Computing zeta functions and L-functions
Lecture 1

Andrew V. Sutherland

June 24, 2019

CMI-HIMR Summer School in Computational Number Theory

Zeta functions of curves and their function fields

Recall that the zeta function of a nice curve $X/\mathbb{F}_q$ is defined by

$$Z_X(T) \coloneqq \exp\left(\sum_{r\geq 1} \frac{N_r}{r} T^r\right),$$

where $N_r \coloneqq \#X(\mathbb{F}_{q^r})$. Equivalently, if we put $K \coloneqq \mathbb{F}_q(X)$ then

$$Z_X(T) = Z_K(T) \coloneqq \sum_{n\geq 1} b_n T^n = \prod_{e\geq 1} (1 - T^e)^{-c_e},$$

where b_n counts effective divisors of degree n, and c_e counts prime divisors (places of K, equivalently, closed points of X) of degree e. Indeed, we have

$$N_r = \#X(\mathbb{F}_{q^r}) = \sum_{e|r} e c_e,$$

$$\log Z_X(T) = -\sum_{e\geq 1} c_e \log(1 - T^e) = \sum_{e\geq 1} c_e \sum_{d\geq 1} \frac{1}{d} T^{de} = \sum_{r\geq 1} \frac{N_r}{r} T^r.$$

Key properties of the zeta function of a curve

From the Weil conjectures for curves (and abelian varieties), we have

1. $Z_X(T) = \frac{L_X(T)}{(1-T)(1-qT)}$ with $L_X \in \mathbb{Z}[T]$ of degree $2g$.
2. $$\begin{aligned} L_X(T) = q^gT^{2g} &+ q^{g-1}a_1T^{2g-1} + \cdots + qa_{g-1}T^{g+1} + a_gT^g \\ 1 &+ \quad a_1T \quad + \cdots + \quad a_{g-1}T^{g-1} + \hookleftarrow \end{aligned}$$
3. $L_X(T) = \prod_{i=1}^{2g}(1-\alpha_i T)$ with $|\alpha_i| = q^{1/2}$;
4. $\#X(\mathbb{F}_{q^r}) = q^r + 1 - \sum_1^{2g}\alpha_i^r$ and $\#\operatorname{Jac}(X)(\mathbb{F}_{q^r}) = \prod_{i=1}^{2g}(1-\alpha_i^r)$.

It follows that $a_1,\ldots,a_g$ determine $N_1,\ldots,N_g$ and conversely, and that both determine $\#X(\mathbb{F}_{q^r})$ and $\#\operatorname{Jac}(X)(\mathbb{F}_{q^r})$ for all $r \geq 1$.

We also have the bounds $|a_i| \leq \binom{2g}{i}q^{i/2}$ (which are not tight in general). Setting all α_i to $\sqrt{q}$, and then to $-\sqrt{q}$, yields the Hasse-Weil bounds

$$(\sqrt{q}-1)^{2g} \leq \#\operatorname{Jac}(\mathbb{F}_q) \leq (\sqrt{q}+1)^{2g}$$

spanning an interval of width $4gq^{g-1/2} + O(q^{g-3/2})$.

The L-function of a curve

Now let X be a nice curve of genus g over a number field K.
The L-function of X is defined the Euler product

$$L(X,s) = L(\mathrm{Jac}(X),s) := \sum_{n\geq 1} a_n n^{-s} := \prod_{\mathfrak{p}} L_{\mathfrak{p}}(N(\mathfrak{p})^{-s})^{-1}.$$

where $\mathfrak{p}$ varies over the primes of K (prime ideals of $\mathcal{O}_K$) and $N(\mathfrak{p}) := \#\mathbb{F}_{\mathfrak{p}}$ is the cardinality of the residue field $\mathbb{F}_{\mathfrak{p}} := \mathcal{O}_K/\mathfrak{p}$.

For primes $\mathfrak{p}$ of good reduction for X we have $L_{\mathfrak{p}}(T) := L_{X_{\mathfrak{p}}}(T)$,
where $X_{\mathfrak{p}}$ denotes the reduction of X to the residue field $\mathbb{F}_{\mathfrak{p}}$.

In every case $L_{\mathfrak{p}} \in \mathbb{Z}[T]$ has degree at most $2g$.
Thus the a_n are integers and $L(X,s)$ is an arithmetic L-function of degree $2g$ with analytic normalization $L_{\mathrm{an}}(X, s+\frac{1}{2})$.

It can happen that X has bad reduction at $\mathfrak{p}$ but $\mathrm{Jac}(X)$ does not;
from the L-function perspective, these are good primes.

The Selberg class with polynomial Euler factors

The Selberg class S^{poly} consists of Dirichlet series $L(s) = \sum_{n \geq 1} a_n n^{-s}$:

1. $L(s)$ has an analytic continuation that is holomorphic at $s \neq 1$;
2. For some $\gamma(s) = Q^s \prod_{i=1}^{r} \Gamma(\lambda_i s + \mu_i)$ and ε, the completed L-function $\Lambda(s) := \gamma(s)L(s)$ satisfies the functional equation

   $$\Lambda(s) = \varepsilon \overline{\Lambda(1 - \bar{s})},$$

   where $Q > 0$, $\lambda_i > 0$, $\mathrm{Re}(\mu_i) \geq 0$, $|\varepsilon| = 1$. Define $\deg L := 2 \sum_i^r \lambda_i$.
3. $a_1 = 1$ and $a_n = O(n^{\epsilon})$ for all $\epsilon > 0$ (Ramanujan conjecture).
4. $L(s) = \prod_p L_p(p^{-s})^{-1}$ for some $L_p \in \mathbb{Z}[T]$ with $\deg L_p \leq \deg L$ (has an Euler product).

The Dirichlet series $L_{\mathrm{an}}(s, X) := L(X, s + \frac{1}{2})$ satisfies (3) and (4), and conjecturally lies in S^{poly}; for $g = 1$ and K totally real this is known.

Strong multiplicity one

Theorem (Kaczorowski-Perelli 2001)

If $A(s)=\sum_{n\geq 1} a_n n^{-s}$ and $B(s)=\sum_{n\geq 1} b_n n^{-s}$ lie in S^{poly} and $a_p=b_p$ for all but finitely many primes p, then $A(s)=B(s)$.

Corollary

If $L_{\mathrm{an}}(s,X)$ lies in S^{poly} then it is completely determined by any choice of all but finitely many coefficients a_p.

Henceforth we assume that $L_{\mathrm{an}}(s,X)\in S^{\mathrm{poly}}$.

Let $\Gamma_{\mathbb{C}}(s)=2(2\pi)^s\Gamma(s)$ and define $\Lambda(X,s):=\Gamma_{\mathbb{C}}(s)^g L(X,s)$. Then

$$\Lambda(X,s)=\varepsilon N^{1-s}\Lambda(X,2-s).$$

where the analytic root number $\varepsilon=\pm 1$ and the analytic conductor $N\in\mathbb{Z}_{\geq 1}$ are determined by the a_p (let us take these as definitions).

Testing the functional equation

Let $G(x)$ be the inverse Mellin transform of $\Gamma_{\mathbb{C}}(s)^g = \int_0^\infty G(x)x^{s-1}dx$, and define

$$S(x) := \frac{1}{x}\sum a_n G(n/x),$$

so that $\Lambda(X,s) = \int_0^\infty S(x)x^{-s}dx$, and for all $x > 0$ we have

$$S(x) = \varepsilon S(N/x).$$

The function $G(x)$ decays rapidly, and for sufficiently large c_0 we have

$$S(x) \approx S_0(x) := \frac{1}{x}\sum_{n\le c_0 x} a_n G(n/x),$$

with an explicit bound on the error $|S(x) - S_0(x)|$.

Effective strong multiplicity one

Fix a finite set of small primes $\mathcal{S}$ (e.g. $\mathcal{S} = \{2\}$) and an integer M that we know is a multiple of the conductor N (e.g. $M = \Delta(X)$).

There is a finite set of possibilities for $\varepsilon = \pm 1$, $N|M$, and the Euler factors $L_p \in \mathbb{Z}[T]$ for $p \in \mathcal{S}$ (the coefficients of $L_p(T)$ are bounded).

Suppose we can compute a_n for $n \le c_1\sqrt{M}$ whenever $p \nmid n$ for $p \in \mathcal{S}$.

We now compute $\delta(x) := |S_0(x) - \varepsilon S_0(N/x)|$ with $x = c_1\sqrt{N}$ for every possible choice of ε, N, and $L_p(T)$ for $p \in \mathcal{S}$. If all but one choice makes $\delta(x)$ larger than our explicit error bound, we know the correct choice.

For a suitable choice of c_1 this is guaranteed to happen.[1] One can explicitly determine a set of $O(N^\epsilon)$ candidate values of c_1, one of which is guaranteed to work; in practice the first one usually works.

---

[1]Subject to our assumptions; if it does not happen then we have found an explicit counterexample to the conjectured Langlands correspondence.

Conductor bounds

The formula of Brumer and Kramer gives explicit bounds on the p-adic valuation of the algebraic conductor N of $\mathrm{Jac}(X)$:

$$v_p(N) \leq 2g + pd + (p-1)\lambda_p(d),$$

where $d = \lfloor \frac{2g}{p-1} \rfloor$ and $\lambda_p(d) = \sum i d_i p^i$, with $d = \sum d_i p^i$, $0 \leq d_i < p$.

g	$p = 2$	$p = 3$	$p = 5$	$p = 7$	$p > 7$
1	8	5	2	2	2
2	20	10	9	4	4
3	28	21	11	13	6

For $g \leq 2$ these bounds are tight (see `www.lmfdb.org` for examples).

For hyperelliptic curves N divides $\Delta(X)$. What about other curves?

Algorithms to compute zeta functions

Given $X/\mathbb{Q}$ of genus g, we want to compute $L_p(T)$ for all good $p \le B$.

	complexity per prime (ignoring factors of $O(\log\log p)$)		
algorithm	$g = 1$	$g = 2$	$g = 3$
point enumeration	$p \log p$	$p^2 \log p$	$p^3 (\log p)^2$
group computation	$p^{1/4} \log p$	$p^{3/4} \log p$	$p(\log p)^2$
p-adic cohomology	$p^{1/2}(\log p)^2$	$p^{1/2}(\log p)^2$	$p^{1/2}(\log p)^2$
CRT (Schoof-Pila)	$(\log p)^5$	$(\log p)^8$	$(\log p)^{14}$
average poly-time	$(\log p)^4$	$(\log p)^4$	$(\log p)^4$

For $L(X, s) = \sum a_n n^{-s}$, we only need a_{p^2} for $p^2 \le B$, and a_{p^3} for $p^3 \le B$. For $1 < r \le g$ we can compute all a_{p^r} with $p^r \le B$ in time $O(B \log B)$ using naive point counting.

The bottom line: it all comes down to computing a_p's at good primes, equivalently, computing $\#X(\mathbb{F}_p) = p + 1 - a_p$ (aka counting points).

The divisor group of a curve (function field)

Recall that we have a (contravariant) equivalence of categories

$$\{\text{nice curves } X/\mathbb{F}_q\} \longleftrightarrow \{\text{function fields } K/\mathbb{F}_q\},$$

which sends X to $\mathbb{F}_q(X)$ and morphisms $\varphi\colon X \to Y$ to field embeddings $\varphi^*\colon \mathbb{F}_q(Y) \to \mathbb{F}_q(X)$ defined by $f \mapsto f \circ \varphi$.

We have a bijection between closed points P of X ($G_{\mathbb{F}_q}$-orbits of $X(\overline{\mathbb{F}}_q)$) and places P of K (equivalence classes of absolute values of K).

The divisor group $\operatorname{Div}(X) = \operatorname{Div}(K)$ is the free abelian group on closed points (places) P. Each $D \in \operatorname{Div}(X)$ has the form

$$D = \sum_P n_P P.$$

For $f \in K^\times$ we define $\operatorname{div}(f) := \sum_P v_P(f)P$, and let $\operatorname{Princ}(X)$ denote the subgroup $\{\operatorname{div}(f) : f \in K^\times\} \cup \{0\}$ of principal divisors.

The divisor class group

Define the homomorphism $\deg\colon \operatorname{Div}(X) \to \mathbb{Z}$ by $D \mapsto \sum_P n_P \deg(P)$, where $\deg(P) = \#P = [\kappa(P) : \mathbb{F}_q]$; note that $\operatorname{Princ}(X) \subseteq \ker \deg$.

We now define $\operatorname{Pic}(X) := \frac{\operatorname{Div}(X)}{\operatorname{Princ}(X)}$, and the divisor class group $\operatorname{Pic}^0(X)$ as the kernel of the degree map $\operatorname{Pic}(X) \to \mathbb{Z}$, yielding the exact sequence

$$0 \longrightarrow \operatorname{Pic}^0(X) \longrightarrow \operatorname{Pic}(X) \longrightarrow \mathbb{Z} \longrightarrow 0.$$

Provided that X has a rational point we have a functorial isomorphism $\operatorname{Pic}^0(X) \simeq \operatorname{Jac}(X)(\mathbb{F}_q)$, meaning $\operatorname{Pic}^0(X_L) \simeq \operatorname{Jac}(X)(L)$ for all $L/\mathbb{F}_q$.[2] We shall henceforth assume $X(\mathbb{F}_q)$ contains a rational point O. W now define the Abel-Jacobi map

$$\begin{aligned} X &\to \operatorname{Pic}^0(X) \\ P &\mapsto [O - P] \end{aligned}$$

When X is an elliptic curve this map is an isomorphism.

[2]This assumption is necessary, $\operatorname{Pic}^0(X_{\overline{\mathbb{F}}_q})^{G_{\mathbb{F}_q}}$ need not equal $\operatorname{Pic}^0(X)$ when $X(\mathbb{F}_q) = \emptyset$

Representing elements of the divisor class group

The Riemann-Roch theorem implies that if we fix $O \in X(\mathbb{F}_q)$, every $\alpha \in \operatorname{Pic}^0(X)$ can be written as $\alpha = [D - gO]$ for some $D \geq 0$.

This allows us to define a birational map between $\operatorname{Sym}^g(X) := X^g/S_g$ and $\operatorname{Jac}(X)$, but this map is not an isomorphism (see the exercises). Explicitly representing elements of $\operatorname{Jac}(X)$ is a hard problem, in general.

Now suppose X is defined by an equation $y^2 = f(x)$ with f monic, squarefree, of degree $2g+1$ and $g \geq 1$. Then X is an elliptic or hyperelliptic curve with a unique rational point ∞ at infinity.

Let $\pi\colon X \to \mathbb{P}^1$ be the x-coordinate projection and let $\phi \in \operatorname{Aut}(X)$ denote the hyperelliptic involution, which operates on the fibers of π by negating the y-coordinate. For each affine closed point P of X the monic polynomial $h_P \in \mathbb{F}_q[x]$ whose roots form $\pi(P)$ is an element of $\mathbb{F}_q(\mathbb{P}^1)$ that we can pullback via π to obtain a principal divisor

$$\operatorname{div}(\pi^*(h_p)) = P + \phi(P) - 2\deg(P)\infty$$

Mumford representation of divisor classes

With $X: y^2 = f(x)$ of genus g, each element of $\mathrm{Pic}^0(X)$ contains a divisor $D - g\infty$ with $D \geq 0$ which we can then write as $\bar{D} = P_1 + \cdots + P_n - n\infty$ with $P_1, \ldots, P_n \in X(\overline{\mathbb{F}}_q) - \{\infty\}$ such that $\phi(P_i) \neq P_j$ for any $j \neq i$ (with $0 \leq n \leq g$). Call such a $\bar{D}$ reduced.

If we put $P_i = (x_i : y_i : 1)$, note that if $x_i = x_j$ then $y_i = y_j$.

We now define $u(x) := \prod_{i=1}^{n}(x - x_i) \in \mathbb{F}_q[x]$, and let $v \in \mathbb{F}_q[x]$ be the unique polynomial of minimal degree such that $v(x_i) = y_i$ with appropriate multiplicity: if $(x - x_i)^k | u$ then $(x - x_i)^k | (v(x) - y_i)$.

The pair $(u, v) \in \mathbb{F}_q[x]^2$ representing a reduced divisor satisfies:

1. u is monic with $\deg(u) \leq g$,
2. $\deg(v) < \deg(u)$,
3. u divides $v^2 - f$ (because $\mathrm{ord}_{x=x_i}(u) \leq \mathrm{ord}_{x=x_i}(v^2 - f)$).

Any such Mumford pair (u, v) determines a reduced divisor $[u, v]$

Theorem: $[u, v] \sim [s, t] \Leftrightarrow (u, v) = (s, t)$.

Cantor's algorithm

The representation $[u, v]$ of reduced divisors is Mumford representation.
To compute in $\mathrm{Pic}^0(X)$ we then rely on Cantor's algorithm.

Input: Pairs $(u_1, v_1), (u_2, v_2) \in \mathbb{F}_q[x]^2$ for reduced divisors D_1, D_2. Output: The pair (u_3, v_3) representing the divisor class $[D_1 + D_2]$

1. Compute $d = \gcd(u_1, u_2, v_1 + v_2) = s_1 u_1 + s_2 u_2 + s_3(v_1 + v_2)$.
2. Compute $u \coloneqq u_1 u_2 / d^2$.
3. Compute $v \coloneqq (s_1 u_1 v_2 + s_2 u_2 v_1 + s_3(v_1 v_2 + f))/d \bmod u$.
4. While $\deg u > g$:
   4.1 Replace u with $(f - v^2)/u$ and then replace v with $-v \bmod u$.
5. Return (u_3, v_3)

To generate elements of $\mathrm{Pic}^0(X)$ we pick random monic $u \in \mathbb{F}_q[x]$ of degree at most g and try to find v such that (u, v) is a Mumford pair.

This is not always possible, but it must succeed with probability $\approx 1/2$, since $\# \mathrm{Pic}^0(\mathbb{F}_q) \approx q^g$.

Remarks and generalizations

Cantor's algorithm has bit-complexity $\tilde{O}(g^2(\log q))$ (near optimal), but the constant factors can be substantially improved for fixed g.

Cantor's algorithm can be generalized to handle arbitrary hyperelliptic curves (in two apparently different but ultimately equivalent ways).

Approximate current state of the art (odd characteristic, affine coords):

	rational WS point		no rational WS point	
genus	add	dbl	add	dbl
1	3**M**+**I**	4**M**+**I**		
2	24**M**+**I**	27**M**+**I**	28**M**+**I**	32**M**+**I**
3	67**M**+**I**	68**M**+**I**	79**M**+**I**	82**M**+**I**

As with elliptic curves, inversions can be avoided by using projective (or other) coordinates, but we prefer affine coordinates; the cost of inversions can be ameliorated with batching.