Computing the image of Galois

Andrew V. Sutherland

Definitions

Let E/K be an elliptic curve, and let $\ell \neq \mathrm{char}(K)$ be prime.

Let $L = K(E[\ell])$ be the Galois extension of K obtained by adjoining the coordinates of the ℓ-torsion points of $E(\bar{K})$ to K.

The Galois group $\mathrm{Gal}(L/K)$ acts linearly on the ℓ-torsion points

$$E[\ell] \simeq \mathbb{Z}/\ell\mathbb{Z} \oplus \mathbb{Z}/\ell\mathbb{Z},$$

yielding a group representation

$$\bar{\rho}_{E,\ell} \colon \mathrm{Gal}(L/K) \longrightarrow \mathrm{Aut}(E[\ell]) \simeq \mathrm{GL}_2(\mathbb{Z}/\ell\mathbb{Z}).$$

This is the *mod-ℓ Galois representation* attached to E.

Definitions

More generally, the Galois group $\text{Gal}(\bar{K}/K)$ acts on the *ℓ-adic Tate module*

$$T_\ell(E) = \varprojlim_n E[\ell^n],$$

yielding a group representation

$$\rho_{E,\ell}\colon \text{Gal}(\bar{K}/K) \longrightarrow \text{Aut}\big(T_\ell(E)\big) \simeq \text{GL}_2(\mathbb{Z}_\ell).$$

This is the *ℓ-adic Galois representation* attached to E.

We may view $\bar{\rho}_{E,\ell}$ as the reduction of $\rho_{E,\ell}$ modulo ℓ.

Surjectivity of $\rho_{E,\ell}$

For E without complex multiplication, $\rho_{E,\ell}$ is usually surjective.

Theorem (Serre)

Let K be a number field and assume E/K does not have CM.

1. *The image of $\rho_{E,\ell}$ has finite index in $\mathrm{GL}_2(\mathbb{Z}_\ell)$ for all ℓ.*
2. $\mathrm{im}\,\rho_{E,\ell} = \mathrm{GL}_2(\mathbb{Z}_\ell)$ *for all sufficiently large ℓ, say $\ell > \ell_{\max}$.*

Conjecturally, there is an $\ell_{\max}$ that depends only on K.

For $K = \mathbb{Q}$, it is believed that $\ell_{\max} = 37$.

For this talk, $K = \mathbb{Q}$.

Reduction modulo ℓ

We shall restrict our attention primarily to $\bar{\rho}_{E,\ell} = \rho_{E,\ell} \bmod \ell$.

Theorem (Serre)

For $K = \mathbb{Q}$ and $\ell > 3$, the map $\rho_{E,\ell}$ is surjective iff $\bar{\rho}_{E,\ell}$ is.

The theorem fails for $\ell = 2$ and $\ell = 3$, but in these cases it suffices to consider $\rho_{E,\ell} \bmod 2^3$ and $\rho_{E,\ell} \bmod 3^2$.

When is $\bar{\rho}_{E,\ell}$ non-surjective?

If $E[\ell](\mathbb{Q})$ is non-trivial, then $\bar{\rho}_{E,\ell}$ is not surjective.
This occurs for $\ell \leq 7$ (Mazur).

If $E/\mathbb{Q}$ admits a rational ℓ-isogeny, then $\bar{\rho}_{E,\ell}$ is not surjective.
For E without CM, this occurs for $\ell \leq 17$ and $\ell = 37$ (Mazur).

However, $\bar{\rho}_{E,\ell}$ may be non-surjective even when $E/\mathbb{Q}$ admits no rational ℓ-isogenies, and $\operatorname{im}\bar{\rho}_{E,\ell}$ may vary in any case.

Classifying the possible subgroups $\operatorname{im}\bar{\rho}_{E,\ell} \subseteq \mathrm{GL}_2(\mathbb{Z}/\ell\mathbb{Z})$ may be viewed as a generalization of Mazur's theorems.

As a first step, we wish to determine which groups do occur.

Main results

A very fast algorithm to compute $\operatorname{im}\bar{\rho}_{E,\ell}$ up to isomorphism, (and essentially up to conjugacy), for small primes ℓ.

If $\bar{\rho}_{E,\ell}$ is surjective, the algorithm proves this unconditionally. If not, its output is heuristically correct with very high probability. (in principle, this can also be made unconditional).

We have tested every elliptic curve in the tables of Cremona and Stein-Watkins (about 137 million curves) for all $\ell < 60$, as well as some 10^{10} curves in various families.

This has yielded what we believe to be a complete classification of $\operatorname{im}\bar{\rho}_{E,\ell}$ for elliptic curves over $\mathbb{Q}$ without CM, at least up to isomorphism (but work is still in progress).

Prior work

Reverter-Vila (2001) determined $\bar{\rho}_{E,\ell}$ for all elliptic curves $E/\mathbb{Q}$ with conductor less than 200 (a total of 739 curves).

Stein (2005) obtained partial results for elliptic curves $E/\mathbb{Q}$ with conductors up to 30000 (a total of 66561 curves).

Zywina (2011) developed an efficient algorithm to determine, given $E/\mathbb{Q}$, the set of primes ℓ for which $\bar{\rho}_{E,\ell}$ is not surjective.

There is a large body of related work (for example, see David-Kisilevsky-Pappalardi, Lang-Trotter, Koblitz-Zywina, ...).

A probabilistic approach

The action of the Frobenius endomorphism on $E[\ell](\mathbb{F}_p)$ corresponds to a conjugacy class A_p in $\operatorname{im} \bar\rho_{E,\ell} \subseteq \mathrm{GL}_2(\mathbb{Z}/\ell\mathbb{Z})$.

We have $\operatorname{tr} A_p = a_p \bmod \ell$ and $\det A_p = p \bmod \ell$, hence we know the characteristic polynomial of A_p.

By varying p, we can "randomly" sample $\operatorname{im} \bar\rho_{E,\ell}$.
The Čebotarev density theorem implies equidistribution.

Unfortunately, this does not give us enough information.

Example: $\ell = 2$

$\mathrm{GL}_2(\mathbb{Z}/2\mathbb{Z}) \simeq S_3$ has 6 subgroups in 4 conjugacy classes.
For $H \subseteq \mathrm{GL}_2(\mathbb{Z}/2\mathbb{Z})$, let $t_i(H) = \#\{A \in H : \operatorname{tr} A = i\}$.
We consider the trace frequencies $t(H) = (t_0(H), t_1(H))$.

1. For $H \simeq S_3$ we have $t(H) = (4, 2)$.
2. The subgroup $H \simeq C_3$ has $t(H) = (1, 2)$.
3. Three conjugate $H \simeq C_2$ have $t(H) = (2, 0)$
4. The trivial subgroup H has $t(H) = (1, 0)$.

1,2 are distinguished from 3,4 by a trace 1 element (easy).
We can distinguish 1 from 2 by comparing frequencies (harder).
We cannot distinguish 3 from 4 at all (impossible).

Unipotent elements are indistinguishable from the identity!

Using the fixed space of A_p

The ℓ-torsion points fixed by the Frobenius endomorphism form the $\mathbb{F}_p$-rational subgroup $E[\ell](\mathbb{F}_p)$ of $E[\ell]$. Thus

$$\ker(A_p - I) \simeq E[\ell](\mathbb{F}_p) = E(\mathbb{F}_p)[\ell].$$

For small p it is easy to compute $E(\mathbb{F}_p)[\ell]$, and this gives information about A_p that *cannot* be derived from a_p.

We can now easily distinguish all 4 subgroups of $\mathrm{GL}_2(\mathbb{Z}/\ell\mathbb{Z})$.

This generalizes nicely.

Subgroup signatures

For each subgroup H of $\mathrm{GL}_2(\mathbb{Z}/\ell\mathbb{Z})$ we define the *extended signature* of H as the multiset

$$S_H = \{(\det A, \operatorname{tr} A, \operatorname{rk}(A_p - I)) : A \in H\}.$$

The *signature* s_H is simply the set S_H, ignoring multiplicities. Note that s_H and S_H are invariant under conjugation.

Lemma

Let $\ell < 60$ be prime, and let G and H be subgroups of $\mathrm{GL}_2(\mathbb{Z}/\ell\mathbb{Z})$ for which the determinant map is surjective.

1. $s_G = s_H \Leftrightarrow S_G = S_H$.
2. $S_G = S_H \Rightarrow G \simeq H$.

The subgroup lattice of $\mathrm{GL}_2(\mathbb{Z}/\ell\mathbb{Z})$

Our strategy is to determine im $\bar{\rho}_{E,\ell}$ by identifying its location in the lattice of subgroups of $\mathrm{GL}_2(\mathbb{Z}/\ell\mathbb{Z})$.

We do not distinguish conjugate subgroups, and we restrict our attention to the upwardly closed set of subgroups $\mathcal{C}_\ell$ for which the determinant map is surjective.

For any $H \in \mathcal{C}_\ell$, we say that a set of signatures s is *minimally covered* by s_H if $s \subset s_H$ and $s \subset s_G \implies s_H \subset s_G$ for all $G \in \mathcal{C}_\ell$.

If s is minimally covered by both s_G and s_H, then $G \simeq H$, by the lemma.

The algorithm

Given an elliptic curve $E/\mathbb{Q}$, a prime ℓ, and $\epsilon > 0$,
set $s \leftarrow \{\}$, $k \leftarrow 0$, and for each prime of good reduction $p \neq \ell$:

1. Compute $E(\mathbb{F}_p)$ to obtain $a = p + 1 - \#E(\mathbb{F}_p)$ and $r = \mathrm{rk}(E(\mathbb{F}_p)[\ell])$.
2. Set $s \leftarrow s \cup (p \bmod \ell, a \bmod \ell, r)$ and increment k.
3. If s is minimally covered by s_H for some $H \in \mathcal{C}_\ell$ and $\delta_H^k < \epsilon$, output H and terminate.

Here δ_H is the maximum over $G \supsetneq H$ of the probability that the signature of a random $A \in G$ lies in s_H (zero if $H = \mathrm{GL}_2(\mathbb{Z}/\ell\mathbb{Z})$).

The values of s_H and δ_H for all $H \in \mathcal{C}_\ell$ are precomputed.

Efficient implementation

If $\bar{\rho}_{E,\ell}$ is surjective, we expect the algorithm to terminate in $O(\log \ell)$ iterations, typically less than 10 for $\ell < 60$.

Otherwise, if $\epsilon = 2^{-n}$ we expect to need $O(\log \ell + n)$ iterations, typically less than $2n$ (we use $n = 100$).

By precomputing tables of $E(\mathbb{F}_p)$ for *all* elliptic curve $E/\mathbb{F}_p$ for small values of p (up to 2^{16}, say), the algorithm is essentially just a sequence of table lookups, which makes it *very fast*.

Precomputing the s_H and δ_H is non-trivial, but this only needs to be done once for each prime ℓ.

Computational results

With $\epsilon = 2^{-100}$ it takes less than a minute to analyze all the curves in Cremona's tables for $\ell < 60$. This includes all curves $E/\mathbb{Q}$ with conductor up to 240,000 (≈1.5 million curves). For curves without CM, this yields 45 distinct signatures of non-surjective Galois images (40 isomorphism classes).

Performing the same analysis on the Stein-Watkins database ($\approx$ 137 million curves with conductors up to 10 million) takes about an hour and yields the same set of signatures.

So far we have analyzed a total of some 10^{10} curves in various families (e.g., bounded coefficients, bounded j-invariants, parametrizations of modular curves). This work is still in progress, but has so far not yielded any new signatures.

Summary of results

Non-surjective images of $\bar{\rho}_{E,\ell}$ for elliptic curves $E/\mathbb{Q}$ without complex multiplication and primes $\ell < 60$.

ℓ	isomorphism	signature	conjugacy[1]
2	3	3	3
3	6	6	7
5	10	12	15
7	10	11	16
11	3	4	7
13	6	7	11
17	1	1	2
37	1	1	2
	40	45	63

[1] There can be up to two conjugacy classes with the same signature, but these must then have equal image in $\mathrm{PGL}_2(\mathbb{Z}/\ell\mathbb{Z})$.

ℓ	GAP	ind	δ_H	$\twoheadrightarrow a_p$	$\to N_p$	isog	type	count
2	1.1	6	0.500	no	no	$(\mathbb{Z}/2\mathbb{Z})^2$	Z	133452
	2.1	3	0.500	no	no	$\mathbb{Z}/2\mathbb{Z}$	B	5281954
	3.1	2	0.333	yes	yes	no	C_{ns}	7412
3	2.1	24	0.250	no	no	$\mathbb{Z}/3\mathbb{Z}$	Z	2189
	4.2	12	0.167	yes	no	yes	C_s	1570
*	6.1	8	0.250	no	no	$\mathbb{Z}/3\mathbb{Z}$	$\subset B$	217794
	8.3	6	0.250	yes	yes	no	$N(C_s)$	205
	12.4	4	0.375	yes	no	yes	B	186668
	16.8	3	0.167	yes	yes	no	$N(C_{ns})$	816
5	4.1	120	0.200	no	no	$\mathbb{Z}/5\mathbb{Z}$	$\subset C_s$	4
	4.1	120	0.200	no	no	yes	$\subset C_s$	4
	8.2	60	0.100	yes	no	yes	$\subset C_s$	4
	16.2	30	0.050	yes	yes	yes	C_2	12
	16.6	30	0.250	yes	yes	no	$< N(C_{ns})$	3
*	20.3	24	0.375	no	no	$\mathbb{Z}/5\mathbb{Z}$	$\subset B$	504
*	20.3	24	0.375	no	no	yes	$\subset B$	520
	32.11	15	0.333	yes	yes	no	$N(C_s)$	15
*	40.12	12	0.250	yes	no	yes	$\subset B$	536
	48.5	10	0.333	yes	yes	no	$N(C_{ns})$	29
	80.30	6	0.417	yes	yes	yes	B	950
	96.67	5	0.217	yes	yes	no	$\twoheadrightarrow S_4$	284

(counts of $\bar{\mathbb{Q}}$-isomorphism classes in the Stein-Watkins database may overlap — im $\bar{\rho}_{E,\ell}$ is not twist invariant)

ℓ	GAP	ind	δ_H	$\twoheadrightarrow a_p$	$\twoheadrightarrow \#E(\mathbb{F}_p)$	isog	type	count
7	18.3	112	0.250	yes	no	no	$\subset N(C_s)$	1
	36.12	56	0.333	yes	no	no	$\subset N(C_s)$	1
*	42.4	48	0.250	no	no	yes	$\subset B$	6
*	42.1	48	0.417	no	no	$\mathbb{Z}/7\mathbb{Z}$	$\subset B$	24
*	42.1	48	0.417	no	no	yes	$\subset B$	24
	72.30	28	0.399	yes	yes	no	$N(C_s)$	1
	84.12	24	0.667	yes	no	yes	$\subset B$	6
*	84.7	24	0.444	yes	no	yes	$\subset B$	24
	96.62	21	0.357	yes	yes	no	$N(C_{ns})$	2
*	126.7	16	0.250	yes	yes	yes	$\subset B$	682
	252.28	8	0.438	yes	yes	yes	B	682
*11	110.1	120	0.450	no	no	yes	$\subset B$	2
*	110.1	120	0.450	no	no	yes	$\subset B$	2
*	220.7	60	0.640	no	no	yes	$\subset B$	2
	240.51	55	0.409	yes	yes	no	$N(C_{ns})$	0
13	288.400	91	0.250	yes	yes	no	$\twoheadrightarrow S_4$	1
*	468.29	56	0.375	yes	yes	yes	$\subset B$	14
*	468.29	56	0.375	yes	yes	yes	$\subset B$	12
	624.155	42	0.667	yes	no	yes	$\subset B$	4
*	624.119	42	0.444	yes	yes	yes	$\subset B$	2
*	936.171	28	0.250	yes	yes	yes	$\subset B$	14
	1872.576	14	0.464	yes	yes	yes	B	42
*17	1088.1674	72	0.375	yes	yes	yes	$\subset B$	2
*37	15984	114	0.444	yes	yes	yes	$\subset B$	2

Generalizations and future work

Working modulo $\ell^e = 4, 8, 9$ finds $5, 2, 1$ cases of non-surjective Galois images where the mod ℓ^{e-1} image is surjective.

Working modulo composite integers *m* is also interesting. (NB: for $E/\mathbb{Q}$ the image in $\mathrm{GL}_2(\hat{\mathbb{Z}})$ is *never* surjective – Serre).

The algorithm works equally well over number fields (we can restrict to degree-1 primes).

In principle we can handle abelian varieties of dimension $g > 1$; e.g., for the Jacobian of a genus 2 curve we can compute the image of Galois in $\mathrm{GSp}_4(\mathbb{Z}/\ell\mathbb{Z})$ for (very) small ℓ.