

NUMERICALLY COMPUTING RESIDUES OF L -FUNCTIONS

DLD

LAST UPDATED 2024.10.14

ABSTRACT. I worked with Eran Assaf to verify calculations in a paper we're writing. One of the things that we wanted to do was to help verify certain inner products and residues.

I record some observations and experiments from our work.

CONTENTS

1. Computing Petersson Inner Products via Hidden Unfolding	1
2. Numerically estimating residues of L -functions	2
2.1. Direct estimation	2
2.2. Partial sums	4
2.3. Smoothed partial sums	4
2.4. Rigorous estimates	6
References	6

1. COMPUTING PETERSSON INNER PRODUCTS VIA HIDDEN UNFOLDING

Suppose we are given two weight k holomorphic cusp forms f and g on $\Gamma_0(N) \subset \mathrm{SL}(2, \mathbb{Z})$, and we want to compute the Petersson inner product

$$\langle f y^{k/2}, g y^{k/2} \rangle = \iint_{\Gamma_0(N) \backslash \mathcal{H}} f(z) \overline{g(z)} y^k d\mu(z). \quad (1)$$

Numerically computing the Petersson inner product in general is hard. PariGP has some functionality here, but it can't do everything.

One approach is to insert the weight 0 Eisenstein series

$$E(z, s) := \sum_{\gamma \in \Gamma_\infty \backslash \Gamma_0(N)} \mathrm{Im}(\gamma z)^s.$$

Consider the inner product

$$\langle f \overline{g} y^k, E(z, \overline{s}) \rangle.$$

Writing $f = \sum a(n)q^n$ and $g = \sum b(n)q^n$, the standard unfolding argument with Eisenstein series shows that

$$\langle f \overline{g} y^k, E(z, \overline{s}) \rangle = \sum_{n \geq 1} a(n) \overline{b(n)} \int_0^\infty e^{-4\pi n y} y^{s+k-1} \frac{dy}{y} = \frac{\Gamma(s+k-1)}{(4\pi)^{s+k-1}} \sum_{n \geq 1} \frac{a(n) \overline{b(n)}}{n^{s+k-1}}.$$

This work was supported by the Simons Collaboration in Arithmetic Geometry, Number Theory, and Computation via the Simons Foundation grant 546235.

The point is that the Eisenstein series has a pole at $s = 1$ with constant residue c . Taking the residue shows that

$$c \langle f y^{k/2}, \bar{g} y^{k/2} \rangle = \operatorname{Res}_{s=1} \langle f \bar{g} y^k, E(z, \bar{s}) \rangle = \frac{\Gamma(k)}{(4\pi)^k} \operatorname{Res}_{s=1} \sum_{n \geq 1} \frac{a(n) \overline{b(n)}}{n^{s+k-1}}.$$

Stated differently, computing (1) can be transformed into a problem about computing a residue of, essentially, a Rankin–Selberg L -function.

Remark 1. This construction is used by Rankin and Selberg in their work introducing the L -function $L(s, f \times g)$. Note that $L(s, f \times g)$ should include a factor $\zeta(2s)$ (i.e. the factor that completes $E(z, s)$) to make it behave well. A very similar construction with a truncated Eisenstein series can be used to *compute* the residue c (and see that it's the volume of the fundamental domain). See for example §1.6 of Bump [Bum98].

I also remark that if f and g are merely cusp forms (and not holomorphic cusp forms), analogous work applies and the result is the same as long as the resulting integral over y is understandable. For example, if f and g are Maass forms, then similar results hold.

If f and g have different eigenvalues, then $\langle f, g \rangle = 0$. If they have the same eigenvalue, then the resulting y integral is understandable (in the sense that one can look it up in a table of integrals, or ask Mathematica).

2. NUMERICALLY ESTIMATING RESIDUES OF L -FUNCTIONS

We now look at the reduced problem. Given a Dirichlet series

$$D(s) = \sum_{n \geq 1} \frac{a(n)}{n^s}$$

that is presumed to have meromorphic continuation to $\operatorname{Re} s > 1 - \delta$ for some $\delta > 0$ and a simple pole at $s = 1$, how can we numerically estimate the residue?

2.1. Direct estimation. We begin very naively.

In a neighborhood of 1, we have

$$D(s) = \frac{r}{s-1} + f(s) \tag{2}$$

for a holomorphic function f . Thus

$$\lim_{s \rightarrow 1} (s-1)D(s) = r,$$

the residue. If we can compute $D(s)$ along a sequence of points tending to 1, then we might just try to estimate this limit.

In practice, $D(s)$ often converges absolutely for $\operatorname{Re} s > 1$, so we choose some small $\epsilon > 0$ and try looking at

$$(1 + \epsilon - 1)D(1 + \epsilon) = \epsilon \cdot D(1 + \epsilon) \approx r.$$

In practice, this behaves very badly. Let's look at $\zeta(s)$ and use the obvious approximation

$$\zeta(1 + \epsilon) = \sum_{n \geq 1} \frac{1}{n^{1+\epsilon}} \approx \sum_{n \leq N} \frac{1}{n^{1+\epsilon}}.$$

```

1 def zeta_approx(N, s):
2     return sum(1/(m**n(s)) for m in range(1, N + 1))
3
4 for epspower in (-2, -3, -4, -5, -6, -7):
5     eps = 10**epspower
6     s = 1 + eps
7     for npower in (3, 4, 5, 6):
8         N = 10**npower
9         print(f"N {N}, eps {eps}: {n(eps * zeta_approx(N, s))}")

```

This produces the following output:

```

N 1000, eps 1/100: 0.0725297980739927
N 10000, eps 1/100: 0.0937690500268027
N 100000, eps 1/100: 0.114528539813696
N 1000000, eps 1/100: 0.134815847783707

N 1000, eps 1/1000: 0.00746173653014950
N 10000, eps 1/1000: 0.00974539312390463
N 100000, eps 1/1000: 0.0120241987617199
N 1000000, eps 1/1000: 0.0142978033378315

N 1000, eps 1/10000: 0.000748309249198581
N 10000, eps 1/10000: 0.000978337305495686
N 100000, eps 1/10000: 0.00120835285756249
N 1000000, eps 1/10000: 0.00143831949805017

N 1000, eps 1/100000: 0.0000748523297425421
N 10000, eps 1/100000: 0.0000978718262077023
N 100000, eps 1/100000: 0.000120894841456149
N 1000000, eps 1/100000: 0.000143917731532429

N 1000, eps 1/1000000: 7.48544707142625e-6
N 10000, eps 1/1000000: 9.78756369334501e-6
N 100000, eps 1/1000000: 0.0000120900799291495
N 1000000, eps 1/1000000: 0.0000143926313619477

N 1000, eps 1/10000000: 7.48546848163299e-7
N 10000, eps 1/10000000: 9.78760180176266e-7
N 100000, eps 1/10000000: 1.20901395097692e-6
N 1000000, eps 1/10000000: 1.43927171867352e-6

```

We cannot usually expect better behavior than from $\zeta(s)$, and this looks *abysmal!* Thinking more about it, we can understand some of this behavior. Clearly the coefficients are all 1, and adding up N terms $1/n^{1+\epsilon}$ will be at most N . It doesn't make sense to choose N and ϵ such that $N\epsilon \ll 1$: the resulting residue approximation will trivially tend to 0. This explains why the estimated residues are getting smaller as $1+\epsilon \rightarrow 1$.

Of course, we know how to handle $\zeta(s)$ very well. But this illustrates the problem: if understanding the behavior at the pole is hard, then understanding behavior near

the pole might also be hard. Similarly, if we don't know how to understand the pole purely abstractly, we presumably don't have particularly good ways to understand the meromorphic continuation. We're then trying to study the continuation in its worst region of convergence.

This isn't the way.

2.2. Partial sums. Let's use a different strategy that avoids estimating $D(s)$ directly. Perron's formula says that

$$\sum_{n \leq X} a(n) = \frac{1}{2\pi i} \int_{(\sigma)} D(s) \frac{X^s}{s} ds,$$

where σ is large enough to be in the region of absolute convergence. We can be more precise and give a truncated Perron formula (see for example §5.3 of Montgomery-Vaughan [MV07]) of the form

$$\sum_{n \leq X} a(n) = \frac{1}{2\pi i} \int_{\sigma-iT}^{\sigma+iT} D(s) \frac{X^s}{s} ds + R(X, t)$$

for a boundable remainder $R(X, t)$. If $|a(n)| \ll 1$ (or indeed, even if this holds *on average*), the remainder is bounded by $X \log X/T$.

Remark 2. If $|a(n)| \ll (\log n)^A$, then the resulting remainder is bounded by $X(\log X)^{A+1}/T$. This probably isn't written down anywhere, but it's easy to show.

Using handwavy analysis, we should expect that we can shift σ slightly left of 1, picking up a pole at $s = 1$ from $D(s)$. The top and bottom contours are approximately bounded by $X \log X/T$. The left contour is annoying and really depends on our understanding of the meromorphic properties of $D(s)$. *Usually*, it has some reasonable growth property that means for some optimal choice of T , the residue dominates the integral and remainder.

Then we (usually) have a relationship of the form

$$\sum_{n \leq X} a(n) = \operatorname{Res}_{s=1} D(s) \frac{X^s}{s} + o(X) = Xr,$$

using r from the Laurent expansion (2). Thus

$$r \approx \frac{1}{X} \sum_{n \leq X} a(n),$$

with convergence depending on the actual bounds for the shifted integrals in Perron's analysis.

For $\zeta(s)$, this is of course perfect. The sum of the integers up to X is exactly X , so this recovers the residue exactly. Nonetheless, sometimes convergence is slow — and proving convergence requires some finicky analysis (very reminiscent of some proofs of the prime number theorem).

2.3. Smoothed partial sums. Alternately, we can use more forgiving approximations than Perron. For example, we can use the Riesz mean (sometimes called the Cesàro integral transform, and I don't know the actual history),

$$\sum_{n \leq X} a(n)(1 - n/X)^A = \frac{1}{2\pi i} \int_{\sigma} D(s) X^s \frac{\Gamma(1+A)\Gamma(s)}{\Gamma(1+A+s)} ds.$$

The left hand side remains easy to compute: it's a finite sum. The point is that if $D(s)$ is polynomially bounded in vertical strips, then one can choose $A \gg 1$ large enough that the integral converges beautifully even after shifting to the left.

The residue theorem then gives

$$\begin{aligned} \sum_{n \leq X} a(n)(1 - n/X)^A &= \operatorname{Res}_{s=1} D(s)X^s \frac{\Gamma(1+A)\Gamma(s)}{\Gamma(1+A+s)} + \int_{(\sigma')} D(s)X^s \frac{\Gamma(1+A)\Gamma(s)}{\Gamma(1+A+s)} ds \\ &= \frac{rX}{1+A} + O_A(X^{\sigma'}), \end{aligned}$$

where we assume A is large enough to guarantee absolute convergence of the shifted integral. (In practice, we take $A > B + 1$, where $|D(\sigma' + it)| \ll (1 + |t|)^B$ describes the polynomial bound for D and $0 < \sigma' < 1$).

For example, consider $\zeta(s)$ for convenience, we have

$$\sum_{n \leq X} (1 - n/X)^2 = \frac{X}{3} + O(X^\epsilon),$$

as we know the trivial convexity estimate

$$|\zeta(\epsilon + it)| \ll (1 + |t|)^{\frac{1}{2}}.$$

Explicitly computing, we find

```

1 tot = 0
2 X = 100000
3 for m in range(1, X + 1):
4     tot += n(1 - m/X)**2
5 residue = 3 / X * tot
6 print(residue)
7 # 0.999985000050005
    
```

That is,

$$\frac{3}{X} \sum_{n \leq X} (1 - n/X)^2 = 0.999985\dots,$$

and so the residue of $\zeta(s)$ at 1 is approximately 1.

Remark 3. This is the conclusion we came to at the end of our meeting. Use this! It works well and is very fast and easy to implement (assuming we can generate lots of coefficients).

After choosing A large enough for initial convergence, choosing larger smoothing powers doesn't necessarily correspond to more rapid convergence. Tail behavior is determined by the growth of the shifted integral. While there is a difference between "quickly growing" and "decaying", there is much less difference between "decaying" and "decaying more".

More technically, the function

$$D(s)X^s \frac{\Gamma(1+A)\Gamma(s)}{\Gamma(1+A+s)}$$

has a pole at $s = 0$ from the gamma functions (and depending on the Dirichlet series $D(s)$, possibly also a pole from $D(s)$ at $s = 0$). If A is sufficiently large and $D(s)$ sufficiently well-behaved that we can shift the line of integration to the left of 0, then the secondary contribution coming from this second pole will control convergence — and this behavior is essentially independent of A .

2.4. Rigorous estimates. To make these estimates rigorous, we would need an absolute bound or understanding of the constants from the $O(\cdot)$ bounds above.

In the case of automorphic L -functions, it would be possible to convert the Riesz mean approach described above into rigorous estimates for the residues (and thus rigorous estimates for the Petersson inner products). To make this fully rigorous, we would either need to have a proved bound for the coefficients $|a(n)|$ or accept a Ramanujan-Petersson type conjecture.

We can do any $GL(2)$ type Hecke modular form, for example: using either Deligne's result on $|a(p)|$ or the Kim–Sarnak bounds.

Coefficient bounds are also usually available for L -functions coming from varieties; the Hecke bound from elliptic curves, for example. (Of course we expect these to all be automorphic in some sense, but what we expect and what we can prove differ).

The structure of making this rigorous would look like

- (1) Give explicit bounds for $|a(n)|$.
- (2) Use this to bound $|L(2 + it)|$.
- (3) Produce the corresponding (weak exponent) convexity estimate on the interval $-1 \leq \operatorname{Re} s \leq 2$. Make the constant dependence explicit!
- (4) Choose $A \gg 1$ large enough to compensate for the weak convexity estimate growth.
- (5) Bound the resulting shifted integral along the line $\operatorname{Re} s = \epsilon$.

I note that I use a wider interval for the convexity argument than the typical $[0, 1]$ or $[-\epsilon, 1 + \epsilon]$ argument. This is because we are not optimizing for the best-possible polynomial growth in the strip: we are instead optimizing for easy-to-understand bounds for $L(2 + it)$ and simple, explicit coefficient growth.

I leave that for another day.

REFERENCES

- [Bum98] Daniel Bump. *Automorphic forms and representations*, volume 55. Cambridge University Press, 1998. (Cited on page 2)
- [MV07] Hugh L. Montgomery and Robert C. Vaughan. *Multiplicative number theory. I. Classical theory*, volume 97 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2007. (Cited on page 4)