S$ be smooth of relative dimension d. Then there is a (natural in $\mathcal{F}$) isomorphism*

$$f^!(\mathcal{F}) \simeq f^*(\mathcal{F}) \otimes f^!(\underline{\Lambda}_S).$$

Moreover,

(1) $f^!(\underline{\Lambda}_S)$ is $\otimes$-invertible (as an object of the derived category!),
(2) for any base change $g : S' \to S$, the natural map

$$(g')^* f^!(\underline{\Lambda}_S) \to (f')^!(\underline{\Lambda}_{S'})$$

is an isomorphism.

Remark 5.10. Scholze would say this theorem tells us that smooth morphisms are *cohomologically smooth*.

Remark 5.11. SGA required f to be separated. We no longer need such hypotheses, thanks to the modern proofs!

It is bad form to say 'there is an isomorphism,' so let me repent my sins by constructing a morphism.

Construction 5.12. Keep the notation of Theorem 5.9. We define

$$\eta : f^*(-) \otimes f^!(\underline{\Lambda}_S) \to f^!(-)$$

as follows.

By adjunction, it suffices to construct a natural transformation

$$Rf_!(f^*(-) \otimes f^!(\underline{\Lambda}_S)) \to \mathrm{id}\,.$$

But the projection formula tells us that

$$Rf_!(f^*\mathcal{F} \otimes f^!(\underline{\Lambda}_S)) = \mathcal{F} \otimes Rf_!f^!(\underline{\Lambda}_S),$$

so we just take $\mathrm{id}_{\mathcal{F}}$ tensored with the counit of the adjunction $f^!, Rf_!$.

When f is étale, we can prove Theorem 5.9 quite easily.

Lemma 5.13. *Theorem 5.9 holds when f is étale. In fact, in this case, we even have $f^!(\underline{\Lambda}_S) = \underline{\Lambda}_X$, so that $f^! = f^*$.*

Proof. For any geometric point $\bar{s}$,

$$(f_!\mathcal{F})_{\bar{s}} = \bigoplus_{\bar{x} \in X_{\bar{s}}} \mathcal{F}_{\bar{x}}.$$

As f is étale, $Rf_! = f_!$. From this stalk formula, it is easy to check that $f_!$ is left adjoint to f^*, so we conclude by uniqueness of adjoints. □

To prove Theorem 5.9, it suffices to prove it when our morphism is $\mathbb{A}^1_S \to S$.

Lemma 5.14. *Assume that, for every base S, the map $\mathbb{A}^1_S \to S$ obeys Theorem 5.9 (in the strong sense, that Construction 5.12 is a natural isomorphism). Then Theorem 5.9 is true in general.*

Proof. We leave it as an exercise to the reader to check that if Theorem 5.9 is true for two maps, it is true for their composition; thus writing $\mathbb{A}^d_S \to S$ as

$$\mathbb{A}^d_S \to \mathbb{A}^1_{\mathbb{A}^{d-1}_S} \to \mathbb{A}^1_{\mathbb{A}^{d-2}_S} \to \cdots \to \mathbb{A}^1_S \to S,$$

we deduce that $\mathbb{A}^d_S \to S$ is always cohomologically smooth.

Next we show that $f : X \to S$ being cohomologically smooth can be checked étale locally on the source X.

To check a morphism of étale sheaves is an isomorphism, it suffices to check étale locally; as η is a morphism of sheaves on X, we can check étale locally on X. It is not obvious, but just as being a line bundle can be checked étale locally, being $\otimes$-invertible in the derived category can be checked étale locally; we leave this to the reader. Lastly, we need to check that $(g')^*f^!(\underline{\Lambda}_S) \to (f')^!(\underline{\Lambda}_{S'})$ is an isomorphism for any base change $g : S' \to S$. But we can check this étale locally on $X \times_S S'$, and hence étale locally on X.

Given any smooth map $f : X \to S$, we can cover X by Zariski opens $j_i : U_i \hookrightarrow X$ such that there exist étale maps $x_i : U_i \to \mathbb{A}^d_S$ with $\pi_S \circ x_i = f \circ j_i$. From here we deduce each $f \circ j_i$ is cohomologically smooth (use Lemma 5.13 on x_i), and hence that f is cohomologically smooth. □

To finish our proof of Poincaré duality, we need to compute $f^!(\underline{\Lambda}_S)$, and to prove Theorem 5.9 for $\mathbb{A}^1_S$. This will be done next time.