Census-taking for curves over finite fields

Kiran S. Kedlaya
(with Yongyuan Huang and Jun Bo Lau)

Department of Mathematics, University of California San Diego*
kedlaya@ucsd.edu
http://kskedlaya.org/slides/

Curves, Abelian VArieties and RElated Topics
Universitat de Barcelona
June 19, 2024

Supported by NSF (grant DMS-1802161) and UC San Diego (Warschawski Professorship), and during the 2023–2024 academic year by IAS and the Simons Foundation.

*The UC San Diego campus occupies unceded ancestral homelands of the Kumeyaay Nation. The Kumeyaay people continue to have an important and thriving presence in the region.

Contents

1. Statement of the problem
2. Motivation: Inverting the zeta function function
3. Motivation: Cohomology of moduli spaces of curves
4. Enumerating curves by Brill–Noether type, I
5. Brill–Noether stratifications
6. Computing points on group quotients

Moduli spaces of curves

For $g > 1$, let $\mathcal{M}_g$ denote the moduli space of smooth curves of genus g; this object exists as a smooth Deligne–Mumford (DM) stack of relative dimension $3g - 3$ over $\mathbb{Z}$.

For any field k, the set $\mathcal{M}_g(k) = \mathrm{Mor}(\mathrm{Spec}(k), \mathcal{M}_g)$ is naturally identified with the set of isomorphism classes of classical[†] curves of genus g over k.

In particular, if k is finite then this set is also finite. However, each element has an automorphism group: the class of a curve C carries the group $\mathrm{Aut}_k(C)$ of automorphisms[‡] of $C \to k$. When we write $\#\mathcal{M}_g(k)$, we count $[C]$ with weight $1/\#\,\mathrm{Aut}_k(C)$.

[†]smooth, projective, geometrically irreducible

[‡]Not to be confused with automorphisms of $C_{\overline{k}}$, which may not be defined over k.

The census problem

Problem

For given q and g, enumerate the set $\mathcal{M}_g(\mathbb{F}_q)$ (or more precisely, one curve representing each isomorphism class).

Given an explicit map $f : S \to \mathcal{M}_g(\mathbb{F}_q)$, in Magma one can easily:[§]

- verify that f is injective;
- compute $\mathrm{Aut}_k(f(s))$ for each $s \in S$.

Thus f witnesses a **lower** bound for $\#\mathcal{M}_g(\mathbb{F}_q)$. If one has an independent **upper** bound, or better yet the exact value (more on this below), these together provide a proof that f is a bijection!

This creates a dependence on closed-source software which it would be desirable to remove (e.g., by porting to SageMath).

[§] At least after applying some bugfixes that Magma has so far refused to accept.

Progress report

A census has been completed in the following cases (all data in LMFDB):

- $g \leq 3$, various q by Sutherland.
- $g = 4, q = 2$ by Xarles.
- $g = 5, q = 2$ by Dragutinović.
- $g = 6, q = 2$ by K–Huang–Lau.

Additional cases of some interest:

- $g = 7, q = 2$: in progress by K–Huang–Lau.
- $g = 2, q \in \{243, 256, 343, 512, 625, 729, 1024\}$; $g = 3, q \in \{3, \ldots, 25\}$; $g = 4, q \in \{3, 4, 5\}$; $g = 5, q = 3$: would be of interest for inverting the zeta function in LMFDB.
- $g = 6, q = 3$: would be of interest for cohomology of moduli spaces (see below). Beware that $3^{15} \gg 2^{18}$!

Contents

1 Statement of the problem

2 Motivation: Inverting the zeta function function

3 Motivation: Cohomology of moduli spaces of curves

4 Enumerating curves by Brill–Noether type, I

5 Brill–Noether stratifications

6 Computing points on group quotients

Inverting the zeta function function

Taking the numerator of the zeta function defines a map $\mathcal{M}_g(\mathbb{F}_q) \to \mathbb{Z}[T]$. It is in general difficult to compute preimages of this map.

When the preimage is **empty**, one can sometimes detect this if the zeta function implies nonsense about point counts (e.g., there is a negative number of places of some degree) or polarizations. Very rarely, one can use polarizations to show that a particular curve is unique for its zeta function (work of Howe).

In general the only available approach is a partial census, perhaps making geometric assumptions consistent with the candidate zeta function(s).

Invariants of abelian varieties

One can stratify moduli spaces of abelian varieties in positive characteristic in various ways (Newton polygon, Ekedahl–Oort). But which strata contain Jacobians?

One way to gather evidence is to compute $\mathcal{M}_g(\mathbb{F}_q)$. For example, all Newton polygons occur for curves of genus 6 over $\mathbb{F}_2$. Maybe by looking at examples, one can find some patterns that extend? (Perhaps ML/AI has a role to play here, but only if we can generate a lot more data...)

The gonality of curves over finite fields

Over an algebraically closed field, a curve of genus g always has gonality at most g. Over a finite field this can increase, but only to $g+1$ (Schmidt).

Theorem (Faber–Grantham–Howe)

A curve of genus g over $\mathbb{F}_q$ has gonality at most g except for:

- 215 *curves with* $g = 3$, $q \in \{2, 3, 4, 5, 7, 8, 9, 11, 13, 17, 19, 23, 29, 32\}$;
- 2 *curves with* $g = 4$, $q \in \{2, 3\}$.

By Riemann–Roch, a curve with gonality $g+1$ cannot have an effective divisor of degree $g-2$. This limits the zeta functions to a finite set; but a partial census was required to rule out some cases with $g \in \{6, 7\}$ (plus a theoretical argument for one case with $g = 9$).

The relative class number one problem for function fields

Theorem (K)

The set of isomorphism classes of pairs (F', F), where F, F' are function fields of curves over finite fields of genera $g < g'$ such that $h_F = h_{F'}$ and F embeds into F', is finite and known (it contains 145 elements). Moreover, if $g > 1$, then F'/F is Galois and cyclic.

Using ideas from bounding $\#C(\mathbb{F}_q)$ for C a curve of genus g over $\mathbb{F}_q$ (as in Serre's 1985 Harvard course), one shows that the zeta functions of F, F' are limited to a finite set with $g \leq 7$. However, one needs at least a partial census to identify curves with a specified zeta function. For example,

$$\mathbb{F}_2(x)[y,z]/(y^2 + (x^3 + x^2 + 1)y + x^2(x^2 + x + 1), z^2 + z + x^2(x+1)y)$$

is the unique function field F of genus 7 over $\mathbb{F}_2$ admitting an unramified quadratic extension F'/F with $h_F = h_{F'}$.

Contents

1 Statement of the problem

2 Motivation: Inverting the zeta function function

3 Motivation: Cohomology of moduli spaces of curves

4 Enumerating curves by Brill–Noether type, I

5 Brill–Noether stratifications

6 Computing points on group quotients

The Lefschetz trace formula for stacks

For X a DM stack of finite type over $\mathbb{F}_q$, the Lefschetz trace formula for Frobenius holds: for any prime ℓ not dividing q,

$$\#X(\mathbb{F}_q) = \sum_{i=0}^{2\dim(X)} (-1)^i \operatorname{Trace}(\operatorname{Frob}_q, H_c^i(X_{\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)).$$

When there exists a polynomial $P(q)$ such that $\#X(\mathbb{F}_q) = P(q)$ for all prime powers q, we say X has **polynomial point count**.

$\mathcal{M}_g$ is known to have polynomial point count for $g \leq 7$ (Canning–Larson). By computing cohomology of $\overline{\mathcal{M}}_6$ using Pixton's relations, Bergstrom–Canning–Petersen–Schmitt obtain

$$\#\mathcal{M}_6(\mathbb{F}_q) = q^{15} + q^{14} + 2q^{13} + q^{12} - q^{10} + q^3 - 1$$

so in particular $\#\mathcal{M}_6(\mathbb{F}_2) = 68615$ (consistent with our data).

Marked points

Let $\mathcal{M}_{g,n}$ be the moduli space of smooth curves of genus g with n marked points. Given the set $\mathcal{M}_g(\mathbb{F}_q)$ for some g and q, it is easy to compute $\#\mathcal{M}_{g,n}(\mathbb{F}_q)$ for any n, or more generally $\#(\mathcal{M}_{g,n}/G)(\mathbb{F}_q)$ for any subgroup G of S_n.

For example, Canning–Larson showed that $\mathcal{M}_{6,n}$ also has polynomial point count for $n = 1, 2$ but the polynomials are not currently known. However, we know their values at $q = 2$, which reduces the number of cohomology groups that would need to be computed. (The value at $q = 3$ would reduce this again!)

Contents

1 Statement of the problem

2 Motivation: Inverting the zeta function function

3 Motivation: Cohomology of moduli spaces of curves

4 Enumerating curves by Brill–Noether type, I

5 Brill–Noether stratifications

6 Computing points on group quotients

Hyperelliptic curves

In odd characteristic, hyperelliptic curves over $\mathbb{F}_q$ are characterized up to quadratic twist by their Weierstrass locus. Enumerating curves thus amounts to computing

$$((\mathrm{Sym}^{2g+2}\,\mathbf{P}^1)/\,\mathrm{PGL}_2)(\mathbb{F}_q);$$

a good algorithm for this has been described by Howe.

In even characteristic, hyperelliptic curves over $\mathbb{F}_q$ are Artin–Schreier covers of $\mathbf{P}^1$. A good algorithm for enumerating these has been described by Xarles for $q = 2$; the general case is similar.

Bielliptic curves

Bielliptic curves can be handled similarly to hyperelliptic curves, by first enumerating elliptic curves and then double covers of each elliptic curve.

In practice, it is convenient to handle these by appealing to explicit geometric class field theory (e.g., in Magma).

Trigonal curves

Every trigonal curve of genus g occurs as an ample divisor in a Hirzebruch surface (the intersection of the quadrics vanishing on the canonical embedding). This surface has the form

$$\mathbf{F}_n := \mathbf{Proj}_{\mathbf{P}^1_k}(\mathcal{O}_{\mathbf{P}^1_k} \oplus \mathcal{O}(n)_{\mathbf{P}^1_k})$$

for some n (the **Maroni invariant**) with $n \leq \lfloor \frac{g+2}{3} \rfloor$, $n \equiv g \pmod 2$.

For $n > 0$, we can represent $\mathbf{F}_n$ as an $(n, 1)$-hypersurface in $\mathbf{P}^1_k \times_k \mathbf{P}^2_k$. For $n = 0$, we instead write $\mathbf{F}_0 \cong \mathbf{P}^1_k \times_k \mathbf{P}^1_k$.

Contents

1 Statement of the problem

2 Motivation: Inverting the zeta function function

3 Motivation: Cohomology of moduli spaces of curves

4 Enumerating curves by Brill–Noether type, I

5 Brill–Noether stratifications

6 Computing points on group quotients

The stratification of $\mathcal{M}_6$

We use a compact but hopefully self-explanatory notation for complete intersections.

Theorem (Mukai+ϵ)

Let C be a curve of genus 6 over a finite field k. Then C is exactly one of:

1. *Hyperelliptic.*
2. *Trigonal of Maroni invariant 2: a $(2,1) \cap (1,3)$ in $\mathbf{P}^1_k \times \mathbf{P}^2_k$.*
3. *Trigonal of Maroni invariant 0: a $(3,4)$ in $\mathbf{P}^1_k \times_k \mathbf{P}^1_k$.*
4. *Bielliptic.*
5. *Smooth quintic curve: a (5) in $\mathbf{P}^2_k$.*
6. *None of the above: a $(1)^4 \cap (2) \cap \mathrm{Gr}(2,5)$ in $\mathbf{P}^9_k$, where $\mathrm{Gr}(2,5)$ is the Grassmannian in its Plücker embedding.*

Warning: We are using that $\mathbf{P}^n_k$ has no twists (e.g., because $\mathrm{Br}(k) = 0$).

The stratification of $\mathcal{M}_7$

Theorem (Mukai+ϵ)

Let C be a curve of genus 7 over a finite field k. Then C is exactly one of:

1. *Hyperelliptic.*
2. *Trigonal of Maroni invariant 3: a (9) in $\mathbf{P}(1:1:3)_k$.*
3. *Trigonal of Maroni invariant 1: a $(1,1) \cap (3,3)$ in $\mathbf{P}^1_k \times_k \mathbf{P}^2_k$.*
4. *Bielliptic.*
5. *Not bielliptic but has a self-adjoint g^2_6: a $(3) \cap (4)$ in $\mathbf{P}(1:1:1:2)_k$.*
6. *Admits a pair of distinct g^2_6's: a $(1,1) \cap (1,1) \cap (2,2)$ in $\mathbf{P}^2_k \times_k \mathbf{P}^2_k$...*
7. **or** *its quadratic twist.*
8. *Admits a g^1_4 but no g^2_6: a $(1,1) \cap (1,2) \cap (1,2)$ in $\mathbf{P}^1_k \times_k \mathbf{P}^3_k$.*
9. *None of the above: a $(1)^9 \cap \mathrm{OG}^+(5,10)$ in $\mathbf{P}^{15}_k$ where OG^+ is a component of the orthogonal Grassmannian in its spinor embedding.*

Can this be pushed further?

Mukai has a similar "flowchart" for curves of genus 8 over an algebraically closed field. It should be easy to adapt this over a finite field.

In genus 9, Mukai describes the general curve but does not give a complete flowchart even over an algebraically closed field. The Betti tables and special linear series were computed by Sagraloff; it should be possible to extract a flowchart from this.

In genus 10, it is known that $\mathcal{M}_{10}$ is unirational, e.g., because a general genus-10 curve has 42 g^1_6's. But this description seems hard to work with; is there a better one out there? (The natural analogue of Mukai's arguments only yields a divisor in $\mathcal{M}_{10}$.)

Contents

1 Statement of the problem

2 Motivation: Inverting the zeta function function

3 Motivation: Cohomology of moduli spaces of curves

4 Enumerating curves by Brill–Noether type, I

5 Brill–Noether stratifications

6 Computing points on group quotients

A standard approach

Many strata in moduli have the form X/G where X is a quasiprojective variety and G is a (connected) linear algebraic group. To enumerate $(X/G)(k) = X(k)/G(k)$, the standard approach is:

- Enumerate $X(k)$.
- Choose a small generating set of $G(k)$ (e.g., by sampling at random).
- Find connected components in the Cayley graph on $X(k)$ for this generating set.

However, in some cases $X(k)$ is too big to enumerate, particularly for the generic strata in genus 6 and 7. We use various tricks to circumvent this, including the following.

Computing orbits of actions on subsets

Suppose G is a finite group acting on a finite set S which is small enough to enumerate. We can then find orbit representatives, together with a function f carrying each $s \in S$ to an element $g \in G$ for which $g^{-1}(s)$ is an orbit representative (sometimes called a **transporter**).

But now suppose we want the same for the induced action on $\binom{S}{n}$, the set of n-element subsets, for some n for which this set is too big to enumerate. Given a transporter for $\binom{S}{n-1}$, we can compute a transporter for the incidence correspondence

$$\Gamma \subset \binom{S}{n-1} \times \binom{S}{n}$$

and then one for $\binom{S}{n}$.

Computing orbits of actions on subspaces

Suppose G is a finite group acting k-linearly on a finite-dimensional k-vector space V which is small enough to enumerate, but we want to compute a transporter for the induced action on $\mathrm{Gr}(n, V)(k)$, the set of n-dimensional subspaces, for some n for which this set is too big to enumerate. Given a transporter for $\mathrm{Gr}(n-1, V)(k)$, we can compute a transporter for the incidence correspondence

$$\Gamma \subset \mathrm{Gr}(n-1, V)(k) \times \mathrm{Gr}(n, V)(k)$$

and then obtain one for $\mathrm{Gr}(n, V)(k)$.