L-functions via deformations: from hyperelliptic curves to hypergeometric motives

Kiran S. Kedlaya

Department of Mathematics, University of California, San Diego
kedlaya@ucsd.edu
http://kskedlaya.org/slides/

Arithmetic of Hyperelliptic Curves
Abdus Salam International Centre for Theoretical Physics (ICTP)
September 8, 2017

Contents

Zeta functions of algebraic varieties

For X an algebraic variety over a finite field $\mathbb{F}_q$, the *zeta function*

$$\zeta(X,T) = \prod_{x \in X^\circ} (1 - T^{[\kappa(x):\mathbb{F}_q]})^{-1} = \exp\left(\sum_{n=1}^\infty \#X(\mathbb{F}_{q^n})\frac{T^n}{n}\right) \in \mathbb{Z}[\![T]\!]$$

is a rational function of T. That is because it is possible to a *spectral interpretation* of $\zeta(X,T)$ consisting of a field K of characteristic 0; finite-dimensional K-vector spaces V_i for $i = 0, 1, \dots, 2\dim(X)$; and K-linear endomorphisms F_i on V_i satisfying the *Lefschetz trace formula*:

$$\#X(\mathbb{F}_{q^n}) = \sum_{i=0}^{2\dim(X)} (-1)^i \operatorname{trace}(F_i^n, V_i) \qquad (n = 1, 2, \dots).$$

This then implies that

$$\zeta(X,T) = \prod_{i=0}^{2\dim(X)} \det(1 - F_i T, V_i)^{(-1)^{i+1}}.$$

Zeta functions of algebraic varieties

For X an algebraic variety over a finite field $\mathbb{F}_q$, the *zeta function*

$$\zeta(X, T) = \prod_{x \in X^\circ} (1 - T^{[\kappa(x):\mathbb{F}_q]})^{-1} = \exp\left(\sum_{n=1}^\infty \#X(\mathbb{F}_{q^n}) \frac{T^n}{n}\right) \in \mathbb{Z}[\![T]\!]$$

is a rational function of T. That is because it is possible to a *spectral interpretation* of $\zeta(X, T)$ consisting of a field K of characteristic 0; finite-dimensional K-vector spaces V_i for $i = 0, 1, \dots, 2\dim(X)$; and K-linear endomorphisms F_i on V_i satisfying the *Lefschetz trace formula*:

$$\#X(\mathbb{F}_{q^n}) = \sum_{i=0}^{2\dim(X)} (-1)^i \operatorname{trace}(F_i^n, V_i) \qquad (n = 1, 2, \dots).$$

This then implies that

$$\zeta(X, T) = \prod_{i=0}^{2\dim(X)} \det(1 - F_i T, V_i)^{(-1)^{i+1}}.$$

Weil cohomology: ℓ-adic versus p-adic

Such data are provided in a systematic way by *Weil cohomology* constructions, of which there are two general types.

- Grothendieck's formalism of *étale cohomology* produces one Weil cohomology theory with coefficients in $\mathbb{Q}_\ell$ for each prime ℓ other than p, the characteristic of $\mathbb{F}_q$. This theory is quite rich, and has been the basis for most new developments on geometric zeta functions.
- Building on Dwork's original proof of rationality (predating étale cohomology!), Berthelot introduced *rigid cohomology* with coefficients in a finite extension[1] of $\mathbb{Q}_p$. (This relates explicitly to *crystalline cohomology* for smooth proper varieties or *Monsky-Washnitzer cohomology* for smooth affine varieties.) Recently, most formalism of étale cohomology has been replicated for rigid cohomology.

[1] The residue field must contain $\mathbb{F}_q$.

Weil cohomology: ℓ-adic versus p-adic

Such data are provided in a systematic way by *Weil cohomology* constructions, of which there are two general types.

- Grothendieck's formalism of *étale cohomology* produces one Weil cohomology theory with coefficients in $\mathbb{Q}_\ell$ for each prime ℓ other than p, the characteristic of $\mathbb{F}_q$. This theory is quite rich, and has been the basis for most new developments on geometric zeta functions.
- Building on Dwork's original proof of rationality (predating étale cohomology!), Berthelot introduced *rigid cohomology* with coefficients in a finite extension[1] of $\mathbb{Q}_p$. (This relates explicitly to *crystalline cohomology* for smooth proper varieties or *Monsky-Washnitzer cohomology* for smooth affine varieties.) Recently, most formalism of étale cohomology has been replicated for rigid cohomology.

[1] The residue field must contain $\mathbb{F}_q$.

Weil cohomology: ℓ-adic versus p-adic

Such data are provided in a systematic way by *Weil cohomology* constructions, of which there are two general types.

- Grothendieck's formalism of *étale cohomology* produces one Weil cohomology theory with coefficients in $\mathbb{Q}_\ell$ for each prime ℓ other than p, the characteristic of $\mathbb{F}_q$. This theory is quite rich, and has been the basis for most new developments on geometric zeta functions.
- Building on Dwork's original proof of rationality (predating étale cohomology!), Berthelot introduced *rigid cohomology* with coefficients in a finite extension[1] of $\mathbb{Q}_p$. (This relates explicitly to *crystalline cohomology* for smooth proper varieties or *Monsky-Washnitzer cohomology* for smooth affine varieties.) Recently, most formalism of étale cohomology has been replicated for rigid cohomology.

[1] The residue field must contain $\mathbb{F}_q$.

Factorization of zeta functions and varieties

Suppose X is smooth proper over $\mathbb{F}_q$. By Deligne's analogue of the Riemann hypothesis, there exists a unique factorization

$$\zeta(X, T) = \prod_{i=0}^{2\dim(X)} P_i(T)^{(-1)^{i+1}}$$

in which $P_i(T) \in 1 + T\mathbb{Z}[T]$ has all $\mathbb{C}$-roots of absolute value $q^{-i/2}$.

More precisely, Deligne (1974, 1980) showed that for the data F_i, V_i arising from ℓ-adic étale cohomology, the polynomial $P_i(T) = \det(1 - F_i T, V_i)$ has all $\mathbb{C}$-roots of absolute value $q^{-i/2}$. A variant of the second proof can be executed with rigid cohomology (K, 2006).

There is a formal process for "factoring" X into pieces that account for the individual P_i; this is the theory of *motives*.

Factorization of zeta functions and varieties

Suppose X is smooth proper over $\mathbb{F}_q$. By Deligne's analogue of the Riemann hypothesis, there exists a unique factorization

$$\zeta(X,T) = \prod_{i=0}^{2\dim(X)} P_i(T)^{(-1)^{i+1}}$$

in which $P_i(T) \in 1 + T\mathbb{Z}[T]$ has all $\mathbb{C}$-roots of absolute value $q^{-i/2}$.

More precisely, Deligne (1974, 1980) showed that for the data F_i, V_i arising from ℓ-adic étale cohomology, the polynomial $P_i(T) = \det(1 - F_i T, V_i)$ has all $\mathbb{C}$-roots of absolute value $q^{-i/2}$. A variant of the second proof can be executed with rigid cohomology (K, 2006).

There is a formal process for "factoring" X into pieces that account for the individual P_i; this is the theory of *motives*.

Factorization of zeta functions and varieties

Suppose X is smooth proper over $\mathbb{F}_q$. By Deligne's analogue of the Riemann hypothesis, there exists a unique factorization

$$\zeta(X, T) = \prod_{i=0}^{2\dim(X)} P_i(T)^{(-1)^{i+1}}$$

in which $P_i(T) \in 1 + T\mathbb{Z}[T]$ has all $\mathbb{C}$-roots of absolute value $q^{-i/2}$.

More precisely, Deligne (1974, 1980) showed that for the data F_i, V_i arising from ℓ-adic étale cohomology, the polynomial $P_i(T) = \det(1 - F_i T, V_i)$ has all $\mathbb{C}$-roots of absolute value $q^{-i/2}$. A variant of the second proof can be executed with rigid cohomology (K, 2006).

There is a formal process for "factoring" X into pieces that account for the individual P_i; this is the theory of *motives*.

Zeta functions and L-functions

Suppose now that X is a smooth proper variety over a number field K. Then for the motive of weight i associated to X, one gets an L-function by taking an Euler product $\prod_{\mathfrak{p}} L_{\mathfrak{p}}(s)$ in which for almost all prime ideals $\mathfrak{p}$ of $\mathfrak{o}_K$, we have $L_{\mathfrak{p}}(s) = P_i(\mathrm{Norm}(\mathfrak{p})^{-s})$ where P_i is the corresponding factor of the zeta function of the reduction of (an integral model of) X modulo $\mathfrak{p}$.

This may be familiar for $X = E$ an elliptic curve. Over $\mathbb{F}_q$, we have

$$P_0(T) = 1 - T, \qquad P_1(T) = 1 - a_E T + qT^2, \qquad P_2(T) = 1 - qT.$$

Over K, for $i = 0, 1, 2$, the resulting L-functions are

$$\zeta_K(s), \qquad L(E, s), \qquad \zeta_K(s-1)$$

where $L(E, s)$ is (almost) $\prod_{\mathfrak{p}}(1 - a_{E,\mathfrak{p}} q^{-s} + q^{1-2s})^{-1}$ for $q = \mathrm{Norm}(\mathfrak{p})$. Similar considerations apply when X is a hyperelliptic (or arbitrary) curve.

Zeta functions and L-functions

Suppose now that X is a smooth proper variety over a number field K. Then for the motive of weight i associated to X, one gets an L-function by taking an Euler product $\prod_{\mathfrak{p}} L_{\mathfrak{p}}(s)$ in which for almost all prime ideals $\mathfrak{p}$ of $\mathfrak{o}_K$, we have $L_{\mathfrak{p}}(s) = P_i(\mathrm{Norm}(\mathfrak{p})^{-s})$ where P_i is the corresponding factor of the zeta function of the reduction of (an integral model of) X modulo $\mathfrak{p}$.

This may be familiar for $X = E$ an elliptic curve. Over $\mathbb{F}_q$, we have

$$P_0(T) = 1 - T, \qquad P_1(T) = 1 - a_E T + qT^2, \qquad P_2(T) = 1 - qT.$$

Over K, for $i = 0, 1, 2$, the resulting L-functions are

$$\zeta_K(s), \qquad L(E, s), \qquad \zeta_K(s-1)$$

where $L(E, s)$ is (almost) $\prod_{\mathfrak{p}}(1 - a_{E,\mathfrak{p}} q^{-s} + q^{1-2s})^{-1}$ for $q = \mathrm{Norm}(\mathfrak{p})$. Similar considerations apply when X is a hyperelliptic (or arbitrary) curve.

Zeta functions and L-functions

Suppose now that X is a smooth proper variety over a number field K. Then for the motive of weight i associated to X, one gets an L-function by taking an Euler product $\prod_{\mathfrak{p}} L_{\mathfrak{p}}(s)$ in which for almost all prime ideals $\mathfrak{p}$ of $\mathfrak{o}_K$, we have $L_{\mathfrak{p}}(s) = P_i(\mathrm{Norm}(\mathfrak{p})^{-s})$ where P_i is the corresponding factor of the zeta function of the reduction of (an integral model of) X modulo $\mathfrak{p}$.

This may be familiar for $X = E$ an elliptic curve. Over $\mathbb{F}_q$, we have

$$P_0(T) = 1 - T, \qquad P_1(T) = 1 - a_E T + qT^2, \qquad P_2(T) = 1 - qT.$$

Over K, for $i = 0, 1, 2$, the resulting L-functions are

$$\zeta_K(s), \qquad L(E, s), \qquad \zeta_K(s-1)$$

where $L(E, s)$ is (almost) $\prod_{\mathfrak{p}}(1 - a_{E,\mathfrak{p}} q^{-s} + q^{1-2s})^{-1}$ for $q = \mathrm{Norm}(\mathfrak{p})$. Similar considerations apply when X is a hyperelliptic (or arbitrary) curve.

Computational aspects of Weil cohomology

With a few exceptions[2], the only methods we know for computing $\zeta(X, T)$ are to explicitly compute the matrices via which F_i act on some basis of V_i, for some choice of Weil cohomology.

The definition of étale cohomology, which quantifies over all covers in the étale topology, is hard to make computationally effective. This can be done for curves of low genus (using the Jacobian as in Schoof's method) and for motives attached to modular forms (Edixhoven et al.).

By contrast, rigid cohomology can be defined[3] more concretely in terms of differential forms on certain p-adic rigid analytic spaces. Correspondingly, it tends to be a better source for algorithms.

[2] One exception is when one can actually count $\#X(\mathbb{F}_{q^n})$ for enough n to pin down the rational function. Another is for curves of low genus over not-too-large fields, where one can diagnose the order of the class group using baby step-giant step.

[3] More precisely, to put rigid cohomology on a sound footing it should *also* be defined in the language of sites (Le Stum), then compared to more concrete constructions.

Computational aspects of Weil cohomology

With a few exceptions[2], the only methods we know for computing $\zeta(X, T)$ are to explicitly compute the matrices via which F_i act on some basis of V_i, for some choice of Weil cohomology.

The definition of étale cohomology, which quantifies over all covers in the étale topology, is hard to make computationally effective. This can be done for curves of low genus (using the Jacobian as in Schoof's method) and for motives attached to modular forms (Edixhoven et al.).

By contrast, rigid cohomology can be defined[3] more concretely in terms of differential forms on certain p-adic rigid analytic spaces. Correspondingly, it tends to be a better source for algorithms.

[2] One exception is when one can actually count $\#X(\mathbb{F}_{q^n})$ for enough n to pin down the rational function. Another is for curves of low genus over not-too-large fields, where one can diagnose the order of the class group using baby step-giant step.

[3] More precisely, to put rigid cohomology on a sound footing it should *also* be defined in the language of sites (Le Stum), then compared to more concrete constructions.

Computational aspects of Weil cohomology

With a few exceptions[2], the only methods we know for computing $\zeta(X, T)$ are to explicitly compute the matrices via which F_i act on some basis of V_i, for some choice of Weil cohomology.

The definition of étale cohomology, which quantifies over all covers in the étale topology, is hard to make computationally effective. This can be done for curves of low genus (using the Jacobian as in Schoof's method) and for motives attached to modular forms (Edixhoven et al.).

By contrast, rigid cohomology can be defined[3] more concretely in terms of differential forms on certain p-adic rigid analytic spaces. Correspondingly, it tends to be a better source for algorithms.

[2] One exception is when one can actually count $\#X(\mathbb{F}_{q^n})$ for enough n to pin down the rational function. Another is for curves of low genus over not-too-large fields, where one can diagnose the order of the class group using baby step-giant step.

[3] More precisely, to put rigid cohomology on a sound footing it should *also* be defined in the language of sites (Le Stum), then compared to more concrete constructions.

Contents

Zeta functions of hyperelliptic curves

For X a curve[4] of genus g over $\mathbb{F}_q$, we have

$$\zeta(X,T) = \frac{P_1(T)}{(1-T)(1-qT)},$$

$$P_1(T) = \sum_{i=0}^{2g} a_i T^i \in \mathbb{Z}[T], \qquad a_0 = 1, \qquad a_{g+i} = q^i a_{g-i}.$$

The roots of $P_1(T)$ in $\mathbb{C}$ lie on the circle $|T| = q^{-1/2}$ (Weil). Aside: the class group of X (a/k/a $\#J(\mathbb{F}_q)$ for J the Jacobian of X) has order $P_1(1)$.

As per the general setup, we wish to compute $P_1(T)$ as $\det(1-FT,V)$ for suitable F acting on suitable V.

[4]As a scheme, we want X to be of finite type over $\mathbb{F}_q$ of dimension 1 and also smooth, proper, and geometrically connected.

Zeta functions of hyperelliptic curves

For X a curve[4] of genus g over $\mathbb{F}_q$, we have

$$\zeta(X, T) = \frac{P_1(T)}{(1-T)(1-qT)},$$

$$P_1(T) = \sum_{i=0}^{2g} a_i T^i \in \mathbb{Z}[T], \qquad a_0 = 1, \qquad a_{g+i} = q^i a_{g-i}.$$

The roots of $P_1(T)$ in $\mathbb{C}$ lie on the circle $|T| = q^{-1/2}$ (Weil). Aside: the class group of X (a/k/a $\#J(\mathbb{F}_q)$ for J the Jacobian of X) has order $P_1(1)$.

As per the general setup, we wish to compute $P_1(T)$ as $\det(1 - FT, V)$ for suitable F acting on suitable V.

[4] As a scheme, we want X to be of finite type over $\mathbb{F}_q$ of dimension 1 and also smooth, proper, and geometrically connected.

Zeta functions of hyperelliptic curves

For X a curve[4] of genus g over $\mathbb{F}_q$, we have

$$\zeta(X,T) = \frac{P_1(T)}{(1-T)(1-qT)},$$

$$P_1(T) = \sum_{i=0}^{2g} a_i T^i \in \mathbb{Z}[T], \qquad a_0 = 1, \qquad a_{g+i} = q^i a_{g-i}.$$

The roots of $P_1(T)$ in $\mathbb{C}$ lie on the circle $|T| = q^{-1/2}$ (Weil). Aside: the class group of X (a/k/a $\#J(\mathbb{F}_q)$ for J the Jacobian of X) has order $P_1(1)$.

As per the general setup, we wish to compute $P_1(T)$ as $\det(1-FT,V)$ for suitable F acting on suitable V.

[4] As a scheme, we want X to be of finite type over $\mathbb{F}_q$ of dimension 1 and also smooth, proper, and geometrically connected.

The direct cohomological method (...)

Suppose $p \neq 2$ and X is a hyperelliptic curve of genus g with a rational Weierstrass point, which then admits an affine model $y^2 = Q(x)$ with Q monic of degree $2g+1$. We may then take V to be the first (algebraic) de Rham cohomology of a smooth lift of X over the unramified extension K of $\mathbb{Q}_p$ with residue field $\mathbb{F}_q$.

Concretely, for $\tilde{Q}$ a monic lift of Q, we have $V = \bigoplus_{i=0}^{2g-1} K \cdot \frac{x^i\,dx}{2y}$, with a quite explicit recipe for rewriting general differentials in terms of these.

The action of F is given by $x \mapsto x^q$ and (as a series)

$$y \mapsto y^q \left(1 + \frac{\tilde{Q}(x^q) - \tilde{Q}(x)^q}{\tilde{Q}(x)^q}\right)^{1/2}.$$

This leads to an algorithm implemented (and extended to cover $p = 2$ and X without a rational Weierstrass point) in `GP/Pari`, `Magma`, and `Sage`.

The direct cohomological method (...)

Suppose $p \neq 2$ and X is a hyperelliptic curve of genus g with a rational Weierstrass point, which then admits an affine model $y^2 = Q(x)$ with Q monic of degree $2g+1$. We may then take V to be the first (algebraic) de Rham cohomology of a smooth lift of X over the unramified extension K of $\mathbb{Q}_p$ with residue field $\mathbb{F}_q$.

Concretely, for $\tilde{Q}$ a monic lift of Q, we have $V = \bigoplus_{i=0}^{2g-1} K \cdot \frac{x^i\,dx}{2y}$, with a quite explicit recipe for rewriting general differentials in terms of these.

The action of F is given by $x \mapsto x^q$ and (as a series)

$$y \mapsto y^q \left(1 + \frac{\tilde{Q}(x^q) - \tilde{Q}(x)^q}{\tilde{Q}(x)^q}\right)^{1/2}.$$

This leads to an algorithm implemented (and extended to cover $p = 2$ and X without a rational Weierstrass point) in `GP/Pari`, `Magma`, and `Sage`.

The direct cohomological method (...)

Suppose $p \neq 2$ and X is a hyperelliptic curve of genus g with a rational Weierstrass point, which then admits an affine model $y^2 = Q(x)$ with Q monic of degree $2g+1$. We may then take V to be the first (algebraic) de Rham cohomology of a smooth lift of X over the unramified extension K of $\mathbb{Q}_p$ with residue field $\mathbb{F}_q$.

Concretely, for $\tilde{Q}$ a monic lift of Q, we have $V = \bigoplus_{i=0}^{2g-1} K \cdot \frac{x^i\,dx}{2y}$, with a quite explicit recipe for rewriting general differentials in terms of these.

The action of F is given by $x \mapsto x^q$ and (as a series)

$$y \mapsto y^q \left(1 + \frac{\tilde{Q}(x^q) - \tilde{Q}(x)^q}{\tilde{Q}(x)^q}\right)^{1/2}.$$

This leads to an algorithm implemented (and extended to cover $p = 2$ and X without a rational Weierstrass point) in `GP/Pari`, `Magma`, and `Sage`.

The direct cohomological method (...)

Suppose $p \neq 2$ and X is a hyperelliptic curve of genus g with a rational Weierstrass point, which then admits an affine model $y^2 = Q(x)$ with Q monic of degree $2g+1$. We may then take V to be the first (algebraic) de Rham cohomology of a smooth lift of X over the unramified extension K of $\mathbb{Q}_p$ with residue field $\mathbb{F}_q$.

Concretely, for $\tilde{Q}$ a monic lift of Q, we have $V = \bigoplus_{i=0}^{2g-1} K \cdot \frac{x^i\,dx}{2y}$, with a quite explicit recipe for rewriting general differentials in terms of these.

The action of F is given by $x \mapsto x^q$ and (as a series)

$$y \mapsto y^q \left(1 + \frac{\tilde{Q}(x^q) - \tilde{Q}(x)^q}{\tilde{Q}(x)^q}\right)^{1/2}.$$

This leads to an algorithm implemented (and extended to cover $p = 2$ and X without a rational Weierstrass point) in `GP/Pari`, `Magma`, and `Sage`.

The deformation method (Lauder): geometric picture

Let X_t be a family of hyperelliptic curves over $\mathbb{F}_q$ in one parameter t, lifted to a family $\tilde{X}_t$ over K. Then the relative de Rham cohomology of $\tilde{X}_t$ over the t-line forms a vector bundle of rank $2g$ away from the bad fibers, with the added structure of a *Gauss-Manin connection*.

Moreover, over a certain rigid-analytic subspace of the t-line, this connection admits a *Frobenius structure* which specializes to the Frobenius matrices described on the previous slide.

This gives an alternate approach to computing zeta functions, which is implemented in `Magma` (Hubrechts, Tuitman). There is also an implementation by Sebastian Pancratz as an optional `Sage` package.

I will *not* demonstrate this here. It is much simpler to do so for hypergeometric motives!

The deformation method (Lauder): geometric picture

Let X_t be a family of hyperelliptic curves over $\mathbb{F}_q$ in one parameter t, lifted to a family $\tilde{X}_t$ over K. Then the relative de Rham cohomology of $\tilde{X}_t$ over the t-line forms a vector bundle of rank $2g$ away from the bad fibers, with the added structure of a *Gauss-Manin connection*.

Moreover, over a certain rigid-analytic subspace of the t-line, this connection admits a *Frobenius structure* which specializes to the Frobenius matrices described on the previous slide.

This gives an alternate approach to computing zeta functions, which is implemented in `Magma` (Hubrechts, Tuitman). There is also an implementation by Sebastian Pancratz as an optional `Sage` package.

I will *not* demonstrate this here. It is much simpler to do so for hypergeometric motives!

The deformation method (Lauder): geometric picture

Let X_t be a family of hyperelliptic curves over $\mathbb{F}_q$ in one parameter t, lifted to a family $\tilde{X}_t$ over K. Then the relative de Rham cohomology of $\tilde{X}_t$ over the t-line forms a vector bundle of rank $2g$ away from the bad fibers, with the added structure of a *Gauss-Manin connection*.

Moreover, over a certain rigid-analytic subspace of the t-line, this connection admits a *Frobenius structure* which specializes to the Frobenius matrices described on the previous slide.

This gives an alternate approach to computing zeta functions, which is implemented in `Magma` (Hubrechts, Tuitman). There is also an implementation by Sebastian Pancratz as an optional `Sage` package.

I will *not* demonstrate this here. It is much simpler to do so for hypergeometric motives!

The deformation method (Lauder): geometric picture

Let X_t be a family of hyperelliptic curves over $\mathbb{F}_q$ in one parameter t, lifted to a family $\tilde{X}_t$ over K. Then the relative de Rham cohomology of $\tilde{X}_t$ over the t-line forms a vector bundle of rank $2g$ away from the bad fibers, with the added structure of a *Gauss-Manin connection*.

Moreover, over a certain rigid-analytic subspace of the t-line, this connection admits a *Frobenius structure* which specializes to the Frobenius matrices described on the previous slide.

This gives an alternate approach to computing zeta functions, which is implemented in `Magma` (Hubrechts, Tuitman). There is also an implementation by Sebastian Pancratz as an optional `Sage` package.

I will *not* demonstrate this here. It is much simpler to do so for hypergeometric motives!

The deformation method: concrete interpretation

In the context of the previous slide, there exist a $2g \times 2g$ matrix N over $K(t)$ with poles at the bad fibers[5] and a $2g \times 2g$ matrix F whose entries are rigid analytic functions defined away from some neighborhoods of the poles of N, satisfying the commutation relation

$$NF - pF\sigma(N) + t\frac{dF}{dt} = 0$$

where σ is the substitution $t \mapsto t^p$.

For any $\lambda \in \mathbb{F}_q$, let $[\lambda] \in K$ be its Teichmüller lift. Then $F([\lambda])$ equals the Frobenius matrix acting on some basis of the rigid cohomology of X_λ.

The matrix N can typically be computed easily. This imposes a differential equation on the entries of F, which can be solved after establishing an initial condition, e.g., by running the direct method on one fiber.

[5] and possibly more points depending on the choice of a basis of the vector bundle.

The deformation method: concrete interpretation

In the context of the previous slide, there exist a $2g \times 2g$ matrix N over $K(t)$ with poles at the bad fibers[5] and a $2g \times 2g$ matrix F whose entries are rigid analytic functions defined away from some neighborhoods of the poles of N, satisfying the commutation relation

$$NF - pF\sigma(N) + t\frac{dF}{dt} = 0$$

where σ is the substitution $t \mapsto t^p$.

For any $\lambda \in \mathbb{F}_q$, let $[\lambda] \in K$ be its Teichmüller lift. Then $F([\lambda])$ equals the Frobenius matrix acting on some basis of the rigid cohomology of X_λ.

The matrix N can typically be computed easily. This imposes a differential equation on the entries of F, which can be solved after establishing an initial condition, e.g., by running the direct method on one fiber.

[5] and possibly more points depending on the choice of a basis of the vector bundle.

The deformation method: concrete interpretation

In the context of the previous slide, there exist a $2g \times 2g$ matrix N over $K(t)$ with poles at the bad fibers[5] and a $2g \times 2g$ matrix F whose entries are rigid analytic functions defined away from some neighborhoods of the poles of N, satisfying the commutation relation

$$NF - pF\sigma(N) + t\frac{dF}{dt} = 0$$

where σ is the substitution $t \mapsto t^p$.

For any $\lambda \in \mathbb{F}_q$, let $[\lambda] \in K$ be its Teichmüller lift. Then $F([\lambda])$ equals the Frobenius matrix acting on some basis of the rigid cohomology of X_λ.

The matrix N can typically be computed easily. This imposes a differential equation on the entries of F, which can be solved after establishing an initial condition, e.g., by running the direct method on one fiber.

[5] and possibly more points depending on the choice of a basis of the vector bundle.

Contents

Why are HGMs better examples than HECs?

To illustrate the deformation method, we will use hypergeometric motives (HGMs) instead of hyperelliptic curves (HECs). Why?

- For HECs, one must compute the Gauss-Manin connection for a suitable one-parameter family. For HGMs, this is replaced with a simple explicit formula.
- For HECs, the connection may have many singularities, which contribute to the complexity of subsequent calculations. For HGMs, the only bad points are $t = 0, 1, \infty$.
- For HECs, we must develop[6] power series solutions at some point. For HGM, these are given explicitly by hypergeometric series.
- For HECs, we need an outside source for the initial condition on Frobenius. For HGMs, (conjecturally) there is an explicit formula.

[6] Silver lining: there is a quadratically convergent method for this.

Why are HGMs better examples than HECs?

To illustrate the deformation method, we will use hypergeometric motives (HGMs) instead of hyperelliptic curves (HECs). Why?

- For HECs, one must compute the Gauss-Manin connection for a suitable one-parameter family. For HGMs, this is replaced with a simple explicit formula.
- For HECs, the connection may have many singularities, which contribute to the complexity of subsequent calculations. For HGMs, the only bad points are $t = 0, 1, \infty$.
- For HECs, we must develop[6] power series solutions at some point. For HGM, these are given explicitly by hypergeometric series.
- For HECs, we need an outside source for the initial condition on Frobenius. For HGMs, (conjecturally) there is an explicit formula.

[6] Silver lining: there is a quadratically convergent method for this.

Why are HGMs better examples than HECs?

To illustrate the deformation method, we will use hypergeometric motives (HGMs) instead of hyperelliptic curves (HECs). Why?

- For HECs, one must compute the Gauss-Manin connection for a suitable one-parameter family. For HGMs, this is replaced with a simple explicit formula.
- For HECs, the connection may have many singularities, which contribute to the complexity of subsequent calculations. For HGMs, the only bad points are $t = 0, 1, \infty$.
- For HECs, we must develop[6] power series solutions at some point. For HGM, these are given explicitly by hypergeometric series.
- For HECs, we need an outside source for the initial condition on Frobenius. For HGMs, (conjecturally) there is an explicit formula.

[6] Silver lining: there is a quadratically convergent method for this.

Why are HGMs better examples than HECs?

To illustrate the deformation method, we will use hypergeometric motives (HGMs) instead of hyperelliptic curves (HECs). Why?

- For HECs, one must compute the Gauss-Manin connection for a suitable one-parameter family. For HGMs, this is replaced with a simple explicit formula.
- For HECs, the connection may have many singularities, which contribute to the complexity of subsequent calculations. For HGMs, the only bad points are $t = 0, 1, \infty$.
- For HECs, we must develop[6] power series solutions at some point. For HGM, these are given explicitly by hypergeometric series.
- For HECs, we need an outside source for the initial condition on Frobenius. For HGMs, (conjecturally) there is an explicit formula.

[6] Silver lining: there is a quadratically convergent method for this.

Why are HGMs better examples than HECs?

To illustrate the deformation method, we will use hypergeometric motives (HGMs) instead of hyperelliptic curves (HECs). Why?

- For HECs, one must compute the Gauss-Manin connection for a suitable one-parameter family. For HGMs, this is replaced with a simple explicit formula.
- For HECs, the connection may have many singularities, which contribute to the complexity of subsequent calculations. For HGMs, the only bad points are $t = 0, 1, \infty$.
- For HECs, we must develop[6] power series solutions at some point. For HGM, these are given explicitly by hypergeometric series.
- For HECs, we need an outside source for the initial condition on Frobenius. For HGMs, (conjecturally) there is an explicit formula.

[6] Silver lining: there is a quadratically convergent method for this.

The hypergeometric trace formula

Magma's HGM package computes Euler factors of the associated L-functions using a trace formula derived from Greene's finite hypergeometric functions. The trace over $\mathbb{F}_q$ for the parameter t equals[7]

$$\frac{1}{1-q}\sum_{r=0}^{q-2}\omega_p(M/t)^r Q_q(r)$$

where ω_p is the Teichmüller character and

$$Q_q(r) = (-1)^{m_0} q^{D+m_0-m_r} G_q(r)$$

where $G_q(r) = \prod_v g_q(rv)^{\gamma_v}$ where

$$g_q(a) = \sum_{u \in \mathbb{F}_q^\times} \omega_p(u)^{-a} \zeta_p^{\mathrm{trace}_{\mathbb{F}_q/\mathbb{F}_p}(u)} \qquad \text{is a Gauss sum.}$$

[7] The quantities M, D, m_r, and $v \mapsto \gamma_v$ are independent of t; we omit the definitions.

The hypergeometric trace formula

Magma's HGM package computes Euler factors of the associated L-functions using a trace formula derived from Greene's finite hypergeometric functions. The trace over $\mathbb{F}_q$ for the parameter t equals[7]

$$\frac{1}{1-q}\sum_{r=0}^{q-2}\omega_p(M/t)^r Q_q(r)$$

where ω_p is the Teichmüller character and

$$Q_q(r) = (-1)^{m_0} q^{D+m_0-m_r} G_q(r)$$

where $G_q(r) = \prod_v g_q(rv)^{\gamma_v}$ where

$$g_q(a) = \sum_{u\in\mathbb{F}_q^\times} \omega_p(u)^{-a}\zeta_p^{\mathrm{trace}_{\mathbb{F}_q/\mathbb{F}_p}(u)} \qquad \text{is a Gauss sum.}$$

[7] The quantities M, D, m_r, and $v \mapsto \gamma_v$ are independent of t; we omit the definitions.

The hypergeometric trace formula

Magma's HGM package computes Euler factors of the associated L-functions using a trace formula derived from Greene's finite hypergeometric functions. The trace over $\mathbb{F}_q$ for the parameter t equals[7]

$$\frac{1}{1-q}\sum_{r=0}^{q-2}\omega_p(M/t)^r Q_q(r)$$

where ω_p is the Teichmüller character and

$$Q_q(r) = (-1)^{m_0} q^{D+m_0-m_r} G_q(r)$$

where $G_q(r) = \prod_v g_q(rv)^{\gamma_v}$ where

$$g_q(a) = \sum_{u \in \mathbb{F}_q^\times} \omega_p(u)^{-a} \zeta_p^{\mathrm{trace}_{\mathbb{F}_q/\mathbb{F}_p}(u)} \qquad \text{is a Gauss sum.}$$

[7] The quantities M, D, m_r, and $v \mapsto \gamma_v$ are independent of t; we omit the definitions.

The hypergeometric trace formula

Magma's HGM package computes Euler factors of the associated L-functions using a trace formula derived from Greene's finite hypergeometric functions. The trace over $\mathbb{F}_q$ for the parameter t equals[7]

$$\frac{1}{1-q}\sum_{r=0}^{q-2}\omega_p(M/t)^r Q_q(r)$$

where ω_p is the Teichmüller character and

$$Q_q(r) = (-1)^{m_0} q^{D+m_0-m_r} G_q(r)$$

where $G_q(r) = \prod_v g_q(rv)^{\gamma_v}$ where

$$g_q(a) = \sum_{u\in\mathbb{F}_q^\times} \omega_p(u)^{-a}\zeta_p^{\mathrm{trace}_{\mathbb{F}_q/\mathbb{F}_p}(u)} \qquad \text{is a Gauss sum.}$$

[7] The quantities M, D, m_r, and $v \mapsto \gamma_v$ are independent of t; we omit the definitions.

Traces vs. deformations

Following GP/Pari (Cohen), Magma (Watkins) computes the Gauss sum $g_v(r)$ very efficiently using the Gross-Koblitz formula

$$g_v(a) = -\pi^{S_p(a)} \prod_{i=0}^{f-1} \Gamma_p\left(\frac{a^{(i)}}{q-1}\right)$$

where $f = \log_p q$; π is the $(p-1)$-st root of $-p$ for which $\zeta_p \equiv 1 + \pi \pmod{\pi^2}$; $S_p(a)$ is the sum of the base-p digits of a; $a^{(i)}$ is the remainder of $p^{-i}a$ modulo $q-1$; and Γ_p is the p-adic Gamma function.

This works well for computing L-functions: if you want all Fourier coefficients up to X, you only need traces for $q \leq X$.

However, if you want all the Euler factors for $p \leq X$ (e.g., to compute Sato-Tate statistics), you need traces for $q = p^1, \ldots, p^{\lfloor d/2 \rfloor}$ where d is the degree of the Euler factor. By contrast, deformation computes the whole Frobenius matrix at once, so has complexity linear in p rather than q.

Traces vs. deformations

Following GP/Pari (Cohen), Magma (Watkins) computes the Gauss sum $g_v(r)$ very efficiently using the Gross-Koblitz formula

$$g_v(a) = -\pi^{S_p(a)} \prod_{i=0}^{f-1} \Gamma_p\left(\frac{a^{(i)}}{q-1}\right)$$

where $f = \log_p q$; π is the $(p-1)$-st root of $-p$ for which $\zeta_p \equiv 1 + \pi \pmod{\pi^2}$; $S_p(a)$ is the sum of the base-p digits of a; $a^{(i)}$ is the remainder of $p^{-i}a$ modulo $q-1$; and Γ_p is the p-adic Gamma function.

This works well for computing L-functions: if you want all Fourier coefficients up to X, you only need traces for $q \leq X$.

However, if you want all the Euler factors for $p \leq X$ (e.g., to compute Sato-Tate statistics), you need traces for $q = p^1, \dots, p^{\lfloor d/2 \rfloor}$ where d is the degree of the Euler factor. By contrast, deformation computes the whole Frobenius matrix at once, so has complexity linear in p rather than q.

Traces vs. deformations

Following `GP/Pari` (Cohen), `Magma` (Watkins) computes the Gauss sum $g_v(r)$ very efficiently using the Gross-Koblitz formula

$$g_v(a) = -\pi^{S_p(a)} \prod_{i=0}^{f-1} \Gamma_p\left(\frac{a^{(i)}}{q-1}\right)$$

where $f = \log_p q$; π is the $(p-1)$-st root of $-p$ for which $\zeta_p \equiv 1 + \pi \pmod{\pi^2}$; $S_p(a)$ is the sum of the base-p digits of a; $a^{(i)}$ is the remainder of $p^{-i}a$ modulo $q-1$; and Γ_p is the p-adic Gamma function.

This works well for computing L-functions: if you want all Fourier coefficients up to X, you only need traces for $q \leq X$.

However, if you want all the Euler factors for $p \leq X$ (e.g., to compute Sato-Tate statistics), you need traces for $q = p^1, \dots, p^{\lfloor d/2 \rfloor}$ where d is the degree of the Euler factor. By contrast, deformation computes the whole Frobenius matrix at once, so has complexity linear in p rather than q.

Contents

Deformation for hypergeometric motives: demo

The remainder of the lecture consists of an explicit calculation of HGM Euler factors using the deformation method. This demo is contained in a Jupyter notebook: click here.

Disclaimer: the correctness of these calculations depends on various missing facts. Some of these should be easy to obtain (e.g., the amount of working p-adic and t-adic precision required to obtain the final answers) and some may be more difficult (e.g., the formula for the initial condition of the Frobenius structure).

Also, we do not claim that deformations can be used to compute bad Euler factors. However, there are only finitely many for any given t, corresponding to primes for which one of $t, t^{-1}, t-1$ has positive valuation, so there is no need to optimize this.

Deformation for hypergeometric motives: demo

The remainder of the lecture consists of an explicit calculation of HGM Euler factors using the deformation method. This demo is contained in a Jupyter notebook: click here.

Disclaimer: the correctness of these calculations depends on various missing facts. Some of these should be easy to obtain (e.g., the amount of working p-adic and t-adic precision required to obtain the final answers) and some may be more difficult (e.g., the formula for the initial condition of the Frobenius structure).

Also, we do not claim that deformations can be used to compute bad Euler factors. However, there are only finitely many for any given t, corresponding to primes for which one of $t, t^{-1}, t-1$ has positive valuation, so there is no need to optimize this.

Deformation for hypergeometric motives: demo

The remainder of the lecture consists of an explicit calculation of HGM Euler factors using the deformation method. This demo is contained in a Jupyter notebook: click here.

Disclaimer: the correctness of these calculations depends on various missing facts. Some of these should be easy to obtain (e.g., the amount of working p-adic and t-adic precision required to obtain the final answers) and some may be more difficult (e.g., the formula for the initial condition of the Frobenius structure).

Also, we do not claim that deformations can be used to compute bad Euler factors. However, there are only finitely many for any given t, corresponding to primes for which one of $t, t^{-1}, t-1$ has positive valuation, so there is no need to optimize this.