The ℓ-adic Fourier Transform

Akshay Venkatesh*

Mary 24, 2017

1 Kloosterman sums

Let $\psi\colon \mathbf{F}_p \to \mathbf{C}^\times$ or $\overline{\mathbf{Q}}_p^\times$ be a character, e.g.

$$x \mapsto e^{2\pi i x/p}.$$

An important property is that

$$\sum_{a\in\mathbf{F}_p} \psi(ay) = \begin{cases} 0 & y \neq 0, \\ p & y = 0 \end{cases} \tag{1.1}$$

For $a \in \mathbf{F}_p^*$, the *Kloosterman sum* $K(a)$ is

$$K(a) = \sum_{xy=a} \psi(x+y) = \sum_{x\in\mathbf{F}_p, x\neq 0} \psi(ax + x^{-1}). \tag{1.2}$$

Example 1.1. $K(a)$ is always a real number, because the sum is symmetric with respect to complex conjugation. For $p = 7$ and $a = 1$, it is

$$\zeta_7^2 + \zeta_7^{-1} s + \zeta_7 + \zeta_7^{-2} + \zeta_7 + \zeta_y^{-1}.$$

Remark 1.2. The analogue of $K(a)$ over $\mathbf{R}$ would be something like

$$\int_{\mathbf{R}} e^{i(ax+1/x)}\, dx.$$

This isn't convergent, but

$$\int_{\mathbf{R}} e^{-(ax+1/x)}\, dx \sim \sqrt{a} K(\sqrt{a})$$

where K is a Bessel function. There is a parallel between these special functions and the special sheaves that will arise later.

*notes by Tony Feng

This was first studied by Kloosterman, in analyzing the Hardy-Littlewood circle method. Obviously $|K(a)| < p$; Kloosterman wanted to show that $|K(a)|$ is *much* less than p. He showed that

$$|K(a)| \leq p^{3/4}.$$

Weil later improved this to $|K(a)| \leq 2\sqrt{p}$.

We'll discuss Kloosterman's proof. He studied the sum

$$\sum_a |K(a)|^4$$

and showed

$$\sum_a |K(a)|^4 \leq cp^3.$$

We start off by writing

$$\sum_a |K(a)|^4 = \sum_a \sum_{x,y,z,w} \psi(a(x+y-z-w) + (x^{-1}+y^{-1}-z^{-1}-w^{-1})$$

Using (1.1), this is

$$= p \sum_{x+y=z+w} \psi(x^{-1}+y^{-1}-z^{-1}-w^{-1})$$

$$= p \left(\sum_{\substack{x+y=z+w \\ x^{-1}+y^{-1} \neq z^{-1}+w^{-1}}} \psi(x^{-1}+y^{-1}-z^{-1}-w^{-1}) + \# \left\{ \substack{x+y=z+w \\ x^{-1}+y^{-1}=z^{-1}+w^{-1}} \right\} \right)$$

Recall that we want to get a bound of p^3, while the trivial bound is about p^4. The second term $\# \left\{ \substack{x+y=z+w \\ x^{-1}+y^{-1}=z^{-1}+w^{-1}} \right\}$ has only about p^2 terms, so that's good. The first term is p times a sum over p^3 things, so we need to do something intelligent there. Luckily, it has a *scaling* symmetry. The sum over each $\mathbf{F}_p^*$-orbits is -1 by (1.1).

So, letting $N = \#\{x+y=z+w\}$ and $A = \# \left\{ \substack{x+y=z+w \\ x^{-1}+y^{-1}=z^{-1}+w^{-1}} \right\}$, we get

$$\sum_a |K(a)|^4 = p\left(\frac{N-A}{p-1}(-1) + A\right).$$

To conclude that, note that N has size about p^3, A has size about p^2.

Remark 1.3. You can evaluate N and A exactly.For odd p, we think $A = 3(p-2)(p-1)$. This gives

$$\sum_a |K(a)|^4 = 2p^3 + \text{ (lower order terms)}.$$

Kloosterman's argument is the earliest instance I know of the following principle:to bound a single value of a function, put that function in a family and raise it to a higher power. This idea was used again by Rankin in the context of modular forms, which Deligne said was an inspiration for his proof of the Weil conjectures.

2 The sheaf-function correspondence

We now want to implement this idea with sheaves. Let X be a variety over $\mathbf{F}_p$. Suppose you have a Weil sheaf $\mathcal{F}$ on X, meaning a sheaf on $X_{\overline{\mathbf{F}}_p}$ with a Frobenius endomorphism. Then we get a function f on $|X|$ or $X(\mathbf{F}_{p^n})$, given by $f(x) =$ trace of geometric Frobenius at x.

Remark 2.1. If for example $\mathcal{F}$ is lisse and semisimple, then the function f determines the sheaf, because the Frobenii are dense in the monodromy group.

2.1 Translation of sheaf-theoretic operations

Operations on sheaves can be translated into operations on functions.

- The tensor product of sheaves translates into product of functions.

$$\mathcal{F} \otimes \mathcal{G} \mapsto f_{\mathcal{F}} \cdot f_{\mathcal{G}}$$

- The pulblack of sheaves translates into pullback of functions.

$$\pi^* \mathcal{F} \mapsto f_{\mathcal{F}} \circ \pi$$

- If $\mathcal{F}$ is pure of weight $w \in \mathbf{Z}$, then $\mathcal{F}^\vee$ correponds to the function $\overline{f} p^{-w \cdot \deg}$. For instance, the Kloosterman sums being real-valued corresponds to the Kloosterman sheaves being self-dual up to Tate twist.

- If $\mathcal{F}$ is in the derived category of Weil sheaves[1] then

$$R\pi_! \mathcal{F} \mapsto \sum_i (-1)^i f_{H^i \mathcal{F}}$$

  With these conventions, the Lefschetz trace formula translates into the the statement that the derived pushforward of sheaf corresponds to the pushforward of f as defined by

$$y \in Y(\mathbf{F}_{p^n}) \mapsto \sum_{x \in X(\mathbf{F}_{p^m}), \pi(x) = y} f(x).$$

[1] Although we have glided over this point in this seminar, the construction of the "derived category of ℓ-adic sheaves" (or Weil sheaves) is actually quite subtle. It is not obtained by the naïve construction taking the derived category of a category of ℓ-adic sheaves, although this is often what one pretends for practical purposes. Suffice it to say that working rigorously with the "derived category of ℓ-adic sheaves" requires a good deal more care than one might think; "arguments" which treat this category as a genuine derived category are merely reasoning by analogy.

For example, given a map $\pi\colon X \to \mathbf{A}^1$, then the function

$$y \in \mathbf{F}_{p^m} \mapsto \# X_y(\mathbf{F}_{p^m})$$

comes from $R\pi_! \underline{\mathbf{Q}_\ell}$.

Example 2.2. We're going to make a sheaf corresponding to the function ψ.

We start out with the Artin-Schreier cover

$$y^p - y = x \subset \mathbf{A}^2.$$

This maps via the x-coordinate to $\mathbf{A}^1$, which is an étale cover. The Galois group is canonically $\mathbf{Z}/p\mathbf{Z}$, generated by $y \mapsto y+1$. In other words, this cover induces a map

$$\pi_1(\mathbf{A}^1_{\mathbf{F}_p}) \to \mathbf{Z}/p\mathbf{Z} \xrightarrow{\psi} \overline{\mathbf{Q}}_\ell^*.$$

Let $\mathcal{L}_\psi$ be the corresponding rank 1 lisse sheaf on $\mathbf{A}^1$. We compute the associated function. We need to figure out where Frobenius goes. The Frobenius at $x \in \mathbf{A}^1(\mathbf{F}_{p^m})$ takes $(x,y) \mapsto (x^{p^m}, y^{p^m})$. Of course $x^{p^m} = x$. We have

$$\begin{aligned} y^p &= y + x \\ y^{p^2} &= y^p + x^p = y + x + x^p \\ &\vdots \\ y^{p^m} &= y + x + x^p + \ldots + x^{p^{m-1}} \end{aligned}$$

Therefore, the Frobenius at x takes

$$(x,y) \mapsto (x, y + \underbrace{x + x^p + x^{p^2} + \ldots + x^{p^{m-1}}}_{\mathrm{Tr}_{\mathbf{F}_{p^m}/\mathbf{F}_p}(x)}).$$

The conclusion is that Frobenius acts on the stalk $\mathcal{L}_\psi$ as multiplication by $\psi(\mathrm{Tr}_{\mathbf{F}_{p^m}/\mathbf{F}_p}(x))$. Therefore *geometric* Frobenius acts as multiplication by $\psi(-\mathrm{Tr}_{\mathbf{F}_{p^m}/\mathbf{F}_p}(x)) = \overline{\psi}(\mathrm{Tr}_{\mathbf{F}_{p^m}/\mathbf{F}_p}(x))$.

2.2 The method of families

Suppose we have an open subset $U \subset \mathbf{A}^1$, and $\mathcal{G}$ is a Weil sheaf on U, associated to a function g. Assume $g \geq 0$. (This can be arranged by taking the sum of $\mathcal{G}$ with its conjugate.)

We can then bound a single value of g by a sum:

$$g(x) \leq \sum_{x \in \mathbf{F}_{p^n}} g(x).$$

Of course this isn't sharp, but if you apply this to large powers of g then it will be sharp.

To simplify things, assume $H_c^0(G) = 0$. Then

$$\sum_{x \in U(\mathbf{F}_{p^m})} g(x) = \sum_{\beta = \text{ eig. of } F \text{ on } H_c^2} \beta^m - \sum_{\alpha = \text{ eig. of } F \text{ on } H_c^1} \alpha^m.$$

We can easily analyze the β's, by using Poincaré duality to relate H_c^2 to the coinvariants of geometric π_1. But the point is that $\max |\alpha| \leq \max |\beta|$, which allows you to ignore α. Why? This is because the expression is positive.

Remark 2.3. This same observation appears elsewhere. For instance, the Weil bound is not optimal. For a curve it gives $p + 1 + 2g\sqrt{p}$, but this can't be attained because it would give a negative number of points over $\mathbf{F}_{p^2}$.

This argument (applied to a high power of g) is the engine that provides the bounds in the proof.

The central part of the proof will be the following statement: if $\mathcal{F}$ is pure of weight w on $U \subset \mathbf{A}^1$, then the weights of $H_c^1(U, \mathcal{F})$ are $\leq w + 1$. In terms of the functions $f_{\mathcal{F}}$ associated to $\mathcal{F}$, this statement translates to the bound

$$\sum_{x \in U(\mathbf{F}_{p^n})} f_{\mathcal{F}}(x) \leq p^{m(w+1)/2}.$$

If $H_c^0 = 0$ and $H_c^2 = 0$ (it is easy to reduce to this case), then

$$\mathrm{Tr}(\mathrm{Frob}^m, H_c^1) = -\sum_{x \in U(\mathbf{F}_{p^n})} f_{\mathcal{F}}(x)$$

so this becomes a question of bounding the eigenvalues of Frobenius on cohomology.

To obtain this estimate, we try to embed it into a family. We want to find a function g on $\mathbf{A}^1$ (associated to a sheaf) such that $g(0 \in \mathbf{F}_{p^m}) = \sum f_{\mathcal{F}}(x)$. Then we'll use the method of families. The punchline is that we take g to be the Fourier transform of f. Si next we'll make a sheaf associated to the function $g = \mathrm{FT}(f)$, defined by

$$\mathrm{FT}(f_{\mathcal{F}})(y) = \sum_{x \in \mathbf{F}_{p^m}} f_{\mathcal{F}}(x) \psi \circ \mathrm{Tr}(yx)$$

for $y \in \mathbf{F}_{p^m}$.

3 Fourier transform

Let $\mathcal{F}$ be a sheaf on $\mathbf{A}^1$. We make a new sheaf $\mathrm{FT}_\psi \mathcal{F}$ such that

$$f_{\mathrm{FT}_\psi}(y) = \sum_{x \in \mathbf{F}_{p^m}} f_{\mathcal{F}}(x) \psi \circ \mathrm{Tr}(yx).$$

We just replicate the Fourier transform step-by-step.

We start with $\mathcal{F}$, pull it back to $\mathbf{A}^2$ via $(x,y)\mapsto x$. Then we tensor with $m^*\mathcal{L}_\psi$, where $m(x,y)=xy$. Finally, to sum over the first variable we push forward via $(x,y)\mapsto y$. The last step is to shift by degree 1, basically to preserve the property of being a sheaf (but it still might not quite).

Denote this functor by FT_ψ.

Theorem 3.1. *We have*

$$\mathrm{FT}_{\overline{\psi}} \circ FT_\psi = \mathrm{Id} \text{ (up to Tate twist).}$$

This mirrors the usual calculation

$$\begin{aligned}\sum_y \psi(-yz)\sum_x f(x)\psi(yx) &= \sum_{x,y} f(x)\psi(y(x-z)) \\ &= p^m f(z).\end{aligned}$$

The proof replicates this calculation at the level of sheaves. The only step that wasn't formal was the calculation

$$\sum_{a\in\mathbf{F}_p}\psi(ay) = \begin{cases}0 & y\neq 0\\ p\end{cases}$$

so we need a sheaf-theoretic analogue of it, which is

$$H_c^*(\mathbf{A}^1_{\overline{\mathbf{F}}_p}, L_\psi)=0.$$

To prove this, recall that the sheaf L_ψ came from the covering

$$y^p - y = x$$

by taking the ψ-component of the pushforward of the constant sheaf. Then H_c^* is the $\psi^{\pm 1}$-component of $H_c^*(C,\overline{\mathbf{Q}}_\ell)$, which is 0 (since $C=\mathbf{A}^1$).

The idea of the proof of the Weil conjectures is to bound F-eigenvalues on$H_c^1(U\subset \mathbf{A}^1,\mathcal{G})$, which is the fiber at 0of $\mathrm{FT}_\psi(\mathcal{G})$.

4 Kloosterman sheaves

Recall that we defined the Kloosterman function

$$K(a)=\sum_{xy=a}\psi(x+y)=\sum_{x\in\mathbf{F}_p, x\neq 0}\psi(ax+x^{-1}).$$

We're going to make a sheaf Kl on $\mathbf{G}_m$ such that

$$f_{\mathrm{Kl}}(a\in\mathbf{F}_{p^m}) = \sum_{x\in\mathbf{F}_{p^m}^*}\psi\circ\mathrm{Tr}(ax+x^{-1}).$$

We start with L_ψ on $\mathbf{G}_m$ to get $\psi(x)$, apply inversion to get $\psi(x^{-1})$, and apply FT_ψ.

This gives a lisse sheaf Kl on $\mathbf{G}_m$, pure of weight 1. Since $\mathrm{rank}(\mathrm{Kl}) = 2$, this corresponds to a representation $\pi_1(\mathbf{G}_m) \to \mathrm{GL}_2(\overline{\mathbf{Q}}_\ell)$ whose Zariski closure is SL_2 (the real-ness suggests the sheaf is self-dual).

Suppose we want to understand Kl_a for $a \in \overline{\mathbf{F}}_p$. Take $a = 1$. It is the ψ-component of $H^1_c(y^p - y = x + x^{-1})$. We will show that $\dim H^1_c(y^p - y = x + x^{-1}) = 2(p-1)+1$. This strongly suggests that, because there are $p-1$ characters ψ, each piece has dimension 2.

If you actually want to compute, you have to understand the behavior of the sheaf at ∞. Consider

$$y^p - y = x + x^{-1} \to \mathbf{G}_m$$

and compactify it to $X \to \mathbf{P}^1$, of degree p.

By Riemann-Hurwitz,

$$2g_X - 2 = p(-2) + \deg(\text{ram. divisor}).$$

The ramification is supported at $0, \infty$. Since the equation $y^p - y = x + x^{-1}$ is symmetric, the answer will be the same at both points, so we just to the calculation 0. Localizing at 0, we need to consider the field extension L/K where $K = \mathbf{F}_p((x))$ and $L = K(y)$ with $y^p - y = x + x^{-1}$. Since $v(x) = 1$, we have $v(y) = -1/p$, $\tau = y^{-1}$ is a uniformizer. The discriminant is the field extension

$$\prod_{i \neq j} (\tau_i - \tau_j)$$

Since the conjugates just add, a typical term is

$$\frac{1}{y+1} - \frac{1}{y} = \frac{-1}{y(y+1)}$$

with valuation $2/p$. So the discriminant has valuation $p(p-1)(2/p) = 2(p-1)$. (This is double what one would expect in characteristic 0.) So the conclusion is that

$$2g_X - 2 = p(-2) + 4(p-1) \implies \dim H^1(X) = 2(p-1).$$

To get C from the compactified guy, you delete 2 points so

$$\dim H^1(C) = 2(p-1) + 1.$$

So we've verified that $\mathrm{rank}\,\mathrm{KL} = 2$.

Remark 4.1. This is related to the fact that the K-Bessel function from Remark 1.2 satisfies a *second*-order differential equation.

Let's return to the estimate:

$$\sum |K(a)|^4 \sim 2p^3.$$

This sum can be interpreted as

$$\sum \mathrm{Tr}(F|H_c^i(\mathrm{Kl} \otimes \mathrm{Kl}^\vee \otimes \mathrm{Kl} \otimes \mathrm{Kl}^\vee)(-1)^i.$$

By Deligne, H_c^1 contributes a second-order term, so the leading term comes from

$$H_c^2 = (V \otimes V^\vee \otimes V \otimes V^\vee)_{\pi_1^{\mathrm{geom}}}(-1) = (V \otimes V^\vee \otimes V \otimes V^\vee)^{\pi_1^{\mathrm{geom}}}$$

We can interpret $V \otimes V^\vee \otimes V \otimes V^\vee = \mathrm{End}(V \otimes V^\vee)$. So we want to compute

$$\dim \mathrm{End}(V \otimes V^\vee)^{\pi_1^{\mathrm{geom}}}.$$

We have an irreducible decomposition of $V \otimes V^\vee$ into the direct sum of a 3-dimensional representation and a 1-dimensional representation, so there are indeed two independent π_1^{geom}-equivariant endomorphisms.