

Building a database of modular curves (using Magma).

Andrew V. Sutherland

Massachusetts Institute of Technology



November 27, 2023

Background and context

Last year the [Simons Collaboration in Arithmetic Geometry, Number Theory, and Computation](#) launched a project to create a database of modular curves to become part of the [L-functions and Modular Forms Database](#). Contributors include:

Nikola Adžaga, Eran Assaf, Jennifer Balakrishnan, Barinder Banwait, Raymond van Bommel, Shiva Chidambaram, Garen Chiloyan, Edgar Costa, Alex Cowan, Juanita Duque-Rosero, Noam Elkies, Sachi Hashimoto, Daniel Hast, Aashraya Jha, Timo Keller, Jean Kieffer, David Lowry-Duda, Alvaro Lozano-Robledo, Kimball Martin, Pietro Mercuri, Philippe Michaud-Jacobs, Grant Molnar, Steffen Müller, Filip Najman, Ekin Ozman, Oana Padurariu, Bjorn Poonen, David Roe, Rakvi, Jeremy Rouse, Ciaran Schembri, Padmavathi Srinivasan, Sam Schiavone, Bianca Viray, John Voight, Borna Vukorepa, and David Zywina.

This project has many disparate components, but they are all tied together by a common group-theoretic scaffold, which is inspired by Mazur's [Program B](#).

Mazur's 1976 lectures on *Rational points on modular curves*

In the course of preparing my lectures for this conference, I found a proof of the following theorem, conjectured by Ogg (conjecture 1 [17b]):

THEOREM 1. Let ϕ be the torsion subgroup of the Mordell-Weil group of an elliptic curve E , over \mathbb{Q} . Then ϕ is isomorphic to one of the following 15 groups:

$$\begin{aligned} & \mathbb{Z}/m \cdot \mathbb{Z} && \text{for } m \leq 10 \text{ or } m = 12 \\ & \mathbb{Z}/2 \cdot \mathbb{Z} \times \mathbb{Z}/2\nu \cdot \mathbb{Z} && \text{for } \nu \leq 4 . \\ & \vdots \end{aligned}$$

Theorem 1 also fits into a general program:

B. Given a number field K and a subgroup H of $GL_2 \widehat{\mathbb{Z}} = \prod_p GL_2 \mathbb{Z}_p$ classify all elliptic curves E/K whose associated Galois representation on torsion points maps $\text{Gal}(\overline{K}/K)$ into $H \subset GL_2 \widehat{\mathbb{Z}}$.

Galois representations attached to elliptic curves

Let E be an elliptic curve over a number field k . The action of Gal_k on $E[N]$ yields

$$\rho_{E,N}: \text{Gal}_k \rightarrow \text{Aut}(E[N]) \simeq \text{GL}_2(\mathbb{Z}/N\mathbb{Z}) =: \text{GL}_2(N).$$

Choosing a compatible system of bases and taking the inverse limit yields

$$\rho_E: \text{Gal}_k \rightarrow \varprojlim \text{GL}_2(N) \simeq \text{GL}_2(\widehat{\mathbb{Z}}) \simeq \prod \text{GL}_2(\mathbb{Z}_\ell).$$

Note that ρ_E and its image are defined only up to GL_2 -conjugacy.

In this talk **we will always work up to GL_2 -conjugacy**.

Theorem (Serre 1972)

If E/k is a non-CM elliptic curve then $\rho_E(\text{Gal}_k)$ is an open subgroup of $\text{GL}_2(\widehat{\mathbb{Z}})$.

When $k = \mathbb{Q}$ the index $[\text{GL}_2(\widehat{\mathbb{Z}}) : \rho_E(\text{Gal}_k)]$ is divisible by 2.

For any fixed k one expects the **index** $[\text{GL}_2(\widehat{\mathbb{Z}}) : \rho_E(\text{Gal}_k)]$ to be bounded for non-CM E/k .

For $k = \mathbb{Q}$ the bound 2736 has been conjectured (see [Zywina 2022](#)).

The modular curve X_H

Definition (Deligne, Rapoport 1973)

For each open $H \leq \mathrm{GL}_2(\widehat{\mathbb{Z}})$. The **modular curves** X_H and Y_H are coarse spaces for the stacks \mathcal{M}_H and \mathcal{M}_H^0 parametrizing elliptic curves E with **H -level structure**: equivalence classes $[\iota]_H$ of isomorphisms $\iota: E[N] \xrightarrow{\sim} \mathbb{Z}(N)^2$, where $\iota \sim \iota'$ if $\iota = h \circ \iota'$ for some $h \in H$.

- X_H is a smooth proper $\mathbb{Z}[\frac{1}{N}]$ -scheme with open subscheme Y_H .
The complement X_H^∞ of Y_H in X_H (the **cusps**) is finite étale over $\mathbb{Z}[\frac{1}{N}]$.
- If $\det(H) = \widehat{\mathbb{Z}}^\times$ the generic fiber of X_H is a nice curve X_H/\mathbb{Q} , and $X_H(\mathbb{C})$ is the Riemann surface $X_{\Gamma_H} := \Gamma_H \backslash \mathcal{H}$, with $\Gamma_H \subseteq \mathrm{SL}_2(\mathbb{Z})$ the preimage of $\pi_N(H) \cap \mathrm{SL}_2(N)$.
If $\det(H) \neq \widehat{\mathbb{Z}}^\times$ then X_H is not geometrically connected, but it is a curve over \mathbb{Q} .
- For E/k with $j(E) \neq 0, 1728$ we have $\rho_{E,N}(\mathrm{Gal}_k) \leq H \iff (E, [\iota]_H) \in Y_H(k)$.

Subgroup inclusions $H \leq H'$ induce morphisms $X_H \rightarrow X_{H'}$.

In particular, every X_H is equipped with a map $j: X_H \rightarrow X(1)$ to the **j -line** $X(1) \simeq \mathbb{P}^1$.

Three fundamental invariants: level, index, genus

For each (conjugacy class of) open $H \leq \mathrm{GL}_2(\widehat{\mathbb{Z}})$ we define the following invariants.

- the **level** $n(H)$ is the least N for which H contains the kernel of $\mathrm{GL}_2(\widehat{\mathbb{Z}}) \rightarrow \mathrm{GL}_2(N)$.
- the **index** $i(H)$ is the positive integer $[\mathrm{GL}_2(\widehat{\mathbb{Z}}) : H] = [\mathrm{GL}_2(N) : H(N)]$.
- the **genus** $g(H)$ is the nonnegative integer

$$g(H) := g(\Gamma) := 1 + \frac{i(\Gamma)}{12} - \frac{e_2(\Gamma)}{4} - \frac{e_3(\Gamma)}{3} - \frac{e_\infty(\Gamma)}{2} \quad (\Gamma := \pm H(N) \cap \mathrm{SL}_2(N)),$$

where $i(\Gamma) := [\mathrm{SL}_2(N) : \Gamma]$ counts right Γ -cosets in $\mathrm{SL}_2(N)$, e_2 and e_3 count cosets fixed by $\begin{pmatrix} 0 & 1 \\ -1 & -1 \end{pmatrix}$ and $\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$, respectively, and $e_\infty(\Gamma)$ counts $\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ -orbits of $\Gamma \backslash \mathrm{SL}_2(N)$.

When $\det(H) = \widehat{\mathbb{Z}}^\times$ and $-I \in H$, the level $n(H)$ controls the bad primes of X_H , the index $i(H)$ is the degree of the map $X_H \rightarrow X(1)$, and $g(H)$ is the genus of X_H/\mathbb{Q} .

If $H' \leq H$ then $n(H) | n(H')$ and $i(H) | i(H')$ and $g(H) \leq g(H')$.

Coarse and fine subgroups

Definition

Open $H \leq \mathrm{GL}_2(\widehat{\mathbb{Z}})$ that contain $-I$ are **coarse groups**; those that do not are **fine groups**. A **quadratic refinement** of a coarse group H is a fine group H' for which $H = \pm H'$.

A typical coarse H has infinitely many quadratic refinements H' , all of which satisfy:

- $n(H) | n(H')$, $i(H') = 2i(H)$, $g(H') = g(H)$.
- $X_{H'} \simeq X_H$ (as curves); in particular $L(X_{H'}, s) = L(X_H, s)$ and $X_{H'}(k) \leftrightarrow X_H(k)$.
- $j(X_{H'}(k)) = j(X_H(k))$ for every k/\mathbb{Q} .

If H' is a quadratic refinement of H and E/k has Galois image $\rho_E(\mathrm{Gal}_k) = H$, the quadratic twist \tilde{E}/k by the fixed field of $\rho_E^{-1}(H')$ has Galois image $\rho_{\tilde{E}}(\mathrm{Gal}_k) = H'$.

Example

The elliptic curve $14.a4$ corresponds to a point on $X_1(3)$, a quadratic refinement of $X_0(3)$. Every **twist** has a rational 3-isogeny, but only $14.a4$ has a rational 3-torsion point.

The determinant map

For E/k the composition $\det \circ \rho_E: \text{Gal}_k \rightarrow \widehat{\mathbb{Z}}^\times$ factors through $\text{Gal}(k^{\text{cyc}}/k)$.

For E/\mathbb{Q} we have $\det \circ \rho_E = \chi_{\text{cyc}}$, where $\chi_{\text{cyc}}: \text{Gal}_{\mathbb{Q}} \rightarrow \widehat{\mathbb{Z}}^\times$ is the cyclotomic character.

For E/k the image $\rho_E(\text{Gal}_k)$ lies in the subgroup $\det^{-1}(\chi_{\text{cyc}}(\text{Gal}_k))$ of index $[k \cap \mathbb{Q}^{\text{cyc}} : \mathbb{Q}]$.

For E/\mathbb{Q} the Kronecker-Weber theorem implies that if $H_E := \rho_E(\text{Gal}_{\mathbb{Q}})$ then

$$[H_E, H_E] = H_E \cap \text{SL}_2(\widehat{\mathbb{Z}})$$

which is a non-trivial constraint: for most $H \in \text{GL}_2(\widehat{\mathbb{Z}})$ we have $[H, H] < H \cap \text{SL}_2(\widehat{\mathbb{Z}})$.

If E/k has image $H_E := \rho_E(\text{Gal}_k)$ then its base change to k^{cyc} has image $H_E \cap \text{SL}_2(\widehat{\mathbb{Z}})$.

If $\Gamma_H := H \cap \text{SL}_2(\widehat{\mathbb{Z}}) = H' \cap \text{SL}_2(\widehat{\mathbb{Z}}) =: \Gamma_{H'}$ then $X_H/\mathbb{Q}^{\text{cyc}} \simeq X_{H'}/\mathbb{Q}^{\text{cyc}}$.

If $H(N) \cap \text{SL}_2(N) = H'(N) \cap \text{SL}_2(N)$ with $n(H), n(H') | N$ then $X_H/\mathbb{Q}(\zeta_N) \simeq X_{H'}/\mathbb{Q}(\zeta_N)$.

Subgroups of $GL_2(\widehat{\mathbb{Z}})$ vs subgroups of $SL_2(\widehat{\mathbb{Z}})$

For any fixed g there are only finitely many open $\Gamma \leq SL_2(\widehat{\mathbb{Z}})$ containing $-I$ with $g(\Gamma) = g$. You can find complete lists for $g \leq 24$ in the [Cummins–Pauli database](#).¹

By contrast, $GL_2(\widehat{\mathbb{Z}})$ contains infinitely many coarse subgroups of every genus.

For open $H \leq GL_2(\widehat{\mathbb{Z}})$ with $\det(H) = \widehat{\mathbb{Z}}^\times$, the index and genus of H depend only on $\Gamma := H \cap SL_2(\widehat{\mathbb{Z}})$, but the levels of H and Γ may differ.

For distinct H, H' of the same level N with common intersection in $SL_2(N)$, the curves $X_H, X_{H'}$ are not isomorphic. They typically have non-isogenous Jacobians and different sets of rational points (in particular, one may be empty when the other is not!).

Example

For the groups $H = 15.60.2.c.1$ and $15.60.2.d.1$, $H \cap SL_2(\widehat{\mathbb{Z}})$ has CP label $15D^2$.

The first X_H has no \mathbb{Q} -points and rank 1 $\text{Jac}(X_H) \sim 75.c \times 225.c$.

The second $X_H = X_{ns}^+(15)$ has 6 rational \mathbb{Q} -points and rank 2 $\text{Jac}(X_H) \sim 225.a \times 225.c$.

¹Cummins and Pauli consider Γ up to $GL_2(\mathbb{Z})$ -conjugacy, not $GL_2(\widehat{\mathbb{Z}})$ -conjugacy.

Rational points on X_H

Let H be an open subgroup of $\mathrm{GL}_2(\widehat{\mathbb{Z}})$ of level N (which we may view as $H \leq \mathrm{GL}_2(N)$).

Definition

The set $Y_H(\bar{k})$ consists of equivalence classes $(E, [\iota]_H)$, where $(E, [\iota]_H) \sim (E', [\iota']_H)$ if there is an isomorphism $\phi: E \rightarrow E'$ for which $\phi_N: E[N] \rightarrow E'[N]$ satisfies $\iota \sim \iota' \circ \phi_N$.

Each $\sigma \in \mathrm{Gal}_K$ induces $\sigma^{-1}: E^\sigma[N] \xrightarrow{\sim} E[N]$ via $(x : y : z) \mapsto (\sigma^{-1}(x) : \sigma^{-1}(y) : \sigma^{-1}(z))$.

We have a Gal_k -action on $Y_H(\bar{k})$: $(E, [\iota]_H) \mapsto (E^\sigma, [\iota \circ \sigma^{-1}]_H)$, and define $Y_H(k) := Y_H(\bar{k})^{\mathrm{Gal}_k}$.

Equivalently, $Y_H(\bar{k})$ is the set of pairs $(j(E), \alpha)$, with $\alpha = Hg \mathrm{Aut}(E_{\bar{k}}) \in H \backslash \mathrm{GL}_2 / \mathrm{Aut}(E_{\bar{k}})$, on which Gal_k acts via $(j(E), \alpha) \mapsto (j(E)^\sigma, \alpha^\sigma)$, where $\alpha^\sigma = Hg\rho_E(\sigma) \mathrm{Aut}(E_{\bar{k}})$.

Gal_k acts on $X_H^\infty(\bar{k}) := \pm H \backslash \mathrm{GL}_2 / \langle \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \rangle$ via $\left(\begin{smallmatrix} \chi_{\mathrm{cyc}}(\sigma) & 0 \\ 0 & 1 \end{smallmatrix} \right)$, and $X_H^\infty(k) := X_H^\infty(\bar{k})^{\mathrm{Gal}_k}$.

We now define $X_H(\bar{k}) := Y_H(\bar{k}) \sqcup X_H^\infty(\bar{k})$, and $X_H(k) := X_H(\bar{k})^{\mathrm{Gal}_k} = Y_H(k) \sqcup X_H^\infty(k)$.

Counting \mathbb{F}_q -points on X_H

Theorem (Duke, Tóth 2002)

Let E/\mathbb{F}_q be an elliptic curve, and let π_E denote its Frobenius endomorphism. Define $a := \text{tr } \pi_E = q + 1 - \#E(\mathbb{F}_q)$ and $R := \text{End}(E) \cap \mathbb{Q}(\pi_E)$, let $\Delta := \text{disc}(R)$ and $\delta := \Delta \bmod 4$, and let $b := \sqrt{(a^2 - 4q)/\Delta}$ if $\Delta \neq 1$ and $b := 0$ otherwise. The integer matrix

$$A_E := \begin{pmatrix} (a + b\delta)/2 & b \\ b(\Delta - \delta)/4 & (a - b\delta)/2 \end{pmatrix}$$

gives the action of π_E on $E[N]$ for all $N \geq 1$.

We compute $A_E = A(t, v, d)$ for all E/\mathbb{F}_q by enumerating solutions (t, v, D) to

$$4q = t^2 - v^2D,$$

and making appropriate adjustments for $j(E) = 0, 1728$ and supersingular E/\mathbb{F}_q . We then count the double cosets fixed by $A(t, v, d)$ with multiplicity $h(D)$.

The algorithm

Given $H \leq \mathrm{GL}_2(N)$ containing $-I$ and a prime power q , compute $X_H(\mathbb{F}_q)$ as follows:

- 1 Compute the **permutation character** $\chi_H: \mathrm{GL}_2(N) \rightarrow \mathbb{Z}$ counting H -cosets fixed by g , which is equal to $[\mathrm{GL}_2(N) : H] \#(H \cap [g]) / \#[g]$ where $[g]$ is the conjugacy class of g .
- 2 Compute $n_\infty := \#X_H^\infty(\mathbb{F}_q)$ by counting elements of $H \backslash \mathrm{GL}_2(N) / \langle \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \rangle$ fixed by $\begin{pmatrix} q & 0 \\ 0 & 1 \end{pmatrix}$.
- 3 Compute $n_0 := \#j_H^{-1}(0)$ and $n_{1728} := \#j_H^{-1}(1728)$ by computing A_π for each twist, summing $\chi_H(A_\pi)$ values, and dividing by $\# \mathrm{Aut}(E_{\bar{k}})$.
- 4 Compute $n_{\mathrm{ord}} := \sum_{t,v,D} \chi_H(A(t,v,D)) h(D)$ with (t,v,D) varying over solutions to $4q = t^2 - v^2 D$ with $t \perp q$ and $D < -4$.
- 5 Similarly compute n_{ss} similarly (omitting $j(E) = 0, 1728$; see [RSZB22] for details).
- 6 Output $\#X_H(\mathbb{F}_q) = n_\infty + n_0 + n_{1728} + n_{\mathrm{ord}} + n_{\mathrm{ss}}$.

As written the running time of this algorithm is $\tilde{O}(N^3) + \tilde{O}(\sqrt{q})$.

The $\tilde{O}(N^3)$ term is independent of q and can be improved to $\tilde{O}(N^{5/2})$ (possibly $\tilde{O}(N^2)$).

Performance comparison

Time to compute $\#X_0(N)(\mathbb{F}_p)$ for all primes $p \leq B$ in seconds.

B	trace formula in Pari/GP v2.11				point-counting via moduli			
	$N = 41$	42	209	210	$N = 41$	42	209	210
2^{12}	0.1	0.4	0.2	0.7	0.0	0.0	0.0	0.0
2^{13}	0.3	1.0	0.5	1.8	0.0	0.0	0.1	0.0
2^{14}	0.6	2.5	1.1	4.8	0.1	0.1	0.1	0.1
2^{15}	1.7	7.1	3.1	12.8	0.2	0.2	0.2	0.2
2^{16}	4.8	19.6	8.9	35.4	0.4	0.4	0.6	0.5
2^{17}	14.4	55.1	25.7	97.8	1.1	0.9	1.5	1.2
2^{18}	43.5	156	74.3	274	2.8	2.6	4.0	3.3
2^{19}	128	442	214	769	7.8	7.0	11.0	9.1
2^{20}	374	1260	610	2169	22.2	19.8	31.1	26.2
2^{21}	1100	3610	1760	6100	69.0	61.3	91.8	77.9
2^{22}	?	?	?	?	213	187	263	228
2^{23}	?	?	?	?	665	579	762	678
2^{24}	?	?	?	?	2060	1790	2220	1990

(? = did not complete within one day; the genus of $X_0(N)$ is 3, 5, 19, 41 for $N = 41, 42, 209, 210$)

Decomposing the Jacobian of X_H

Let H be an open subgroup of $\mathrm{GL}_2(\widehat{\mathbb{Z}})$ of level N and let J_H denote the Jacobian of X_H .

Theorem (Rouse, S, Voight, Zureick-Brown 2021)

Each simple factor of J_H is isogenous to A_f for a weight-2 eigenform f on $\Gamma_0(N^2) \cap \Gamma_1(N)$.

If we know the q -expansions of the eigenforms in $S_2(\Gamma_0(N^2) \cap \Gamma_1(N))$ we can uniquely determine the decomposition of J_H up to isogeny using linear algebra and point-counting.

It suffices to work with **trace forms** $\mathrm{Tr}(f)$ (the sum of the Galois conjugates of f)

$$\mathrm{Tr}(f)(q) := \sum_{n=1}^{\infty} \mathrm{Tr}_{\mathbb{Q}(f)/\mathbb{Q}}(a_n(f))q^n,$$

since the integers $a_n(\mathrm{Tr}(f))$ uniquely determine $L(A_f, s)$ and the isogeny class of A_f .

By strong multiplicity one (**Soundararajan 2004**), the $a_p(\mathrm{Tr}(f))$ for enough $p \nmid N$ suffice.

Decomposing the Jacobian of X_H

Let $\{[f_1], \dots, [f_m]\}$ be the Galois orbits of the weight-2 eigenforms for $\Gamma_0(N^2) \cap \Gamma_1(N)$. Then

$$L(J_H, s) = \prod_{i=1}^m L(A_{f_i}, s)^{e_i}$$

for some unique vector of nonnegative integers $e(H) := (e_1, \dots, e_m)$.

Let $T(B) \in \mathbb{Z}^{n \times m}$ have columns $[a_1(\text{Tr}(f_i)), a_2(\text{Tr}(f_i)), \dots, a_p(\text{Tr}(f_i)), \dots]$ for good $p \leq B$.
Let $a(H; B) := [g(H), a_2(H), \dots, a_p(H), \dots]$, where $a_p(H)p + 1 - \#X_H(\mathbb{F}_p)$, for good $p \leq B$.

For all sufficiently large B the \mathbb{Q} -linear system

$$T(B)x = a(H; B),$$

has the unique solution $x = e(H)$.

We can then compute the analytic rank of J_H as $\text{rk}(J_H) = \sum e_i \text{rk}(f_i)$ using the [LMFDB](#).

Gassmann classes

For subgroups H_1 and H_2 of a finite group G the following are equivalent:

- $\#(H_1 \cap C) = \#(H_2 \cap C)$ for every conjugacy class $C \subseteq G$.
- There is a conjugacy-class-preserving bijection of sets $H_1 \leftrightarrow H_2$.
- The permutation characters $\chi_{H_1}: G \rightarrow \mathbb{Z}$ and $\chi_{H_2}: G \rightarrow \mathbb{Z}$ coincide.
- The G -sets $[H_1 \backslash G]$ and $[H_2 \backslash G]$ are isomorphic as K -sets for every cyclic $K \leq G$.
- The permutation modules $\mathbb{Q}[H_1 \backslash G]$ and $\mathbb{Q}[H_2 \backslash G]$ are isomorphic as $\mathbb{Q}[G]$ -modules.

Subgroups that satisfy any of these equivalent conditions are **Gassmann equivalent**.²

Open $H_1, H_2 \leq \mathrm{GL}_2(\widehat{\mathbb{Z}})$ are Gassmann equivalent if $H_1(N), H_2(N) \leq \mathrm{GL}_2(N)$ are Gassmann equivalent for any N divisible by the levels of H_1 and H_2 .

Proposition

For Gassmann equivalent $H_1, H_2 \leq \mathrm{GL}_2(\widehat{\mathbb{Z}})$ we have $\mathrm{Jac}(X_{H_1}) \sim \mathrm{Jac}(X_{H_2})$.

²See [S21] for more on arithmetic equivalence.

Labels

Coarse groups $H \leq \mathrm{GL}_2(\widehat{\mathbb{Z}})$ with $\det(H) = \widehat{\mathbb{Z}}^\times$ have labels of the form **N.i.g.c.n**:

- N, i, g are the level, index, genus of H , respectively;
- c identifies the Gassmann class of H among those with label prefix **N.i.g**;
- n identifies the conjugacy class of H for those with label prefix **N.i.g.c**.

Fine groups $H \leq \mathrm{GL}_2(\widehat{\mathbb{Z}})$ with $\det(H) = \widehat{\mathbb{Z}}^\times$ have labels of the form **N.i.g-M.c.m.n**:

- N, i, g are the level, index, genus of H , respectively;
- M, c, m are components of the label **M.j.g.c.m** of $\pm H$;
- n identifies the conjugacy class of H for those with label prefix **N.i.g-M.c.m**.

Gassmann classes are ordered by lexicographically sorting characters via their values on conjugacy classes of elements ordered by **similarity invariant**.

Conjugacy classes of subgroups are ordered by their **canonical generators**.

These also play a key role in our algorithm for enumerating open subgroups of $\mathrm{GL}_2(\widehat{\mathbb{Z}})$.

Similarity invariants

Let p^e be prime power. Each $A \in M_2(p^e)$ is similar³ to a matrix of the form

$$zI + p^j \begin{pmatrix} 0 & 1 \\ -d & t \end{pmatrix},$$

where the tuple of integers $\text{inv}(A) := (j, z, d, t)$ is uniquely determined by

- $j \leq e$ is the largest integer such that $A \bmod p^j$ is a scalar matrix;
- $z \in [0, p^j - 1]$ satisfies $zI = A \bmod p^j$.
- $d, t \in [0, p^{e-j} - 1]$ satisfy $d = \det p^{-j}(A - zI)$ and $t = \text{tr } p^{-j}(A - zI)$.

We extend this to general moduli $N = p_1^{e_1} \cdots p_n^{e_n}$ with $p_1 < \cdots < p_n$ prime via

$$\text{inv}(A) := (\text{inv}(A \bmod p_1^{e_1}), \dots, \text{inv}(A \bmod p_n^{e_n})).$$

Lemma

Matrices $A, B \in \text{GL}_2(N)$ are conjugate if and only if $\text{inv}(A) = \text{inv}(B)$.

³ A and B are similar if $EA = BE$ for some $E \in \text{GL}_2(p^e)$. See [AOPV09] for a proof of the claims above.

Canonical generators

Given an open $H \leq \mathrm{GL}_2(\widehat{\mathbb{Z}})$ we wish to choose a representative of the conjugacy class $[H]$ that H represents, and generators for it in a way that depends only on $[H]$.

We first fix an ordering of $\mathrm{GL}_2(N)$ -conjugacy classes $[g]$ (rather than sorting by similarity invariant it is better to sort by decreasing $|g|$, decreasing $\#[g]$, then by similarity invariant).

The **canonical generators** for coarse $H \leq \mathrm{GL}_2(\widehat{\mathbb{Z}})$ of level N are the lexicographically minimal sequence $h_1, \dots, h_n \in \mathrm{GL}_2(N)$ such that

- $H(N) \cap \mathrm{SL}_2(N) = \langle h_1, \dots, h_m \rangle$ for some $m \leq n$ and $H(N) = \langle h_1, \dots, h_n \rangle$.
- $\langle h_1, \dots, h_i \rangle < \langle h_1, \dots, h_{i+1} \rangle$ for $1 \leq i < n$;
- $[h_1], \dots, [h_m]$ and $[h_{m+1}], \dots, [h_n]$ are nondecreasing (under our fixed ordering);

The **canonical generators** for fine $H \leq \mathrm{GL}_2(\widehat{\mathbb{Z}})$ are the sequence $\varepsilon_1 h_1, \dots, \varepsilon_n h_n$ where h_1, \dots, h_n are canonical generators for $\pm H$ and $\varepsilon_1, \dots, \varepsilon_n \in \{\pm 1\}^n$ minimize $\sum_{\varepsilon_i=1} 2^{i-1}$.

Computing canonical generators

Given $H \leq \mathrm{SL}_2(N)$ proceed as follows:

- 1 Replace N with N/p if $H \bmod N/p$ has the same index as $H \bmod N$ until $N = n(H)$.
- 2 Compute the ordered list L of primitive similarity classes C supporting $\chi_H: G \rightarrow \mathbb{Z}$.
- 3 Let $K = \langle h \rangle$ where $h = \min(C)$ for the first $C \in L$ (if $\#H = \#K$ return h).
- 4 Compute a right transversal T of $N(H)$ in $\mathrm{GL}_2(n)$ and remove $g \in T$ for which $h \notin H^g$.
- 5 For each $C \in L$:
 - a Replace C with $\{c \in C : c \in H^g \text{ some } g \in T \text{ and } c \notin K\}$.
If C is empty proceed to the next C .
 - b Replace K with $\langle K, h \rangle$. If $\#K = \#H$ then return generators for K .
 - c Remove $g \in T$ for which $K \not\subset H^g$.

For $H \in \mathrm{GL}_2(N)$ we apply this to $H \cap \mathrm{SL}_2(N)$, lift to $\mathrm{GL}_2(n(H))$, and continue from there.

The computation of the list of primitive similarity classes (and their intersections with H) can be optimized, as can the computations of T and of $\#K$ and $\#H$.

Subgroup lattice enumeration

- 1 Compute canonical generators for $GL_2(N)$, let $V_0^c = (GL_2(N))$, $V_0^f = \emptyset$, and $i = 0$.
- 2 Compute V_{i+1}^c , V_{i+1}^f , and E_{i+1}^c as follows:
 - a For each $H \in V_i^c$ compute the maximal subgroups $H' < H$ with $\det(K) = \widehat{\mathbb{Z}}^\times$.
 - b Compute signs ε_i for each fine maximal $F < H$ and compute canonical generators.
 - c Add distinct F to V_{i+1}^f along with generators for $F \cap K$ for each coarse maximal $K < H$.
 - d Add coarse maximal $K < H$ to V_{i+1}^c and coarse edges (K, H) to V_{i+1}^c .
- 3 Compute canonical generators for $H \in V_{i+1}^c$, remove duplicates, update E_{i+1}^c .
- 4 Increment i and return to step 2 if V_i^c is nonempty.
- 5 Compute E^f using signs from 2b and intersections from 2c, group by coarse parent.
- 6 Output $V^c := \bigcup_i V_i^c$, $V^f := \bigcup_i V_i^f$, $E^c := \bigcup_i E_i^c$, and E^f .

Steps 2, 3, 5 are designed to be highly parallelizable.

This description omits many details (conjugators, level-lifting, hashing, etc...).

Subgroup lattice enumeration timings

N	coarse		fine		Magma	New alg (threads)			
	groups	edges	groups	edges		1	2	4	8
2	4	4	0	0	0.0	0.0	0.5	0.5	0.5
3	6	6	3	2	0.0	0.1	1.0	1.0	1.0
4	22	41	21	30	0.2	0.2	1.5	1.5	1.6
5	13	19	6	4	0.1	0.2	1.3	1.3	1.3
6	44	104	26	56	0.4	0.3	1.9	1.9	2.0
7	14	20	13	18	0.1	0.1	1.3	1.3	1.4
8	285	964	981	4764	939.6	3.4	4.6	4.0	3.9
9	48	97	52	104	6.5	0.5	2.1	2.0	2.1
10	98	280	48	104	1.8	0.8	2.4	2.4	2.4
11	21	34	20	29	0.3	0.2	1.4	1.4	1.5
12	767	3030	2064	9710	4066.1	13.2	9.6	6.6	5.3
13	30	58	24	34	0.9	0.4	1.9	2.0	2.1
14	117	326	127	375	11.0	1.7	3.0	2.6	2.7
15	235	649	360	910	211.3	5.4	5.3	4.1	3.7
16	1737	7000	8317	46944	256112.2	60.8	36.4	21.0	13.9

Modular curves X_H/\mathbb{Q} of level $N < 120$

level	coarse X_H/\mathbb{Q}	fine X_H/\mathbb{Q}	X_H/\mathbb{Q}
80	1290003	17680972	18970975
96	1092407	16341502	17433909
112	731883	9785095	10516978
72	296277	3702663	3998940
48	260576	3130003	3390579
104	270822	2728804	2999626
84	243511	2190301	2433812
88	141944	1659920	1801864
60	230103	1529643	1759746
56	97836	1120338	1218174
40	93630	869934	963564
24	35951	353577	389528
⋮	⋮	⋮	⋮
	≈ 5 million	≈ 62 million	≈ 67 million

Modular curves X_H/\mathbb{Q} of level $N \leq 400$ and genus $g \leq 24$

level	coarse X_H/\mathbb{Q}	fine X_H/\mathbb{Q}	X_H/\mathbb{Q}
240	275 184	5 113 941	5 389 125
120	251 423	2 938 971	3 190 394
336	233 684	4 367 741	4 601 425
168	161 247	2 499 153	2 660 400
312	157 819	2 188 045	2 345 864
264	148 031	2 140 707	2 288 738
280	82 433	947 340	1 029 773
48	43 910	486 297	530 207
360	28 184	455 652	483 836
24	23 102	210 057	233 159
⋮	⋮	⋮	⋮
	≈ 2 million	≈ 23 million	≈ 25 million

Coarse modular curves X_H/\mathbb{Q} of level $N \leq 70$ and genus $g \leq 24$



