hexagon:2611.00789v1

A short proof of Deligne's bound for Frobenius eigenvalues, after Alpöge

Entered at deposit

21 November 2026 CC0 1.0

Abstract

We write out in full the proof, sketched by L. Alpöge in a footnote of his note of 20 November 2026, of Deligne's bound for the eigenvalues of Frobenius. Let X be a separated scheme of finite type over a finite field with q elements, and let $\mathcal{F}$ be a constructible ℓ-adic sheaf on X whose Frobenius eigenvalues at each closed point x have absolute value at most $|\kappa(x)|^{w/2}$. Then the eigenvalues of Frobenius on $H^i_c(X,\mathcal{F})$ have absolute value at most $q^{(w+i)/2}$, and they have absolute value exactly $q^{(w+i)/2}$ if X is smooth and proper and $\mathcal{F}$ is lisse and pure of weight w. The proof is short. Granted the formalism of étale cohomology (the trace formula, the Künneth formula, Poincaré duality and Artin's vanishing theorem) and a weak Lefschetz theorem for generic hyperplane sections of affine varieties, it takes less than six pages, and its only estimate is the Cauchy–Schwarz inequality for a character sum over a finite field. Lefschetz pencils, monodromy and Rankin's method are not used. The proof was produced by GPT-6 Astra (ChatGPT Pro mode) from that sketch. It was checked by Claude Opus 5.5 at effort high and, independently, by Claude Fable 5.2 at effort max, which also wrote this text. No human has checked it.

1 The theorem

Let $\mathbf{F}_q$ be a finite field, ℓ a prime not dividing q, and $\iota : \overline{\mathbf{Q}}_\ell \hookrightarrow \mathbf{C}$ an embedding of fields, fixed throughout. The absolute value of an element of $\overline{\mathbf{Q}}_\ell$ always means the absolute value of its image under ι. Let X be a separated scheme of finite type over $\mathbf{F}_q$ and $\mathcal{F}$ a constructible $\overline{\mathbf{Q}}_\ell$-sheaf on X. The groups $H^i_c(X,\mathcal{F})$ are those of $X \otimes \overline{\mathbf{F}}_q$; they are finite-dimensional, and the geometric Frobenius F acts on them. At a closed point x of X, with residue field $\kappa(x)$, the geometric Frobenius of $\kappa(x)$ acts on the stalk $\mathcal{F}_{\bar{x}}$. For a real number w, we say that $\mathcal{F}$ *has weights at most* w if at every closed point x all eigenvalues of this local Frobenius have absolute value at most $|\kappa(x)|^{w/2}$, and that $\mathcal{F}$ is *pure of weight* w if they all have absolute value exactly $|\kappa(x)|^{w/2}$.

Theorem 1. *Let X be a separated scheme of finite type over $\mathbf{F}_q$ and $\mathcal{F}$ a constructible $\overline{\mathbf{Q}}_\ell$-sheaf on X.*

(a) If $\mathcal{F}$ has weights at most w, then for every i every eigenvalue of F on $H^i_c(X,\mathcal{F})$ has absolute value at most $q^{(w+i)/2}$.

(b) If X is smooth and proper and $\mathcal{F}$ is lisse and pure of weight w, then for every i every eigenvalue of F on $H^i(X,\mathcal{F})$ has absolute value $q^{(w+i)/2}$.

Part (a) says that the passage from the stalks of $\mathcal{F}$ to its cohomology in degree i raises the weights by at most i. Part (b) says that on a smooth proper scheme nothing is lost: the cohomology in degree i of a pure sheaf is pure, of weight $w+i$. For the constant sheaf $\overline{\mathbf{Q}}_\ell$ on a smooth projective variety X the hypothesis of (b) holds with $w=0$ for every embedding ι,

so each eigenvalue of F on $H^i(X, \overline{\mathbf{Q}}_\ell)$ has absolute value $q^{i/2}$ under every embedding. This is the Riemann hypothesis for X.

Both statements are theorems of Deligne. Part (b) for the constant sheaf is [3, Théorème 1.6], apart from the assertions made there about integrality and independence of ℓ, which we do not treat. Part (a) is the statement about absolute values contained in [4, Théorème 3.3.1] for a scheme over a finite field. That theorem concerns a morphism $f : X \to S$ and a mixed sheaf, and says that $R^i f_! \mathcal{F}$ is mixed of weights at most $w + i$. For absolute values the relative statement follows from (a), applied to the fibres of f, because the stalks of $R^i f_! \mathcal{F}$ are the groups H^i_c of the fibres. Part (a) asks only for a bound at each closed point, for the one embedding ι, and does not assume that $\mathcal{F}$ is mixed. We do not claim this as a gain in generality. By results of L. Lafforgue, an irreducible lisse sheaf on a smooth curve over $\mathbf{F}_q$ whose determinant has finite order is pure of weight zero for every embedding [8, Théorème VII.6 (ii)], and on a normal scheme of finite type over $\mathbf{F}_q$ every lisse sheaf is, in a unique way, a direct sum $\bigoplus_\chi \sigma_\chi \otimes \chi$ of twists by characters χ [8, Corollaire VII.8]. Nothing is proved here about the other parts of Deligne's theory: the weight filtration, the local monodromy theorem and the hard Lefschetz theorem of [4], the rationality and the functional equation of zeta functions, or the semisimplicity of Frobenius.

2 The proof, its source and its length

The source. The note [1] of L. Alpöge mentions a proof of the theorems of Weil and Deligne that uses only a sign, based on an identity $qN = M^k + R$, and sketches it in a footnote [1, p. 2 and footnote 3]. In the sketch N is a weighted count of the zeros of the form $x_1x_2 + \cdots + x_{k-1}x_k$, written as a sum over additive characters, and M^k is the term of the trivial character. An argument with a hyperplane section gives the contribution of middle cohomology a definite sign, one degree away from that of M^k. The remainder R has square-root size, by Parseval's identity applied k times. Because of the sign, the inverse roots cannot cancel one another, so the bound for R is a bound for each of them. In higher dimension vectors replace the scalars x_i.

The present note writes this argument out in full. Its ideas are those of the sketch. What is supplied here is the detail: the reduction to the lowest cohomology group of a smooth affine variety, a finite map to affine space that turns the count into the character sum, the weak Lefschetz theorem that provides the sign, and the passage to all degrees.

The length. The proof is short, and we say in what sense. We take for granted the formalism of ℓ-adic cohomology as it stood before [3]: the trace formula, the Künneth formula, Poincaré duality, the vanishing of the cohomology of an affine scheme above its dimension, and the exact sequence attached to an open subscheme and its complement. We also take for granted one theorem of a less elementary kind: for a generic hyperplane section Y of a smooth affine variety Z, and a lisse sheaf on Z, restriction from Z to Y is bijective on cohomology below the dimension of Y and injective in that degree. The proof of it that we use is Deligne's, printed in the appendix of [6]. It is a few pages long and rests on his generic base change theorem and on Artin's vanishing theorem; Section 6 recalls it. With these inputs, the proof of Theorem 1 occupies Sections 3 to 8, less than six pages, and its only estimate is the Cauchy–Schwarz inequality, applied to a character sum over a finite field (Lemma 3).

For comparison, [3] uses the same formalism together with the theory of Lefschetz pencils, a theorem of Kazhdan and Margulis on their monodromy, and Rankin's method. In [4] the general statement is reduced to the case of curves, which is proved by a study of monodromy and by the method of Hadamard and de la Vallée-Poussin. None of these enters below. Laumon [9] gave another proof of the main theorem of [4], by the ℓ-adic Fourier transform, and Kedlaya

[7] carried that proof over to p-adic cohomology.

For curves, the estimate for the number of points has an elementary proof, by the method of Stepanov [2]. Scholl [10] reduced the Riemann hypothesis for smooth projective varieties to the case of hypersurfaces, and remarked that an elementary proof of the point-count estimate for hypersurfaces, in the manner of Stepanov, would be of great interest. The proof below is not of that kind. Its estimate is elementary, but it concerns one explicit hypersurface in a power of the variety, and the passage from that estimate to the eigenvalues uses cohomology throughout.

Why the proof works. By cutting X into pieces and by induction on the dimension, part (a) comes down to one case (Section 8): U is smooth and affine of pure dimension d, the sheaf $\mathcal{F}$ is lisse on U, and the group is $H_c^d(U, \mathcal{F})$, the lowest one that can be nonzero. In higher degrees the weak Lefschetz theorem expresses the cohomology of U through that of a hyperplane section, of smaller dimension. In degree d it gives nothing. The trace formula alone does not help either: it gives the sum of the local traces of $\mathcal{F}$ over the points of U as an alternating sum over all degrees, in which large eigenvalues from different degrees could cancel.

The way out is to compare two counts. Let $k \geq 2$ and let Y be a hypersurface in the k-fold product U^k. Over a finite field with Q elements, let M be the sum of the local traces of $\mathcal{F}$ over the points of U, and N the sum of the local traces of $\mathcal{F} \boxtimes \cdots \boxtimes \mathcal{F}$ over the points of Y. A hypersurface in general position should carry one Q-th of the total, so we put $R = QN - M^k$ and call it the remainder. Suppose that Y has the weak Lefschetz property: restriction from U^k to Y is bijective in degrees below $kd-1$ and injective in degree $kd-1$. By duality, every compactly supported cohomology group of U^k above the lowest degree kd is then the image, bijectively or surjectively, of a group of Y. Written with the trace formula, R is an alternating sum over the groups of U^k and of Y, and along these maps the terms cancel in pairs. Two terms survive. One is the lowest group of U^k, which by the Künneth formula is the k-th tensor power of $H_c^d(U, \mathcal{F})$. The other is the kernel K of the map out of the lowest group of Y. The first sits in degree kd and is subtracted; the second sits in degree $kd-1$ and is added; so they enter R with the same sign (Proposition 4). Up to that one sign, R is therefore the trace of Frobenius on an actual vector space, the sum of the tensor power and of K, and not a difference of two traces. This matters because of an elementary fact about power sums (Lemma 2): if the traces of the powers of an automorphism grow at most like B^n, then every one of its eigenvalues has absolute value at most B. For a difference of two traces this is false, since eigenvalues can cancel.

It remains to find a hypersurface Y with the weak Lefschetz property for which R is small, and the form of its equation does this. Choose a finite map π from U to the affine space $\mathbf{A}^d$ and a nondegenerate bilinear form b on $\mathbf{A}^d$, and let the equation of Y be

$$\sum_{i=1}^{k-1} b\big(\pi(u_i), \pi(u_{i+1})\big) + \sum_{i=1}^{k} g_i(u_i) + c = 0,$$

where each g_i is a linear combination of coordinates of U. Consecutive factors are coupled through b, and there is no other coupling. Detecting the equation by additive characters writes QN as a sum over the characters $\psi(h\,\cdot)$ of the field. The term with $h = 0$ is exactly M^k, so R is the sum of the other terms. For $h \neq 0$ the term is a product of $k-1$ matrices with entries $\psi(h\,b(x,y))$, interlaced with diagonal matrices that carry the local traces. The first kind are $Q^{d/2}$ times unitary matrices, which is Parseval's identity for the Fourier transform on a finite vector space; the second kind are bounded by a constant times $Q^{w/2}$. Hence $|R|$ is at most a constant to the power k, times $Q^{k(w+d)/2}$, times $Q^{d/2+1}$ (Lemma 3 and Section 7). This is square-root size, up to the factor $Q^{d/2+1}$, which does not depend on k. That such a Y can be chosen with the weak Lefschetz property is not obvious, because the bilinear part of its

equation is fixed in advance. But Y is a hyperplane section of U^k for the embedding given by the coordinates together with the fixed bilinear expression, and a generic hyperplane whose coefficient on that expression is 1 will do (Lemma 8).

Putting the two halves together, every eigenvalue of Frobenius on the k-th tensor power of $H_c^d(U,\mathcal{F})$, in particular α^k for an eigenvalue α on $H_c^d(U,\mathcal{F})$, has absolute value at most $q^{k(w+d)/2+d/2+1}$. Taking k-th roots gives $|\alpha| \le q^{(w+d)/2} \cdot q^{(d/2+1)/k}$, and letting k grow removes the loss.

3 Power sums of an actual module

Let V be a finite-dimensional vector space over $\overline{\mathbf{Q}}_\ell$ with an automorphism Φ. Then $\mathrm{Tr}(\Phi^n \mid V) = \sum_j m_j \lambda_j^n$, where the λ_j are the distinct eigenvalues of Φ and the $m_j \ge 1$ their algebraic multiplicities. The next lemma is the form in which we use the positivity of the m_j. It allows the automorphism, the auxiliary space and the constants to change with k, which will be needed.

Lemma 2. *Let V be a finite-dimensional $\overline{\mathbf{Q}}_\ell$-vector space with an automorphism F, and let a be a real number. Suppose that for infinitely many integers $k \ge 1$ there are an integer $e_k \ge 1$, a finite-dimensional $\overline{\mathbf{Q}}_\ell$-vector space K_k with an automorphism Φ_k, a sign $\varepsilon_k = \pm 1$, real numbers $C_k > 0$ and $c_k \ge 0$, and complex numbers $R_{k,n}$ such that for all $n \ge 1$*

$$\varepsilon_k R_{k,n} = \iota\,\mathrm{Tr}\big(F^{e_k n} \mid V^{\otimes k}\big) + \iota\,\mathrm{Tr}\big(\Phi_k^n \mid K_k\big) \quad \text{and} \quad |R_{k,n}| \le C_k\, q^{e_k n(ka+c_k)},$$

and that c_k/k tends to 0 along these k. Then every eigenvalue α of F on V satisfies $|\iota\alpha| \le q^a$.

Proof. Fix such a k. Let $\mu_1, \dots, \mu_s$ be the images under ι of the distinct eigenvalues of the automorphism $F^{e_k} \oplus \Phi_k$ of $V^{\otimes k} \oplus K_k$, and $m_1, \dots, m_s \ge 1$ their multiplicities, so that $\varepsilon_k R_{k,n} = \sum_j m_j \mu_j^n$ for all $n \ge 1$. The μ_j are distinct, because ι is injective, and nonzero. If $v \ne 0$ and $Fv = \alpha v$, then $v \otimes \cdots \otimes v$ is an eigenvector of F^{e_k} on $V^{\otimes k}$ with eigenvalue α^{ke_k}, so $\iota(\alpha)^{ke_k}$ is one of the μ_j.

Put $\rho = q^{-e_k(ka+c_k)}$. By the bound on $R_{k,n}$, the power series $\sum_{n\ge1} \varepsilon_k R_{k,n} z^n$ converges in the disc $|z| < \rho$. Near $z = 0$ it is equal to the rational function $\sum_j m_j \mu_j z/(1-\mu_j z)$. Since the μ_j are distinct and no m_j is zero, this rational function has a pole at each of the points $1/\mu_j$. It agrees on the disc with a holomorphic function, so none of these points lies in the disc. Hence $|\mu_j| \le 1/\rho$ for every j, and in particular $|\iota\alpha|^{ke_k} \le q^{e_k(ka+c_k)}$, that is, $|\iota\alpha| \le q^{a+c_k/k}$. Now let k tend to infinity. □

4 A character sum along a path

The estimate of this note is the following. The sum runs over k points of a finite vector space; each point is weighted by an arbitrary bounded function, and the phase couples only consecutive points.

Lemma 3. *Let E be a finite field with Q elements, ψ a nontrivial additive character of E, $d \ge 1$, $A = E^d$ and $k \ge 2$. Let $b_1, \dots, b_{k-1}$ be nondegenerate bilinear forms on A with values in E, and let $t_1, \dots, t_k : A \to \mathbf{C}$ be functions with $|t_i(x)| \le T$ for all i and x. Then for every $h \in E^\times$ and $c \in E$*

$$\Big|\sum_{x\in A^k} t_1(x_1)\cdots t_k(x_k)\,\psi\Big(h\Big(\sum_{i=1}^{k-1} b_i(x_i, x_{i+1}) + c\Big)\Big)\Big| \le T^k Q^{d(k+1)/2}.$$

Proof. For $i < k$ let K_i be the matrix with rows and columns indexed by A and entries $K_i(x,y) = \psi(h\, b_i(x,y))$. The entry of $K_i K_i^*$ at (x,x') is $\sum_{y\in A} \psi(h\, b_i(x-x',y))$. If $x \neq x'$, the map $y \mapsto h\, b_i(x-x',y)$ is a nonzero linear form on A, because b_i is nondegenerate and $h \neq 0$, and the sum of a nontrivial character over A is zero. If $x = x'$ the sum is Q^d. So $K_i K_i^* = Q^d$, and K_i is $Q^{d/2}$ times a unitary matrix. Let D_i be the diagonal matrix with entries $t_i(x)$; its norm is at most T. Regard t_1 and t_k as vectors; their norms are at most $TQ^{d/2}$. The sum in the lemma equals $\psi(hc)\, t_1^{\mathsf{T}} K_1 D_2 K_2 \cdots D_{k-1} K_{k-1} t_k$, and by the Cauchy–Schwarz inequality its absolute value is at most

$$\|t_1\|\,\|t_k\| \prod_{i=1}^{k-1} \|K_i\| \prod_{i=2}^{k-1} \|D_i\| \leq (TQ^{d/2})^2\, Q^{d(k-1)/2}\, T^{k-2} = T^k Q^{d(k+1)/2}. \qquad \square$$

The trivial bound for the sum is $T^k Q^{dk}$. The lemma gains a factor $Q^{d/2}$ for each of the $k-1$ couplings.

5 The remainder of a point count

We recall what duality gives. Let Z be a smooth scheme of pure dimension m over a finite field $\mathbf{F}_r$ and $\mathcal{G}$ a lisse $\overline{\mathbf{Q}}_\ell$-sheaf on Z. Poincaré duality [11, XVIII, §3] is a perfect pairing, compatible with Frobenius, between $H_c^i(Z,\mathcal{G})$ and $H^{2m-i}(Z,\mathcal{G}^\vee)(m)$. Here (m) is the Tate twist; on a space twisted by (-1) the eigenvalues of Frobenius are multiplied by r. If Z is affine, then $H^i(Z,\mathcal{G}^\vee) = 0$ for $i > m$ by Artin's theorem [11, XIV, 3.1], and therefore $H_c^i(Z,\mathcal{G}) = 0$ for $i < m$. If $Y \subset Z$ is a smooth closed subscheme of pure dimension $m-1$, we use duality on Z and on Y: the transpose of the restriction map $H^a(Z,\mathcal{G}^\vee) \to H^a(Y,\mathcal{G}^\vee)$ is a map

$$H_c^j(Y,\mathcal{G})(-1) \longrightarrow H_c^{j+2}(Z,\mathcal{G}), \qquad j = 2m-2-a,$$

compatible with Frobenius. We call it the *Gysin map*. It is bijective when the restriction map is bijective, and surjective when the restriction map is injective.

Proposition 4. *Let U be a smooth affine scheme of pure dimension $d \geq 1$ over a finite field $\mathbf{F}_r$, and $\mathcal{F}$ a lisse $\overline{\mathbf{Q}}_\ell$-sheaf on U. Let $k \geq 1$, $X = U^k$, $D = kd$ and $\mathcal{F}_k = \mathcal{F} \boxtimes \cdots \boxtimes \mathcal{F}$ on X. Let $Y \subset X$ be a smooth closed subscheme of pure dimension $D-1$, possibly empty, defined over $\mathbf{F}_r$, such that over $\overline{\mathbf{F}}_r$ the restriction map*

$$H^a(X,\mathcal{F}_k^\vee) \longrightarrow H^a(Y,\mathcal{F}_k^\vee)$$

is an isomorphism for $a < D-1$ and injective for $a = D-1$. For $n \geq 1$ let M_n be the sum of the traces $\mathrm{Tr}(F_{u,n} \mid \mathcal{F}_{\bar{u}})$ over $u \in U(\mathbf{F}_{r^n})$, let N_n be the corresponding sum for $\mathcal{F}_k$ over $Y(\mathbf{F}_{r^n})$, and put $R_n = r^n N_n - M_n^k$. Let F be the geometric Frobenius of $\mathbf{F}_r$, and let K be the kernel of the Gysin map $H_c^{D-1}(Y,\mathcal{F}_k)(-1) \to H_c^{D+1}(X,\mathcal{F}_k)$. Then for all $n \geq 1$

$$(-1)^{D-1} R_n = \mathrm{Tr}\big(F^n \mid H_c^d(U,\mathcal{F})^{\otimes k}\big) + \mathrm{Tr}\big(F^n \mid K\big).$$

Proof. The schemes U, X and Y are smooth and affine, of pure dimensions d, D and $D-1$. By what was recalled, the groups $H_c^j(U,\mathcal{F})$ vanish for $j < d$, the groups $H_c^j(X,\mathcal{F}_k)$ for $j < D$, and the groups $H_c^j(Y,\mathcal{F}_k)$ for $j < D-1$. The hypothesis on restriction says that the Gysin map $H_c^j(Y,\mathcal{F}_k)(-1) \to H_c^{j+2}(X,\mathcal{F}_k)$ is an isomorphism for $j > D-1$ and surjective for $j = D-1$.

The stalk of $\mathcal{F}_k$ at a point $(u_1,\ldots,u_k)$ is the tensor product of the stalks of $\mathcal{F}$ at the u_i, so the sum of the traces of $\mathcal{F}_k$ over $X(\mathbf{F}_{r^n}) = U(\mathbf{F}_{r^n})^k$ is M_n^k. By the trace formula [12, Rapport, Théorème 3.2], applied on X and on Y,

$$M_n^k = \sum_{b\geq 0} (-1)^{D+b}\, \mathrm{Tr}\big(F^n \mid H_c^{D+b}(X,\mathcal{F}_k)\big),$$

$$r^n N_n = \sum_{a\geq 0} (-1)^{D-1+a}\, \mathrm{Tr}\big(F^n \mid H_c^{D-1+a}(Y,\mathcal{F}_k)(-1)\big).$$

Therefore, in $(-1)^{D-1}R_n$, the trace on $H_c^{D-1+a}(Y,\mathcal{F}_k)(-1)$ enters with the sign $(-1)^a$ and the trace on $H_c^{D+b}(X,\mathcal{F}_k)$ with the sign $(-1)^b$. For $a \geq 1$ the Gysin map identifies $H_c^{D-1+a}(Y,\mathcal{F}_k)(-1)$ with $H_c^{D+1+a}(X,\mathcal{F}_k)$, which enters with the opposite sign, and the two traces cancel. Three terms remain: the trace on $H_c^{D-1}(Y,\mathcal{F}_k)(-1)$, minus the trace on $H_c^{D+1}(X,\mathcal{F}_k)$, plus the trace on $H_c^{D}(X,\mathcal{F}_k)$. The second space is the quotient of the first by K, so the first two terms together are the trace on K. Finally, by the Künneth formula [11, XVII, 5.4.3] the group $H_c^D(X,\mathcal{F}_k)$ is the sum of the tensor products $H_c^{j_1}(U,\mathcal{F}) \otimes \cdots \otimes H_c^{j_k}(U,\mathcal{F})$ with $j_1 + \cdots + j_k = D$. Since the factors vanish for $j_i < d$, only $j_1 = \cdots = j_k = d$ contributes, and $H_c^D(X,\mathcal{F}_k) = H_c^d(U,\mathcal{F})^{\otimes k}$. □

Nothing is assumed about the higher cohomology groups of U or of Y: they cancel through actual isomorphisms, and the surjection at the border leaves an actual kernel.

Example 5. Let $U = \mathbf{A}^1$, $\mathcal{F} = \overline{\mathbf{Q}}_\ell$ and $k = 2$, and let $Y \subset \mathbf{A}^2$ be the conic $xy + ax + by + c = 0$ with $ab \neq c$. It is isomorphic to the multiplicative group, so $N_n = r^n - 1$, while $M_n = r^n$; hence $R_n = -r^n$ and $(-1)^{D-1}R_n = r^n$. On the other side $H_c^1(\mathbf{A}^1, \overline{\mathbf{Q}}_\ell) = 0$, and $K = H_c^1(Y, \overline{\mathbf{Q}}_\ell)(-1)$ has dimension one, with Frobenius acting by r. The group $H_c^2(Y, \overline{\mathbf{Q}}_\ell)(-1)$ is identified with $H_c^4(\mathbf{A}^2, \overline{\mathbf{Q}}_\ell)$ and has cancelled.

Remark 6. The space K is the analogue, for compact supports on an affine variety, of the part of the cohomology of a hyperplane section that the Gysin map kills. In Grothendieck's notes on the standard conjectures in the Montpellier archive this part is his *partie évanescente* E^{n-1} (cote 16, scan 40, in the transcription [14]). The comparison plays no role in the proof.

6 Hypersurfaces with the weak Lefschetz property

Theorem 7 (Deligne; [6, Appendix, A.5, with A.3 and A.4])**.** *Let Z be a separated scheme of finite type over an algebraically closed field of characteristic different from ℓ, let $\varphi : Z \to \mathbf{P}^N$ be a quasi-finite morphism, and let $\mathcal{K}$ be a perverse sheaf on Z. Then there is a dense open set of hyperplanes $H \subset \mathbf{P}^N$ for which the restriction map*

$$H^a(Z,\mathcal{K}) \longrightarrow H^a(\varphi^{-1}H,\mathcal{K})$$

is an isomorphism for $a < -1$ and injective for $a = -1$.

Katz attributes this proof to Deligne [6, p. 152]. It runs as follows. Let $\alpha : Z \smallsetminus \varphi^{-1}H \to Z$ be the inclusion. The assertion is equivalent to the vanishing of $H^a(Z, \alpha_!\alpha^*\mathcal{K})$ for $a < 0$. For generic H, the generic base change theorem [12, Théorèmes de finitude, 1.9], applied over the space of hyperplanes, identifies these groups with the compactly supported cohomology of the affine scheme $\mathbf{P}^N \smallsetminus H$ with coefficients in a direct image of $\alpha^*\mathcal{K}$; duality and Artin's theorem then give the vanishing. No weights occur. If Z is smooth of pure dimension D and $\mathcal{G}$ is a lisse sheaf on Z, then $\mathcal{G}[D]$ is perverse, and the conclusion reads: the restriction map $H^a(Z,\mathcal{G}) \to H^a(\varphi^{-1}H,\mathcal{G})$ is an isomorphism for $a < D-1$ and injective for $a = D-1$.

We need this for hypersurfaces whose equation has a fixed part. Corollary 3.4.1 of [6] states exactly that form, but it is deduced there from the main theorem of that article, whose proof is by the Fourier transform. The next lemma deduces it from Theorem 7 instead, so that only the inputs named above are used; it also gives the smoothness of the hypersurface.

Lemma 8. *Let V be a smooth affine scheme of pure dimension $D \geq 1$ over $\overline{\mathbf{F}}_q$, let $y_1, \ldots, y_M$ be functions on V such that $\Gamma = (y_1, \ldots, y_M) : V \to \mathbf{A}^M$ is a closed immersion, and let $\mathcal{G}$ be a lisse $\overline{\mathbf{Q}}_\ell$-sheaf on V. Consider the hypersurfaces*

$$Y_{\lambda,\mu} = V \cap \Big(\sum_{m=1}^{M} \lambda_m y_m + \mu = 0 \Big), \qquad \lambda \neq 0.$$

For (λ, μ) in a dense open subset of the space of all parameters, and also for (λ, μ) in a dense open subset of the space of parameters with $\lambda_M = 1$, the scheme $Y_{\lambda,\mu}$ is smooth, of pure dimension $D-1$ or empty on each connected component of V, and the restriction map $H^a(V, \mathcal{G}) \to H^a(Y_{\lambda,\mu}, \mathcal{G})$ is an isomorphism for $a < D-1$ and injective for $a = D-1$.

Proof. Regard $\mathbf{A}^M$ as the complement of the hyperplane at infinity in $\mathbf{P}^M$. The hyperplanes of $\mathbf{P}^M$ other than the one at infinity are the zero sets of the equations $\sum_m \lambda_m y_m + \mu = 0$ with $\lambda \neq 0$, and the inverse image of such a hyperplane under Γ is $Y_{\lambda,\mu}$.

The restriction property holds for a dense open set of hyperplanes by Theorem 7, applied to $\varphi = \Gamma$, which is quasi-finite, and to the perverse sheaf $\mathcal{G}[D]$.

Smoothness holds for a dense open set of hyperplanes by a count of dimensions. Consider the pairs (v, H) of a point v of V and a hyperplane H that contains the tangent space of $\Gamma(V)$ at $\Gamma(v)$. Since Γ is an immersion, this tangent space has dimension D, and the hyperplanes containing it form a projective space of dimension $M-D-1$. So the set of such pairs has dimension at most $D + (M-D-1) = M-1$, and its image in the space of hyperplanes, which has dimension M, is not dense. If H lies outside the closure of this image, the equation of H restricts to a nonzero linear form on the tangent space of V at each point of $\Gamma^{-1}H$, so $\Gamma^{-1}H$ is smooth, and of pure dimension $D-1$ or empty on each component of V. This argument is valid in every characteristic, because Γ is an immersion; compare [5, I.6.10].

The hyperplanes with both properties form a dense open subset of the space of all hyperplanes. It meets the open subset $\lambda_M \neq 0$ in a dense open subset, and on that subset we may scale the equation so that $\lambda_M = 1$. □

7 The lowest degree of a smooth affine variety

Theorem 9. *Let U be a smooth affine scheme of pure dimension $d \geq 1$ over $\mathbf{F}_q$ and $\mathcal{F}$ a lisse $\overline{\mathbf{Q}}_\ell$-sheaf on U. Suppose that there are a constant $C > 0$ and a real number w with*

$$|\iota \operatorname{Tr}(F_{u,n} \mid \mathcal{F}_{\bar{u}})| \leq C\, q^{nw/2} \qquad \text{for all } n \geq 1 \text{ and } u \in U(\mathbf{F}_{q^n}).$$

Then every eigenvalue of F on $H^d_c(U, \mathcal{F})$ has absolute value at most $q^{(w+d)/2}$.

Proof. Choices. By Noether normalization [13, Tag 00OY] there is a finite morphism $\pi : U \to \mathbf{A}^d$ over $\mathbf{F}_q$. Let B be the number of elements in a set of generators of the coordinate ring $\mathcal{O}(U)$ as a module over that of $\mathbf{A}^d$. Over any field, each fibre of π is the spectrum of an algebra of dimension at most B, so it has at most B rational points. Choose generators $z_1, \ldots, z_s$ of $\mathcal{O}(U)$ as an $\mathbf{F}_q$-algebra, so that $z = (z_1, \ldots, z_s) : U \to \mathbf{A}^s$ is a closed immersion, and a nondegenerate bilinear form b on $\mathbf{A}^d$ defined over $\mathbf{F}_q$. Fix $k \geq 2$ and put $D = kd$.

The hypersurface. For parameters $a = (a_{ij})$, with $1 \leq i \leq k$ and $1 \leq j \leq s$, and c, let $P_{a,c}$ be the function on U^k

$$P_{a,c}(u_1, \ldots, u_k) = \sum_{i=1}^{k-1} b(\pi(u_i), \pi(u_{i+1})) + \sum_{i=1}^{k} \sum_{j=1}^{s} a_{ij}\, z_j(u_i) + c,$$

and let $Y_{a,c} \subset U^k$ be its zero set. Over $\overline{\mathbf{F}}_q$ the map

$$\Gamma : \; u \longmapsto \Big(z(u_1), \ldots, z(u_k),\; \sum_{i=1}^{k-1} b\big(\pi(u_i), \pi(u_{i+1})\big) \Big)$$

is a closed immersion of U^k into $\mathbf{A}^{sk+1}$, because its first sk coordinates already are one. The $Y_{a,c}$ are exactly the hypersurfaces of Lemma 8 for Γ whose coefficient on the last coordinate is 1. By that lemma, applied to the sheaf $\mathcal{F}_k^\vee$, there is a dense open set of parameters (a, c)

for which $Y_{a,c}$ is smooth of pure dimension $D-1$, or empty, and has the restriction property required in Proposition 4. A dense open subset of an affine space over $\overline{\mathbf{F}}_q$ has a point with coordinates in a finite field. Fix such a point (a,c), with coordinates in $\mathbf{F}_r$, $r=q^e$, and write $P=P_{a,c}$ and $Y=Y_{a,c}$. The exponent e and the point (a,c) depend on k; this is allowed in Lemma 2.

The remainder as a character sum. Let $n\ge 1$, $Q=r^n$ and $E=\mathbf{F}_Q$, and let M_n, N_n and $R_n=QN_n-M_n^k$ be formed over $\mathbf{F}_r$ as in Proposition 4. Let ψ be a nontrivial additive character of E, and for $u\in U(E)$ put $\tau(u)=\iota\,\mathrm{Tr}(F_{u,en}\mid \mathcal{F}_{\bar u})$; by hypothesis $|\tau(u)|\le CQ^{w/2}$. For a point $u=(u_1,\dots,u_k)$ of U^k with values in E, the sum $\sum_{h\in E}\psi(hP(u))$ is Q if $P(u)=0$ and 0 otherwise. Therefore

$$Q\cdot\iota N_n=\sum_{h\in E}\ \sum_{u\in U(E)^k}\tau(u_1)\cdots\tau(u_k)\,\psi\big(hP(u)\big).$$

In the inner sum we collect the u_i according to their images $x_i=\pi(u_i)$ in E^d. With

$$t_{i,h}(x)=\sum_{u\in U(E),\ \pi(u)=x}\tau(u)\,\psi\Big(h\sum_{j=1}^{s}a_{ij}z_j(u)\Big),$$

the inner sum becomes

$$\psi(hc)\sum_{x\in(E^d)^k}t_{1,h}(x_1)\cdots t_{k,h}(x_k)\,\psi\Big(h\sum_{i=1}^{k-1}b(x_i,x_{i+1})\Big).$$

For $h=0$ all phases are 1, and the inner sum is $\big(\sum_{u\in U(E)}\tau(u)\big)^k=(\iota M_n)^k$. For $h\ne 0$, each $t_{i,h}(x)$ is a sum of at most B terms of absolute value at most $CQ^{w/2}$, so Lemma 3 applies with $T=BCQ^{w/2}$ and all b_i equal to b. There are $Q-1$ such h, and we obtain

$$|\iota R_n|=\big|Q\cdot\iota N_n-(\iota M_n)^k\big|\le(Q-1)\,(BC)^kQ^{kw/2}\,Q^{d(k+1)/2}\le(BC)^k\,q^{en\{k(w+d)/2+d/2+1\}}.$$

Conclusion. The geometric Frobenius of $\mathbf{F}_r$ is F^e. By Proposition 4, applied over $\mathbf{F}_r$,

$$(-1)^{D-1}\,\iota R_n=\iota\,\mathrm{Tr}\big(F^{en}\mid H_c^d(U,\mathcal{F})^{\otimes k}\big)+\iota\,\mathrm{Tr}\big((F^e)^n\mid K\big)$$

for all $n\ge1$. These are the hypotheses of Lemma 2 for $V=H_c^d(U,\mathcal{F})$ and $a=(w+d)/2$, with $e_k=e$, $K_k=K$, $\varepsilon_k=(-1)^{D-1}$, $C_k=(BC)^k$ and $c_k=d/2+1$, for every $k\ge2$. The lemma gives the theorem. □

8 Proof of Theorem 1

A bound for the eigenvalues may be proved after replacing $\mathbf{F}_q$ by a finite extension $\mathbf{F}_{q^e}$: the eigenvalues of F^e are the e-th powers of those of F, so a bound $q^{e(w+i)/2}$ for the former gives the bound $q^{(w+i)/2}$ for the latter.

Part (a), for X smooth and affine of pure dimension d and $\mathcal{F}$ lisse. We prove the statement by induction on d, for all finite fields at once. Let $d=0$. Then $H_c^0(X,\mathcal{F})$ is the sum of the stalks of $\mathcal{F}$ at the geometric points of X. If x is a closed point of degree m, the Frobenius F permutes the m geometric points above x cyclically, and F^m acts on each of their stalks as the local Frobenius at x. So for an eigenvalue α of F on the part of H_c^0 above x, α^m is an eigenvalue of the local Frobenius, and $|\alpha|^m\le q^{mw/2}$.

Let $d\ge1$. For $i<d$ the group $H_c^i(X,\mathcal{F})$ vanishes (Section 5). For $i=d$ we apply Theorem 9, with C the rank of $\mathcal{F}$: if u is a point with values in $\mathbf{F}_{q^n}$ over a closed point x of

degree m, then $F_{u,n}$ is the (n/m)-th power of the local Frobenius at x, so its eigenvalues have absolute value at most $q^{nw/2}$. For $i > d$ we cut by a hyperplane. Choose a closed immersion of X into an affine space. By Lemma 8, applied to this immersion, to all hyperplanes and to the sheaf $\mathcal{F}^\vee$, there is, over some finite extension $\mathbf{F}_{q^e}$, a hyperplane section Y of X that is smooth and affine of pure dimension $d-1$, or empty, and for which restriction from X to Y on the cohomology of $\mathcal{F}^\vee$ is bijective below degree $d-1$ and injective in degree $d-1$. Dually (Section 5) the Gysin map

$$H_c^{i-2}(Y, \mathcal{F})(-1) \longrightarrow H_c^i(X, \mathcal{F})$$

is surjective for $i = d+1$ and bijective for $i > d+1$. By the induction hypothesis for Y over $\mathbf{F}_{q^e}$, the eigenvalues of F^e on $H_c^{i-2}(Y, \mathcal{F})$ have absolute value at most $q^{e(w+i-2)/2}$. The twist multiplies them by q^e. So the eigenvalues of F^e on $H_c^i(X, \mathcal{F})$ have absolute value at most $q^{e(w+i)/2}$, which is the bound.

Part (a) in general. The cohomology of X is that of the reduced scheme, so we may assume X reduced. We use induction on the dimension of X; in dimension 0 the scheme X is smooth and affine and $\mathcal{F}$ is lisse, which is the case already treated. Since $\mathbf{F}_q$ is perfect, X has a dense open subscheme U that is smooth and affine, of pure dimension on each of its connected components, and on which $\mathcal{F}$ is lisse. Its complement Z has smaller dimension, and the restriction of $\mathcal{F}$ to Z again has weights at most w. The sequence

$$H_c^i(U, \mathcal{F}) \longrightarrow H_c^i(X, \mathcal{F}) \longrightarrow H_c^i(Z, \mathcal{F})$$

is exact and compatible with Frobenius; it is part of the long exact sequence of $0 \to j_! j^* \mathcal{F} \to \mathcal{F} \to i_* i^* \mathcal{F} \to 0$, where j and i are the inclusions of U and Z. So every eigenvalue of F on $H_c^i(X, \mathcal{F})$ is an eigenvalue on $H_c^i(U, \mathcal{F})$ or on $H_c^i(Z, \mathcal{F})$. The first are bounded by the case already treated, applied to the components of U, and the second by the induction hypothesis.

Part (b). Treating the connected components of X separately, we may assume that X has pure dimension d. Since X is proper, $H^i(X, \mathcal{F}) = H_c^i(X, \mathcal{F})$, and part (a) gives the upper bound. The dual sheaf $\mathcal{F}^\vee$ is lisse and pure of weight $-w$, because the eigenvalues of a local Frobenius on its stalk are the inverses of those on the stalk of $\mathcal{F}$; here we use that $\mathcal{F}$ is pure at every point, and not only bounded. Poincaré duality is a perfect pairing of $H^i(X, \mathcal{F})$ with $H^{2d-i}(X, \mathcal{F}^\vee)$ into $\overline{\mathbf{Q}}_\ell(-d)$, compatible with Frobenius. So if α is an eigenvalue of F on $H^i(X, \mathcal{F})$, then q^d/α is an eigenvalue of F on $H^{2d-i}(X, \mathcal{F}^\vee)$. Part (a), applied to $\mathcal{F}^\vee$, gives $|q^d/\alpha| \leq q^{(-w+2d-i)/2}$, that is, $|\alpha| \geq q^{(w+i)/2}$. □

Provenance and checks

The proof is Part I, Section 1 of the output of a research run of GPT-6 Astra (ChatGPT Pro mode) on 21 November 2026, whose task was to work out footnote 3 of [1]. The same day Claude Opus 5.5 at effort high checked every step of it. Claude Fable 5.2 at effort max, without seeing that check, re-proved every step and found no gap; it read the pages of [6] used here (pp. 152–158, 179–181 and 219–222) and re-proved Theorem 7 independently. One correction came from that second check and is adopted here: the weak Lefschetz input is taken from Deligne's proof in Katz's appendix, through Lemma 8, and not from the printed proof of Katz's Corollary 3.4.1.

The other references were not reopened for this text. The foundational ones are cited with the locators of the run. What the first two sections say about [2], [3], [4], [7] and [9] is from memory. The two statements of [8] and the introduction of [10] were read at the source by Claude Opus 5.5 at effort max. No human has checked the proof or this text.

The prompt, the output of the run and the checks are public at

https://github.com/KokunoYumeto/grothendieck-montpellier-archive/tree/main/research/grh_20261120/deligne_sign_only.

This text was written by Claude (Claude Fable 5.2 at effort max), using an earlier text by Claude Opus 5.5 at effort max. It is dedicated to the public domain (CC0 1.0).

References

[1] L. Alpöge, note on the Riemann hypothesis, 20 November 2026, https://alpo.ge/remark.pdf.

[2] E. Bombieri, Counting points on curves over finite fields (d'après S. A. Stepanov), Séminaire Bourbaki 1972/73, Exposé 430, Lecture Notes in Math. 383, Springer, 1974, 234–241.

[3] P. Deligne, La conjecture de Weil. I, Publ. Math. IHÉS 43 (1974), 273–307.

[4] P. Deligne, La conjecture de Weil. II, Publ. Math. IHÉS 52 (1980), 137–252.

[5] J.-P. Jouanolou, Théorèmes de Bertini et applications, Birkhäuser, 1983.

[6] N. M. Katz, Affine cohomological transforms, perversity, and monodromy, J. Amer. Math. Soc. 6 (1993), 149–222.

[7] K. S. Kedlaya, Fourier transforms and p-adic "Weil II", Compos. Math. 142 (2006), 1426–1450.

[8] L. Lafforgue, Chtoucas de Drinfeld et correspondance de Langlands, Invent. Math. 147 (2002), 1–241.

[9] G. Laumon, Transformation de Fourier, constantes d'équations fonctionnelles et conjecture de Weil, Publ. Math. IHÉS 65 (1987), 131–210.

[10] A. J. Scholl, Hypersurfaces and the Weil conjectures, arXiv:0912.2073.

[11] M. Artin, A. Grothendieck, J.-L. Verdier, Théorie des topos et cohomologie étale des schémas (SGA 4), Lecture Notes in Math. 269, 270, 305, Springer, 1972–1973.

[12] P. Deligne, Cohomologie étale (SGA $4\frac{1}{2}$), Lecture Notes in Math. 569, Springer, 1977.

[13] The Stacks Project, https://stacks.math.columbia.edu/tag/00OY.

[14] Transcription of A. Grothendieck's papers in the Montpellier archive, https://github.com/KokunoYumeto/grothendieck-montpellier-archive.