The fine art of point counting

Andrew V. Sutherland

Massachusetts Institute of Technology

April 2, 2024

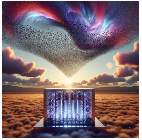

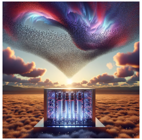

Simons Collaboration in Arithmetic Geometry, Number Theory, and Computation

What is arithmetic geometry?

Arithmetic geometers study solutions to polynomial equations like

Babylonians
circa 1800 BCE

$$x^2 + y^2 = z^2, \qquad\qquad x^2 - 4729494y^2 = 1,$$

Archimedes
251 BCE

$$x^3 + y^3 + z^3 = 3,$$

and even “cursed” equations like

$$y^4 + 5x^4 - 6x^2y^2 + 6x^3z + 26x^2yz + 10xy^2z - 10y^3z$$
$$-32x^2z^2 - 40xyz^2 + 24y^2z^2 + 32xz^3 - 16yz^3 = 0.$$

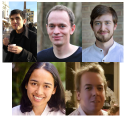

Balakrishnan et al.
2017

Which solutions?

There is a robust theory (classical algebraic geometry) to address this problem over $\mathbb{C}$.

But number theorists are typically interested in solutions over rings like $\mathbb{Z}$ or $\mathbb{Q}$.
These can be very hard to find. Indeed, it took sixty-five years to find

$$569936821221962380720^3 + (-569936821113563493509)^3 + (-472715493453327032)^3 = 3.$$

This problem is undecidable, in general (this is Hilbert's 10th problem).
We can simplify the problem by looking for solutions modulo primes.
In other words, we can count points over finite fields $\mathbb{F}_p$ (or $\mathbb{F}_q$).

These point counts can be used to define zeta functions and L-functions that encode the underlying structure of an equation (or system of equations) in a canonical way.

Counting points modulo p

Let's count points (x, y) on the curve $C\colon x^2 + y^2 = 1$ modulo primes p:

p	2	3	5	7	11	13	17	19	23	29	$\dots$
	2	4	4	8	12	12	16	20	24	28	$p \pm 1$

Better, count points $(x, y, z) \sim (cx, cy, cz)$ on $x^2 + y^2 = z^2 \bmod p$:

p	2	3	5	7	11	13	17	19	23	29	$\dots$
	3	4	6	8	12	14	18	20	24	30	$p + 1$

We always get $p + 1$. The L-function of C is

$$L(C, s) = \prod (1 - p^{-s})^{-1} = \sum n^{-s} = \zeta(s),$$

and we will get the same L-function whenever C has genus 0.

Elliptic curves

Let E be an elliptic curve over $\mathbb{Q}$, which we can write as

$$E\colon y^2 = x^3 + ax + b.$$

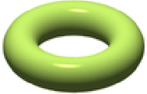

Every curve of genus 1 with rational points has this form.
You use elliptic curves every day! (on your phone/tablet/laptop)

The number of points on E modulo p can be written as

$$\#E_p(\mathbb{F}_p) = p + 1 - a_p,$$

where the trace of Frobenius a_p satisfies $|a_p| \leq 2\sqrt{p}$.

Frobenius

Let us now consider the sequence of real numbers $x_p := -t_p/\sqrt{p} \in [-2, 2]$.

Example: $y^2 = x^3 + x + 1$

p	t_p	x_p	p	t_p	x_p	p	t_p	x_p
3	0	**0.000000**	71	13	**−1.542816**	157	−13	**1.037513**
5	−3	**1.341641**	73	2	**−0.234082**	163	−25	**1.958151**
7	3	**−1.133893**	79	−6	**0.675053**	167	24	**−1.857176**
11	−2	**0.603023**	83	−6	**0.658586**	173	2	**−0.152057**
13	−4	**1.109400**	89	−10	**1.059998**	179	0	**0.000000**
17	0	**0.000000**	97	1	**−0.101535**	181	−8	**0.594635**
19	−1	**0.229416**	101	−3	**0.298511**	191	−25	**1.808937**
23	−4	**0.834058**	103	17	**−1.675060**	193	−7	**0.503871**
29	−6	**1.114172**	107	3	**−0.290021**	197	−24	**1.709929**
37	−10	**1.643990**	109	−13	**1.245174**	199	−18	**1.275986**
41	7	**−1.093216**	113	−11	**1.034793**	211	−11	**0.757271**
43	10	**−1.524986**	127	2	**−0.177471**	223	−20	**1.339299**
47	−12	**1.750380**	131	4	**−0.349482**	227	0	**0.000000**
53	−4	**0.549442**	137	12	**−1.025229**	229	−2	**0.132164**
59	−3	**0.390567**	139	14	**−1.187465**	233	−3	**0.196537**
61	12	**−1.536443**	149	14	**−1.146925**	239	−22	**1.423062**
67	12	**−1.466033**	151	−2	**0.162758**	241	22	**−1.417145**

a1 histogram of y^2 = x^3 + x + 1 for p <= 2^10
170 data points in 13 buckets, z1 = 0.029, out of range data has area 0.018

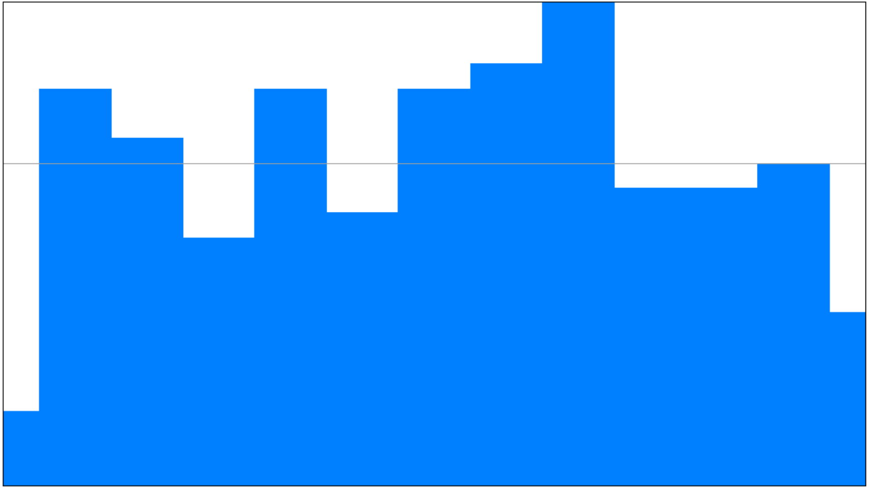

Moments: 1 0.051 1.039 0.081 2.060 0.294 4.971 1.134 13.278 4.308 37.954

a1 histogram of y^2 = x^3 + x + 1 for p <= 2^40
41203088794 data points in 202985 buckets

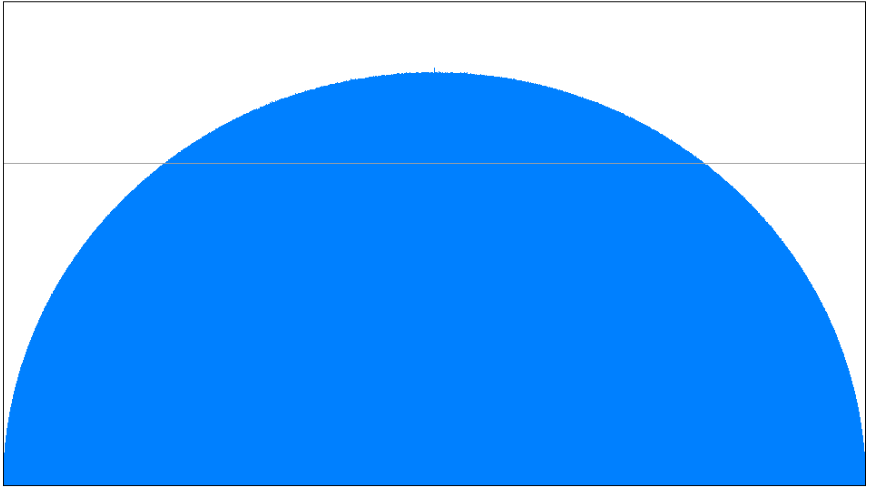

Moments: 1 0.000 1.000 0.000 2.000 0.000 5.000 0.000 14.000 0.000 41.999

a1 histogram of y^2 + xy + y = x^3 - x^2 - 200677624155755265850332082093385427509302303121789565 02x
+ 3448161179503055646703298569039072037485594435931918036126600829629193944873224342 9 for p <= 2^10
172 data points in 13 buckets, z1 = 0.023, out of range data has area 0.250

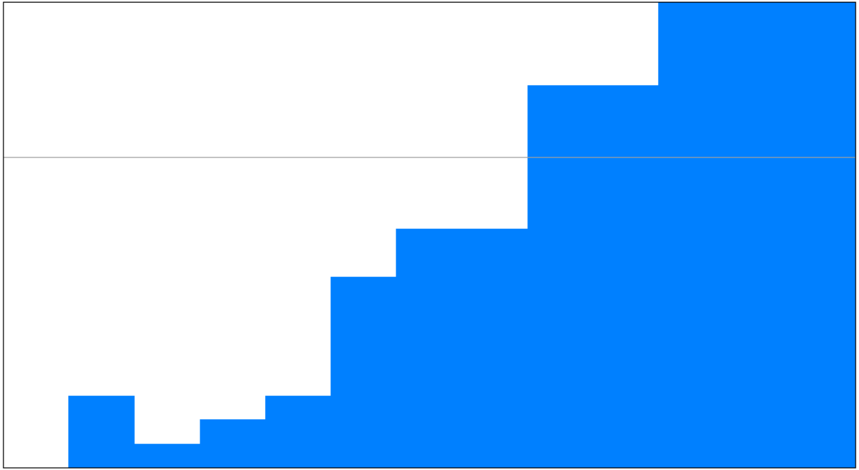

Moments: 1 1.034 1.716 2.532 4.446 7.203 13.024 22.220 40.854 72.100 133.961

a1 histogram of y^2 + xy + y = x^3 - x^2 - 20067762415575526585033208209338542750930230312178956502x
+ 34481611795030556467032985690390720374855944359319180361266008296291939448732243429 for p <= 2^40
41203088796 data points in 202985 buckets

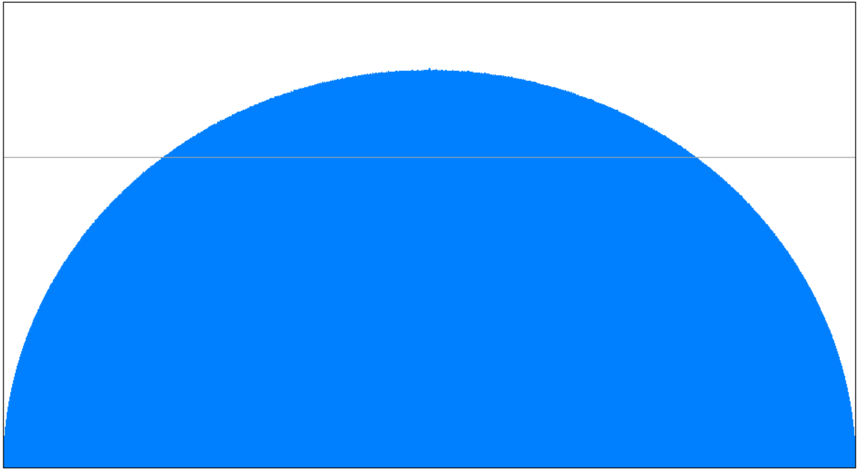

Moments: 1   0.000   1.000   0.000   2.000   0.000   5.000   0.001   14.000   0.003   42.000

The Sato-Tate conjecture

The Sato-Tate conjecture states that, except for certain families of well understood exceptions, we will always get the same limiting distribution as $p \to \infty$.

Mikio Sato

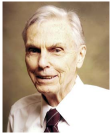

John Tate

Theorem (Taylor et al. 2008)

Let $E/\mathbb{Q}$ be an elliptic curve without any extra endomorphisms. The sequence x_p converges to the semi-circular distribution.

Richard Taylor

Richard Taylor received the 2014 Breakthrough Prize in Mathematics for this work.

Ranks of elliptic curves

Elliptic curves are not just curves, they are also abelian varieties. For $E/\mathbb{Q}$ the rational points are finitely generated: $E(\mathbb{Q}) \simeq T \oplus \mathbb{Z}^r$, where T is finite and r is the rank.

There are many things we do not know about r:

- Is there an algorithm to compute r?
- Which values of r can occur? Is there an upper limit?
- How often does each possible value of r occur, on average?

Theorem (Elkies 1990)

The value of r can be as large as 28.

Theorem (Bhargava-Shankar 2012)

The average value of r lies strictly between 0 and 1.

Conjectures of Birch and Swinnerton-Dyer

Based on extensive computer experiments (in the early 1960s!), Birch and Swinnerton-Dyer conjectured that

$$\lim_{x\to\infty} \prod_{p\leq x} \frac{\#E(\mathbb{F}_p)}{p} = c_E(\log x)^r$$

for some constant c_E, where r is the rank. This implies

$$\lim_{x\to\infty} \frac{1}{\log x} \sum_{p\leq x} \frac{a_p \log p}{p} = -r + \tfrac{1}{2},$$

and suggests that one can compute (or at least predict) r by counting points.

Birch and Swinnerton-Dyer eventually formulated a weaker conjecture relating the order of vanishing of $L(E,s)$ at $s=1$ to r (part of the BSD conjecture).

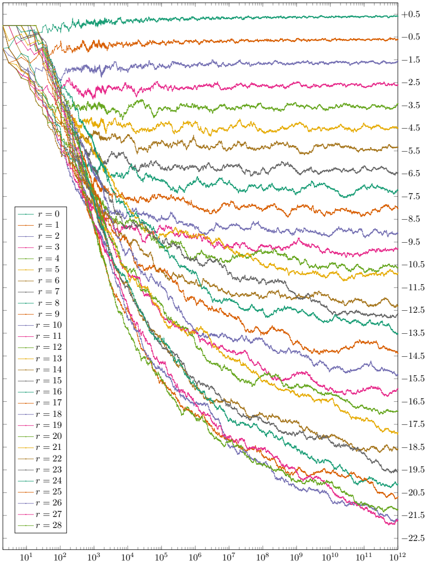

+0.5
−0.5
−1.5
−2.5
−3.5
−4.5
−5.5
−6.5
−7.5
−8.5
−9.5
−10.5
−11.5
−12.5
−13.5
−14.5
−15.5
−16.5
−17.5
−18.5
−19.5
−20.5
−21.5
−22.5
10^1
10^2
10^3
10^4
10^5
10^6
10^7
10^8
10^9
10^{10}
10^{11}
10^{12}
$r = 0$
$r = 1$
$r = 2$
$r = 3$
$r = 4$
$r = 5$
$r = 6$
$r = 7$
$r = 8$
$r = 9$
$r = 10$
$r = 11$
$r = 12$
$r = 13$
$r = 14$
$r = 15$
$r = 16$
$r = 17$
$r = 18$
$r = 19$
$r = 20$
$r = 21$
$r = 22$
$r = 23$
$r = 24$
$r = 25$
$r = 26$
$r = 27$
$r = 28$

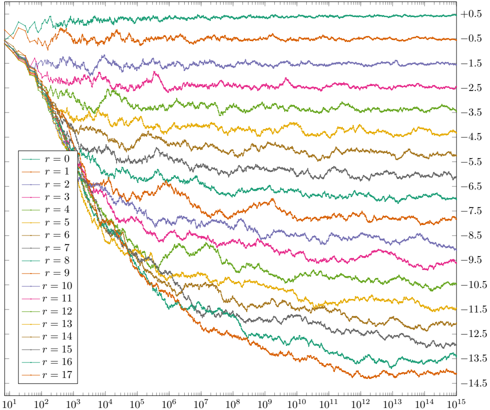

+0.5
−0.5
−1.5
−2.5
−3.5
−4.5
−5.5
−6.5
−7.5
−8.5
−9.5
−10.5
−11.5
−12.5
−13.5
−14.5
10^1
10^2
10^3
10^4
10^5
10^6
10^7
10^8
10^9
10^{10}
10^{11}
10^{12}
10^{13}
10^{14}
10^{15}
$r = 0$
$r = 1$
$r = 2$
$r = 3$
$r = 4$
$r = 5$
$r = 6$
$r = 7$
$r = 8$
$r = 9$
$r = 10$
$r = 11$
$r = 12$
$r = 13$
$r = 14$
$r = 15$
$r = 16$
$r = 17$

Murmurations of elliptic curves

In 2022, He, Lee, Oliver, and Pozdnyakov ran a series of machine learning experiments in an attempt to predict ranks of elliptic curves over $\mathbb{Q}$ using Frobenius traces.

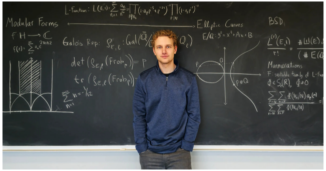

Their efforts to predict ranks were not particularly successful, but in the process they discovered a new and as yet unexplained oscillation in average Frobenius traces in families of elliptic curves ordered by conductor and separated by rank.

You can read more about their discovery in this recent Quanta article.

Murmurations of elliptic curves

Elliptic curves of conductor $N \in (2^n, 2^{n+1}]$ for $11 \le n \le 18$. Blue/red/purple dots at $(p, \bar{a}_p$ or $\bar{m}_p)$ are averages of a_p or $m_p := (-1)^r a_p(E)$ over even/odd/all $E/\mathbb{Q}$.

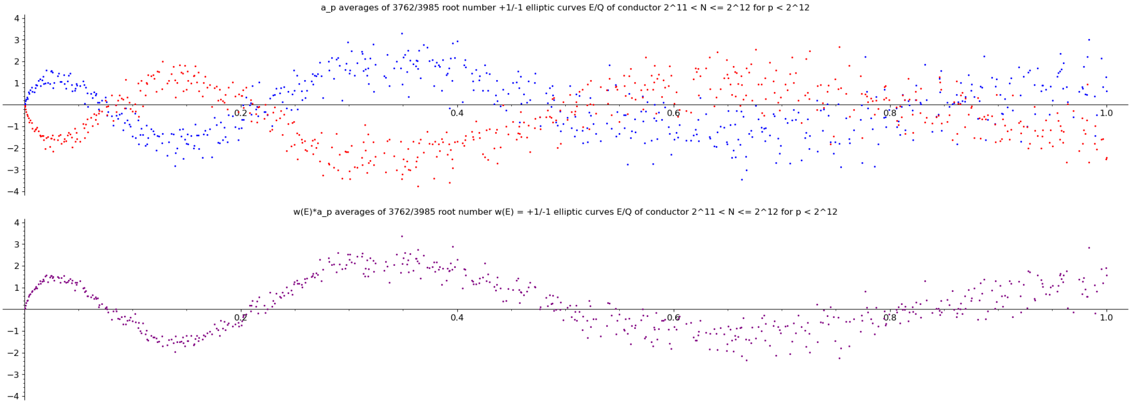

a_p averages of 3762/3985 root number +1/-1 elliptic curves E/Q of conductor 2^11 < N <= 2^12 for p < 2^12

w(E)*a_p averages of 3762/3985 root number w(E) = +1/-1 elliptic curves E/Q of conductor 2^11 < N <= 2^12 for p < 2^12

Murmurations of elliptic curves

Elliptic curves of conductor $N \in (2^n, 2^{n+1}]$ for $11 \leq n \leq 18$. Blue/red/purple dots at $(p, \bar{a}_p$ or $\bar{m}_p)$ are averages of a_p or $m_p := (-1)^r a_p(E)$ over even/odd/all $E/\mathbb{Q}$.

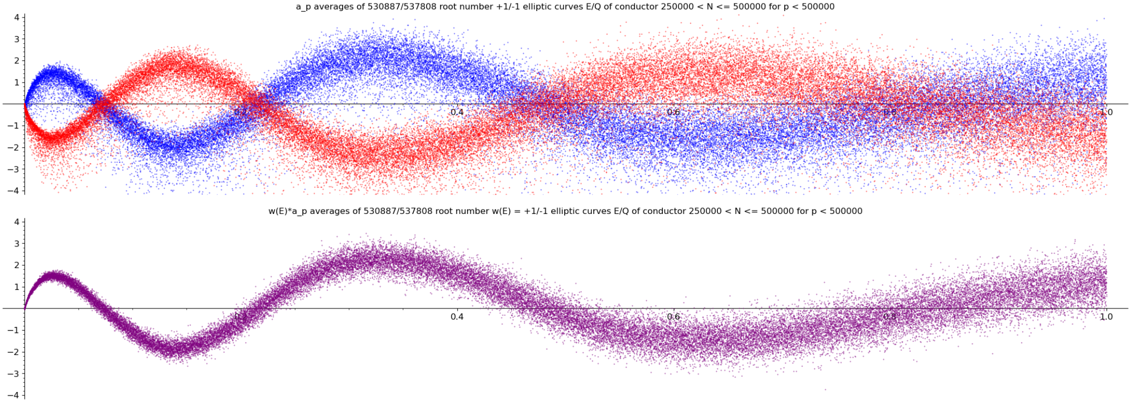

Modularity

The proof of the Sato-Tate conjecture is built on the Modularity Theorem.

Theorem (Taylor-Wiles 1995, Breuil-Conrad-Diamond-Taylor 2001)

For every elliptic curve $E/\mathbb{Q}$ there is a modular form f_E for which $L(E,s) = L(f_E,s)$.
The q-expansion $f_E(q) = \sum a_n q^n$ of f_E is determined by the Frobenius traces a_p of E.

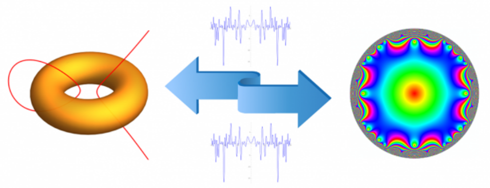

Corollary (Wiles 1995)

The equation $x^n + y^n = z^n$ has no nontrivial integer solutions for $n > 2$.

The L-functions and Modular Forms Database (LMFDB)

The relationship between elliptic curves and modular forms established by the Modularity Theorem was conjectured fifty years earlier.

Compelling evidence for this conjecture was obtained over many decades by tabulating elliptic curves and modular forms and computing their L-functions.

Extensive tables of these (and many other mathematical objects) are now available in the L-functions and Modular Forms Database.

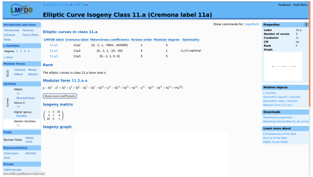

Elliptic Curve Isogeny Class 11.a (Cremona label 11a)

Elliptic curves in class 11.a

LMFDB label	Cremona label	Weierstrass coefficients	Torsion order	Modular degree	Optimality
11.a1	11a2	[0, -1, 1, -7820, -263580]	1	5	
11.a2	11a1	[0, -1, 1, -10, -20]	5	1	$\Gamma_0(N)$-optimal
11.a3	11a3	[0, -1, 1, 0, 0]	5	5	

Rank

The elliptic curves in class 11.a have rank 0.

Modular form 11.2.a.a

$q - 2q^{2} - q^{3} + 2q^{4} + q^{5} + 2q^{6} - 2q^{7} - 2q^{9} - 2q^{10} + q^{11} - 2q^{12} + 4q^{13} + 4q^{14} - q^{15} - 4q^{16} - 2q^{17} + 4q^{18} + O(q^{20})$

Isogeny matrix

$$\begin{pmatrix} 1 & 5 & 25 \\ 5 & 1 & 5 \\ 25 & 5 & 1 \end{pmatrix}$$

Isogeny graph

Properties

Label	11.a
Number of curves	3
Conductor	11
CM	no
Rank	0
Graph	

Related objects

- L-function
- Symmetric square L-function
- Symmetric cube L-function
- Modular form 11.2.a.a

Downloads

- Download q-expansion
- Download stored data for all curves

Learn more about

- Completeness of the data
- Source of the data
- Elliptic Curve labels

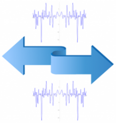

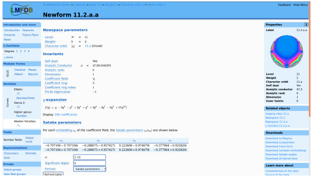

Newform 11.2.a.a

Newspace parameters

Level: $N = 11$
Weight: $k = 2$
Character orbit: $[\chi] =$ 11.a (trivial)

Invariants

Self dual:		Yes
Analytic Conductor	$A =$	47.3311945276
Analytic rank:		0
Dimension:		1
Coefficient field:		$\mathbb{Q}$
Coefficient ring:		$\mathbb{Z}$
Coefficient ring index:		1
Fricke eigenvalue:		-1

q-expansion

$f(q) = q - 2q^{2} - q^{3} + 2q^{4} + q^{5} + 2q^{6} - 2q^{7} - 2q^{9} + O(q^{10})$

Display 100 coefficients

Satake parameters

For each embedding ι_m of the coefficient field, the Satake parameters $\iota_m(\alpha_p)$ are shown below.

α_2	α_3	α_5	α_7
$-0.707106 - 0.707106i$	$-0.288675 - 0.957427i$	$0.223606 - 0.974679i$	$-0.377964 - 0.925820i$
$-0.707106 + 0.707106i$	$-0.288675 + 0.957427i$	$0.223606 + 0.974679i$	$-0.377964 + 0.925820i$

n: 2-10
Significant digits: 6
Format: Satake parameters

Properties

Label	11.2.a.a
Level	11
Weight	2
Character orbit	11.a
Self dual	Yes
Analytic conductor	47.3
Analytic rank	0
Dimension	1
Inner twists	0

Related objects

- Isogeny class 11.a
- Newspace 11.2
- Newspace 11.2.a
- L-function 11.2.a.a

Downloads

- Download to Magma
- Download q-expansion
- Download trace form
- Download complex embeddings
- Download Satake angles
- Download all stored data

Learn more about

- Completeness of the data
- Source of the data

The Langlands Program

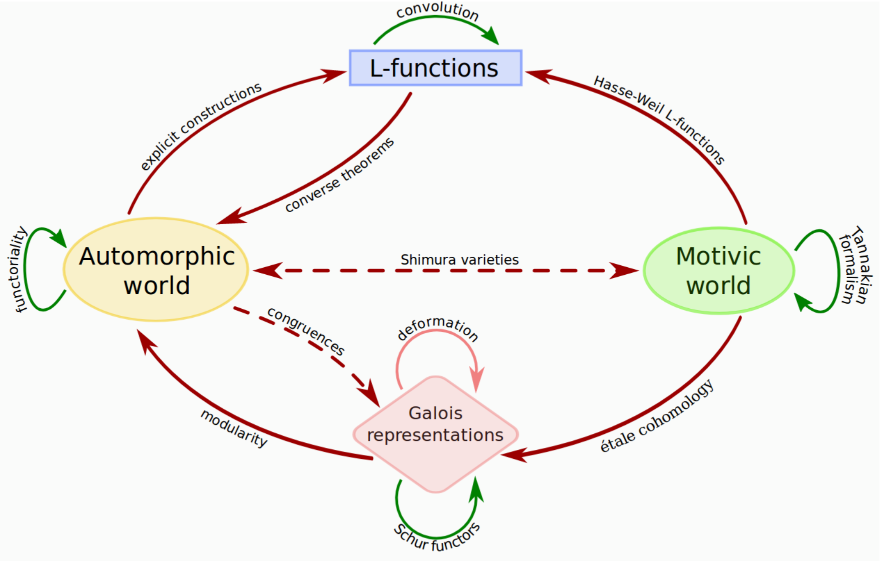

L-functions and zeta functions

The L-function of a (nice) curve $X/\mathbb{Q}$ can be written as

$$L(X,s) := \prod L_p(p^{-s})^{-1} = \sum a_n n^{-s}.$$

The completed L-function $\Lambda(s) := \Gamma_{\mathbb{C}}(s)L(X,s)$ is expected to satisfy

$$\Lambda(s) = \pm N^{1-s}\Lambda(2-s).$$

For good primes p the L-polynomial $L_p(T)$ appears in the zeta function

$$Z(X_p;T) := \exp\left(\sum_{r\geq 1} \#X_p(\mathbb{F}_{p^r})\frac{T^r}{r}\right) = \frac{L_p(T)}{(1-T)(1-pT)}.$$

The Langlands conjectures imply an effective multiplicity one result for $L(X,s)$, which is determined by $\#X(\mathbb{F}_p)$ good primes $p \leq B$, where B depends on the conductor N.

Algorithms to compute zeta functions

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

	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)	$(\log p)^5$	$(\log p)^8$	$(\log p)^{12}$
average poly-time	$(\log p)^4$	$(\log p)^4$	$(\log p)^4$

The bottom row is due to a 2014 breakthrough by David Harvey, and many later refinements [HS 2016, HS 2018, S 2020, CHS 2023].

Timings for genus 3 curves

Time to compute $L_p(T) \bmod p$ for all good $p \leq B$.

B	plane quartic (old)	plane quartic (new)	hyperelliptic (old)	hyperelliptic (new)
2^{12}	18	1.4	1.3	0.1
2^{13}	49	2.4	2.6	0.2
2^{14}	142	4.6	5.4	0.5
2^{15}	475	8.4	12	1.0
2^{16}	1,670	18	29	2.1
2^{17}	5,880	34	74	5.3
2^{18}	22,300	71	192	14
2^{19}	78,100	159	532	37
2^{20}	297,000	421	1,480	97
2^{21}	1,130,000	996	4,170	244
2^{22}	4,280,000	2,430	12,200	617
2^{23}	16,800,000	5,930	36,800	1,500
2^{24}	66,800,000	13,700	113,000	3,520
2^{25}	244,000,000	31,100	395,000	8,220
2^{26}	972,000,000	73,400	1,060,000	19,700

(Intel Xeon E7-8867v3 3.3 GHz CPU seconds).

Sato-Tate groups and their distributions

There are two Sato-Tate distributions that arise for elliptic curves $E/\mathbb{Q}$, depending on whether E is exceptional or not, but over general number fields there are three:

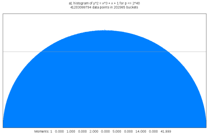

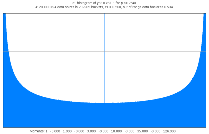

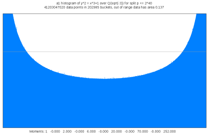

Each corresponds to the distributions of traces in a compact subgroup of SU(2), the Sato–Tate group of E. Katz–Sarnak and Serre extended this random matrix model to curves of genus g and abelian varieties of dimension g using subgroups of USp($2g$).

The Sato-Tate conjecture is open for genus $g > 1$, but Sato-Tate groups have been completely classified for for $g \leq 3$. There are 52 in genus 2 and 410 in genus 3 [Fité-Kedlaya-Rotger-S 2012, Fité-Kedlaya-S 2019].

These classifications involved more than a thousand CPU-years of computation.

Sato-Tate trace distributions of genus 2 curves:

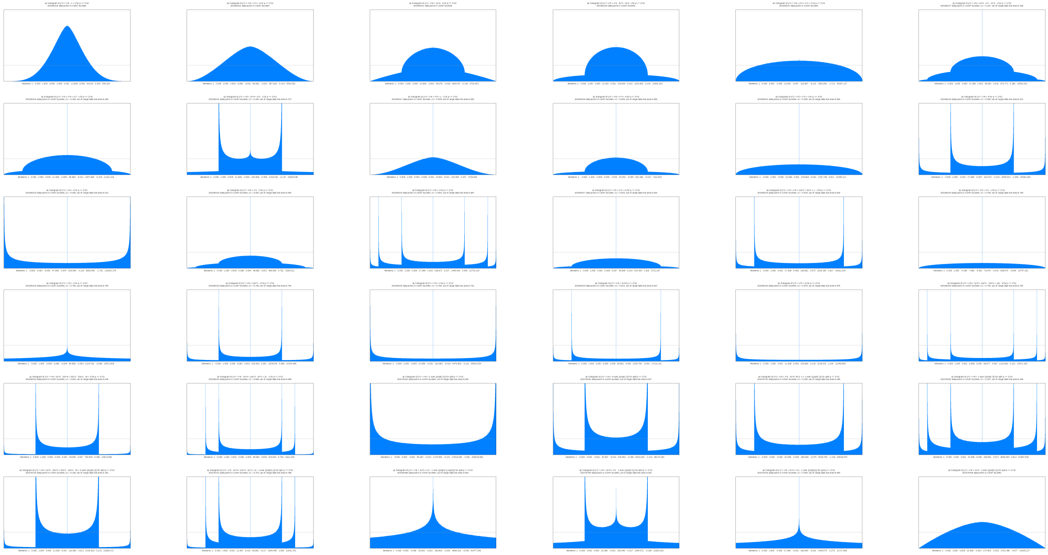

Building a database of low genus curves

To make it feasible to compute L-functions, to make the Langlands correspondence explicit and to investigate murmurations, a we want to tabulate curves by conductor.

No one knows how to do this for curves of genus $g > 1$, not even in principle!
We instead "sieve the sky". We enumerate vast numbers of curves with small coefficients along with their discriminants (which bound the conductor).

In our genus 2 and genus 3 searches we enumerated a total of more than 10^{18} curves, keeping only about 10^6 curves of interest in each case.

To make such a computation feasible requires:

- efficient enumeration algorithms (equations/discriminants in parallel);
- code optimization down to the metal (about 10 nanoseconds per curve);
- massive parallelism (thousands or even millions of cores).

Computation of Jan 9, 2022

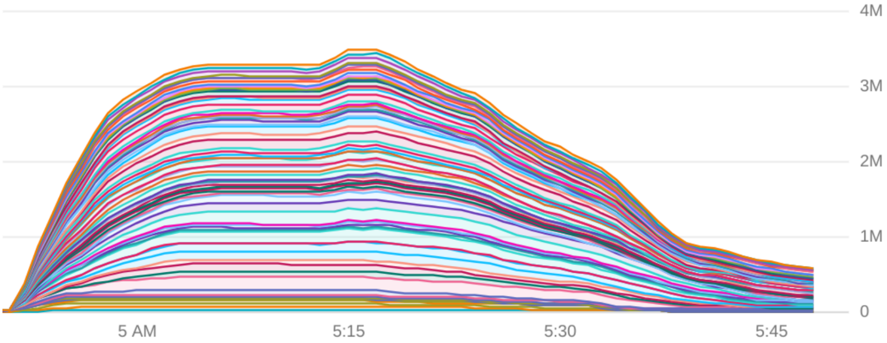

We used 126080 32-vCPU instances in 73 data centers in 23 different locations, including Taiwan, Hong Kong, Tokyo, Osaka, Mumbai, Singapore, Sydney, Finland, Belgium, London, Frankfurt, the Netherlands, Zurich, Montréal, São Paulo, Iowa, South Carolina, Virginia, Oregon, Los Angeles, Salt Lake City, and Las Vegas.

How much carbon does a 300 vCPU-year computation emit?

This is a question http://www.green-algorithms.org/ can help answer.

300 vCPU-years is about 1 314 900 core-hours (2 vCPUs per core).

CPU	cores	platform	location	energy	carbon
i9-9900K (64GB)	1	desktop	Massachusetts	46.99 MWh	19 750 Kg
i9-9900K (64GB)	16	desktop	Massachusetts	17,61 MWh	7 400 Kg
Ryzen 3990X (256GB)	64	desktop	Massachusetts	7.44 MWh	3 260 Kg
Ryzen 3990X (256GB)	64	cloud	Virginia	8.60 MWh	2 650 Kg
Ryzen 3990X (256GB)	64	cloud	Montreal	8.60 MWh	13 Kg

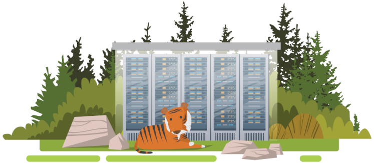

Some preliminary results

We found more than 500 genus 2 curves with conductors $N \leq 1000$, including the curve

$$C_{903} : y^2 + (x^2+1)y = x^5 + 3x^4 - 13x^3 - 25x^2 + 61x - 28$$

of conductor 903 and whose L-function coefficients match the a_n of the paramodular form of conductor 903 computed by Poor-Yuen and extended by Mellit to $n \leq 100$.

We also found curves of conductor 657, 760, 775, 924 not previously known to occur for Jacobians, examples of arithmetic phenomena that do not arise among the curves in the LMFDB, and many more curves of small conductor:

conductor bound	1000	10000	100000	1000000
curves in LMFDB	159	3069	20265	66158
curves found	807	25438	426396	5059724
L-functions in LMFDB	109	2807	19775	65534
L-functions found	200	9409	186811	2315328

Murmurations of genus 2 curves over $\mathbb{Q}$

Genus 2 curves $X/\mathbb{Q}$ of conductor $N \in (2^n, 2^{n+1}]$ for $12 \leq n \leq 19$. Blue/red/purple dots at $(p, \bar{a}_p$ or $\bar{m}_p)$ are averages of a_p or $m_p := (-1)^r a_p(X)$ over even/odd/all $X/\mathbb{Q}$.

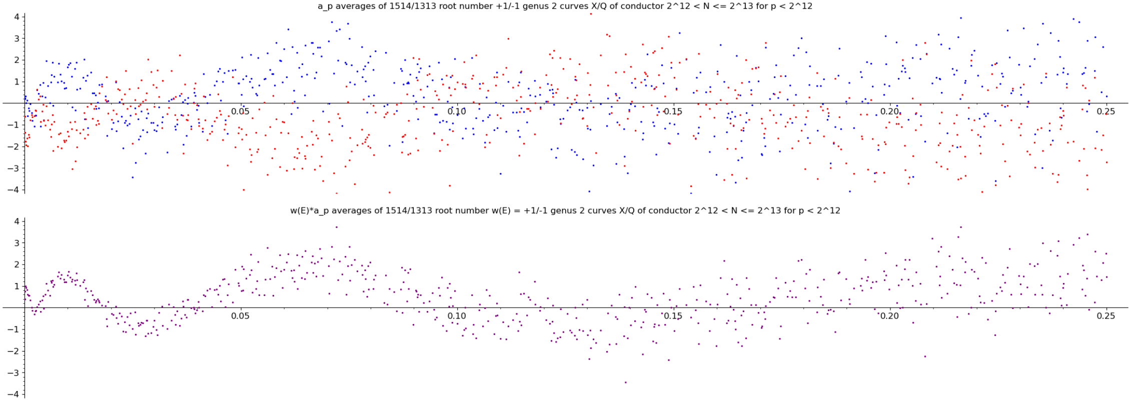

Murmurations of genus 2 curves over $\mathbb{Q}$

Genus 2 curves $X/\mathbb{Q}$ of conductor $N \in (2^n, 2^{n+1}]$ for $12 \le n \le 19$. Blue/red/purple dots at $(p, \bar{a}_p$ or $\bar{m}_p)$ are averages of a_p or $m_p := (-1)^r a_p(X)$ over even/odd/all $X/\mathbb{Q}$.

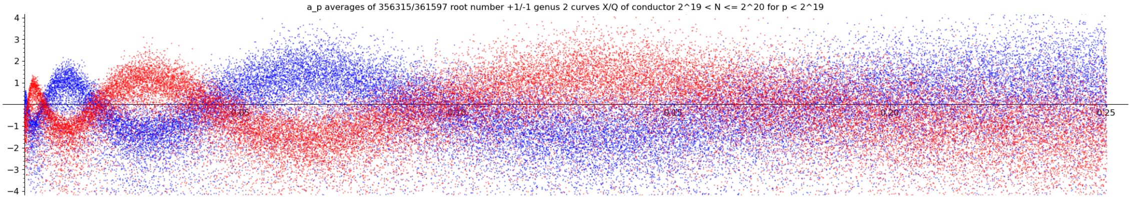

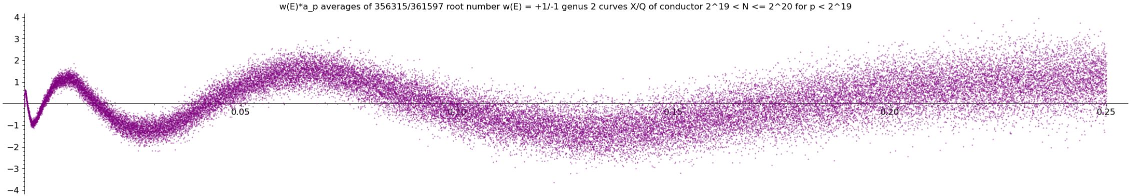

Sixteenth Algorithmic Number Theory Symposium

Massachusetts Institute of Technology
July 15–19, 2024

ANTS XVI