Computation of zeta and L-functions: feasibility and applications

Kiran S. Kedlaya

Department of Mathematics, University of California, San Diego
School of Mathematics, Institute for Advanced Study (2018–2019)
kedlaya@ucsd.edu
http://kskedlaya.org/slides/

Arithmetic Geometry, Number Theory and Computation
Simons Collaboration Annual Meeting
Simons Foundation, New York City
January 11, 2019

Contents

Zeta functions of algebraic varieties

For X an algebraic variety over a finite field $\mathbb{F}_q$, the *zeta function* of X is

$$Z(X, T) = \prod_{x \in X^\circ} (1 - T^{\deg(x/\mathbb{F}_q)})^{-1} = \exp\left(\sum_{n=1}^\infty \frac{T^n}{n} \#X(\mathbb{F}_{q^n})\right),$$

where X° denotes the set of closed points of X. This is a rational function of T each of whose zeroes and poles in $\mathbb{C}$ has absolute value $q^{-i/2}$ for some $i \in \{0, \ldots, 2\dim(X)\}$.

For example, if X is an elliptic curve, then

$$Z(X, T) = \frac{1 - aT + qT^2}{(1 - T)(1 - qT)}$$

where $a \in \mathbb{Z} \cap [-2\sqrt{q}, 2\sqrt{q}]$ (Hasse's bound). The poles are at $T = q^{-0/2}, q^{-2/2}$ and the zeroes have absolute value $q^{-1/2}$.

Zeta functions of algebraic varieties

For X an algebraic variety over a finite field $\mathbb{F}_q$, the *zeta function* of X is

$$Z(X, T) = \prod_{x \in X^\circ} (1 - T^{\deg(x/\mathbb{F}_q)})^{-1} = \exp\left(\sum_{n=1}^{\infty} \frac{T^n}{n} \# X(\mathbb{F}_{q^n})\right),$$

where X° denotes the set of closed points of X. This is a rational function of T each of whose zeroes and poles in $\mathbb{C}$ has absolute value $q^{-i/2}$ for some $i \in \{0, \ldots, 2 \dim(X)\}$.

For example, if X is an elliptic curve, then

$$Z(X, T) = \frac{1 - aT + qT^2}{(1 - T)(1 - qT)}$$

where $a \in \mathbb{Z} \cap [-2\sqrt{q}, 2\sqrt{q}]$ (Hasse's bound). The poles are at $T = q^{-0/2}, q^{-2/2}$ and the zeroes have absolute value $q^{-1/2}$.

L-functions of algebraic varieties

For X a smooth proper variety over a number field K, its i-th L-function is a certain Euler product $L_i(X, s) = \prod_{\mathfrak{p}} L_{\mathfrak{p}}(X, \operatorname{Norm}(\mathfrak{p})^{-s})^{-1}$ ranging over prime ideals of $\mathfrak{o}_K$. For $\mathfrak{p}$ at which X has good reduction, $L_{\mathfrak{p}}(X, T)$ is the polynomial with constant term 1 whose roots are the zeroes/poles of $Z(X_{\mathfrak{p}}, T)$ of absolute value $\operatorname{Norm}(\mathfrak{p})^{-i/2}$.

It is expected that $L_i(X, s)$ admits a meromorphic continuation to $\mathbb{C}$ with functional equation relating s to $i + 1 - s$, but when this is known it is often a deep theorem (e.g., modularity of elliptic curves).

It is further expected that $L_i(X, s)$ has all of its zeroes on the line $\operatorname{Real}(s) = \frac{i+1}{2}$, but this is not known for any X. This question includes the Riemann hypothesis because

$$L_0(\operatorname{Spec}(\mathbb{Q}), s) = \prod_p (1 - p^{-s})^{-1} = \sum_{n=1}^{\infty} n^{-s} = \zeta(s).$$

L-functions of algebraic varieties

For X a smooth proper variety over a number field K, its i-th *L-function* is a certain Euler product $L_i(X, s) = \prod_{\mathfrak{p}} L_{\mathfrak{p}}(X, \text{Norm}(\mathfrak{p})^{-s})^{-1}$ ranging over prime ideals of $\mathfrak{o}_K$. For $\mathfrak{p}$ at which X has good reduction, $L_{\mathfrak{p}}(X, T)$ is the polynomial with constant term 1 whose roots are the zeroes/poles of $Z(X_{\mathfrak{p}}, T)$ of absolute value $\text{Norm}(\mathfrak{p})^{-i/2}$.

It is expected that $L_i(X, s)$ admits a meromorphic continuation to $\mathbb{C}$ with functional equation relating s to $i + 1 - s$, but when this is known it is often a deep theorem (e.g., modularity of elliptic curves).

It is further expected that $L_i(X, s)$ has all of its zeroes on the line $\text{Real}(s) = \frac{i+1}{2}$, but this is not known for any X. This question includes the Riemann hypothesis because

$$L_0(\text{Spec}(\mathbb{Q}), s) = \prod_p (1 - p^{-s})^{-1} = \sum_{n=1}^{\infty} n^{-s} = \zeta(s).$$

L-functions of algebraic varieties

For X a smooth proper variety over a number field K, its i-th *L-function* is a certain Euler product $L_i(X, s) = \prod_{\mathfrak{p}} L_{\mathfrak{p}}(X, \operatorname{Norm}(\mathfrak{p})^{-s})^{-1}$ ranging over prime ideals of $\mathfrak{o}_K$. For $\mathfrak{p}$ at which X has good reduction, $L_{\mathfrak{p}}(X, T)$ is the polynomial with constant term 1 whose roots are the zeroes/poles of $Z(X_{\mathfrak{p}}, T)$ of absolute value $\operatorname{Norm}(\mathfrak{p})^{-i/2}$.

It is expected that $L_i(X, s)$ admits a meromorphic continuation to $\mathbb{C}$ with functional equation relating s to $i + 1 - s$, but when this is known it is often a deep theorem (e.g., modularity of elliptic curves).

It is further expected that $L_i(X, s)$ has all of its zeroes on the line $\operatorname{Real}(s) = \frac{i+1}{2}$, but this is not known for any X. This question includes the Riemann hypothesis because

$$L_0(\operatorname{Spec}(\mathbb{Q}), s) = \prod_p (1 - p^{-s})^{-1} = \sum_{n=1}^{\infty} n^{-s} = \zeta(s).$$

Overview of the talk

This talk has two main goals.

- To survey some examples of questions arising from arithmetic and/or algebraic geometry in which machine computation of the zeta function or L-function associated to a algebraic variety plays (or has the potential to play) an important role.

- To give an indication of what types of computations along these lines are feasible. We will not say much about the techniques involved, except to mention the important role played by p-adic analytic methods (in the style of Dwork).

As a baseline, keep in mind that one can in principle compute a zeta function $Z(X, T)$ by counting $X(\mathbb{F}_{q^n})$ for enough n (once one has a bound on its degree as a rational function). We refer to this as the *brute force method* hereafter.

Overview of the talk

This talk has two main goals.

- To survey some examples of questions arising from arithmetic and/or algebraic geometry in which machine computation of the zeta function or L-function associated to a algebraic variety plays (or has the potential to play) an important role.

- To give an indication of what types of computations along these lines are feasible. We will not say much about the techniques involved, except to mention the important role played by p-adic analytic methods (in the style of Dwork).

As a baseline, keep in mind that one can in principle compute a zeta function $Z(X, T)$ by counting $X(\mathbb{F}_{q^n})$ for enough n (once one has a bound on its degree as a rational function). We refer to this as the *brute force method* hereafter.

Overview of the talk

This talk has two main goals.

- To survey some examples of questions arising from arithmetic and/or algebraic geometry in which machine computation of the zeta function or L-function associated to a algebraic variety plays (or has the potential to play) an important role.

- To give an indication of what types of computations along these lines are feasible. We will not say much about the techniques involved, except to mention the important role played by p-adic analytic methods (in the style of Dwork).

As a baseline, keep in mind that one can in principle compute a zeta function $Z(X, T)$ by counting $X(\mathbb{F}_{q^n})$ for enough n (once one has a bound on its degree as a rational function). We refer to this as the *brute force method* hereafter.

Overview of the talk

This talk has two main goals.

- To survey some examples of questions arising from arithmetic and/or algebraic geometry in which machine computation of the zeta function or L-function associated to a algebraic variety plays (or has the potential to play) an important role.

- To give an indication of what types of computations along these lines are feasible. We will not say much about the techniques involved, except to mention the important role played by p-adic analytic methods (in the style of Dwork).

As a baseline, keep in mind that one can in principle compute a zeta function $Z(X, T)$ by counting $X(\mathbb{F}_{q^n})$ for enough n (once one has a bound on its degree as a rational function). We refer to this as the *brute force method* hereafter.

Contents

Equalities of zeta/L-functions

Let X, Y be smooth projective varieties over $\mathbb{F}_q$ (resp. K). Under what conditions do the zeta functions (resp. L-functions) of X and Y coincide?

Obviously this occurs if X and Y are isomorphic, but the converse is false. For example, if X and Y are abelian varieties, a theorem of Tate (resp. Faltings) implies that X and Y have the same zeta function (resp. the same L-functions) if and only if they are isogenous.

In higher dimensions, Orlov conjectures that over a finite field, if X, Y are *derived equivalent*[*] (i.e., the derived categories of coherent sheaves on X and Y are equivalent), then $Z(X, T) = Z(Y, T)$; this is known up to dimension 3 (Honigs et al.). This should work similarly over a number field.

Question: are there interesting instances of Orlov's conjecture in dimension 4? If so, it should be feasible to test them numerically for small p.

---

[*]Even for abelian varieties, this is stronger than being isogenous; however, an abelian variety is always derived equivalent to its dual.

Equalities of zeta/L-functions

Let X, Y be smooth projective varieties over $\mathbb{F}_q$ (resp. K). Under what conditions do the zeta functions (resp. L-functions) of X and Y coincide?

Obviously this occurs if X and Y are isomorphic, but the converse is false. For example, if X and Y are abelian varieties, a theorem of Tate (resp. Faltings) implies that X and Y have the same zeta function (resp. the same L-functions) if and only if they are isogenous.

In higher dimensions, Orlov conjectures that over a finite field, if X, Y are *derived equivalent*[*] (i.e., the derived categories of coherent sheaves on X and Y are equivalent), then $Z(X, T) = Z(Y, T)$; this is known up to dimension 3 (Honigs et al.). This should work similarly over a number field.

Question: are there interesting instances of Orlov's conjecture in dimension 4? If so, it should be feasible to test them numerically for small p.

---

[*]Even for abelian varieties, this is stronger than being isogenous; however, an abelian variety is always derived equivalent to its dual.

Equalities of zeta/L-functions

Let X, Y be smooth projective varieties over $\mathbb{F}_q$ (resp. K). Under what conditions do the zeta functions (resp. L-functions) of X and Y coincide?

Obviously this occurs if X and Y are isomorphic, but the converse is false. For example, if X and Y are abelian varieties, a theorem of Tate (resp. Faltings) implies that X and Y have the same zeta function (resp. the same L-functions) if and only if they are isogenous.

In higher dimensions, Orlov conjectures that over a finite field, if X, Y are *derived equivalent*[*] (i.e., the derived categories of coherent sheaves on X and Y are equivalent), then $Z(X, T) = Z(Y, T)$; this is known up to dimension 3 (Honigs et al.). This should work similarly over a number field.

Question: are there interesting instances of Orlov's conjecture in dimension 4? If so, it should be feasible to test them numerically for small p.

---

[*]Even for abelian varieties, this is stronger than being isogenous; however, an abelian variety is always derived equivalent to its dual.

Equalities of zeta/L-functions

Let X, Y be smooth projective varieties over $\mathbb{F}_q$ (resp. K). Under what conditions do the zeta functions (resp. L-functions) of X and Y coincide?

Obviously this occurs if X and Y are isomorphic, but the converse is false. For example, if X and Y are abelian varieties, a theorem of Tate (resp. Faltings) implies that X and Y have the same zeta function (resp. the same L-functions) if and only if they are isogenous.

In higher dimensions, Orlov conjectures that over a finite field, if X, Y are *derived equivalent*[*] (i.e., the derived categories of coherent sheaves on X and Y are equivalent), then $Z(X, T) = Z(Y, T)$; this is known up to dimension 3 (Honigs et al.). This should work similarly over a number field.

Question: are there interesting instances of Orlov's conjecture in dimension 4? If so, it should be feasible to test them numerically for small p.

---

[*]Even for abelian varieties, this is stronger than being isogenous; however, an abelian variety is always derived equivalent to its dual.

Modularity of L-functions: elliptic curves

For X an elliptic curve over $\mathbb{Q}$, the analytic behavior of $L_1(X, s)$ is established by equating it with the L-function associated to a cuspidal weight-2 GL_2 eigenform (Wiles, Breuil–Conrad–Diamond–Taylor).

This comparison was confirmed numerically in many cases beforehand, especially by Cremona. His tabulation of elliptic curves over $\mathbb{Q}$ by conductor (found in LMFDB) depends on this theorem for its completeness.

The same strategy extends (partially) to elliptic curves over totally real fields and CM fields (many authors). This also implies the Sato-Tate conjecture for these elliptic curves, more on which later.

Modularity of *L*-functions: elliptic curves

For X an elliptic curve over $\mathbb{Q}$, the analytic behavior of $L_1(X, s)$ is established by equating it with the *L*-function associated to a cuspidal weight-2 GL_2 eigenform (Wiles, Breuil–Conrad–Diamond–Taylor).

This comparison was confirmed numerically in many cases beforehand, especially by Cremona. His tabulation of elliptic curves over $\mathbb{Q}$ by conductor (found in LMFDB) depends on this theorem for its completeness.

The same strategy extends (partially) to elliptic curves over totally real fields and CM fields (many authors). This also implies the Sato-Tate conjecture for these elliptic curves, more on which later.

Modularity of L-functions: elliptic curves

For X an elliptic curve over $\mathbb{Q}$, the analytic behavior of $L_1(X, s)$ is established by equating it with the L-function associated to a cuspidal weight-2 GL_2 eigenform (Wiles, Breuil–Conrad–Diamond–Taylor).

This comparison was confirmed numerically in many cases beforehand, especially by Cremona. His tabulation of elliptic curves over $\mathbb{Q}$ by conductor (found in LMFDB) depends on this theorem for its completeness.

The same strategy extends (partially) to elliptic curves over totally real fields and CM fields (many authors). This also implies the Sato-Tate conjecture for these elliptic curves, more on which later.

Modularity of *L*-functions: beyond elliptic curves

The Langlands program predicts that *L*-functions associated to algebraic varieties should also arise from automorphic forms (which again would explain their analytic behavior).

- There is a precise conjecture for abelian surfaces, checked in some cases (Brumer–Pacetti–Poor–Tornaría–Voight–Yuen, Brumer–Kramer, Boxer–Calegari–Gee–Pilloni).

- There are scattered examples of K3 surfaces and Calabi-Yau threefolds for which the middle-degree *L*-functions have been matched with GL_2-modular forms (many authors).

Some of these statements depend on the *Faltings–Serre criterion* for certifying an equality of two *L*-functions based on comparison of Euler factors; this depends on the existence of Galois representations on both sides. Here computing the geometric *L*-function is the easy part...

Modularity of *L*-functions: beyond elliptic curves

The Langlands program predicts that *L*-functions associated to algebraic varieties should also arise from automorphic forms (which again would explain their analytic behavior).

- There is a precise conjecture for abelian surfaces, checked in some cases (Brumer–Pacetti–Poor–Tornaría–Voight–Yuen, Brumer–Kramer, Boxer–Calegari–Gee–Pilloni).

- There are scattered examples of K3 surfaces and Calabi-Yau threefolds for which the middle-degree *L*-functions have been matched with GL_2-modular forms (many authors).

Some of these statements depend on the *Faltings–Serre criterion* for certifying an equality of two *L*-functions based on comparison of Euler factors; this depends on the existence of Galois representations on both sides. Here computing the geometric *L*-function is the easy part...

Modularity of *L*-functions: beyond elliptic curves

The Langlands program predicts that *L*-functions associated to algebraic varieties should also arise from automorphic forms (which again would explain their analytic behavior).

- There is a precise conjecture for abelian surfaces, checked in some cases (Brumer–Pacetti–Poor–Tornaría–Voight–Yuen, Brumer–Kramer, Boxer–Calegari–Gee–Pilloni).
- There are scattered examples of K3 surfaces and Calabi-Yau threefolds for which the middle-degree *L*-functions have been matched with GL_2-modular forms (many authors).

Some of these statements depend on the *Faltings–Serre criterion* for certifying an equality of two *L*-functions based on comparison of Euler factors; this depends on the existence of Galois representations on both sides. Here computing the geometric *L*-function is the easy part...

Modularity of *L*-functions: beyond elliptic curves

The Langlands program predicts that *L*-functions associated to algebraic varieties should also arise from automorphic forms (which again would explain their analytic behavior).

- There is a precise conjecture for abelian surfaces, checked in some cases (Brumer–Pacetti–Poor–Tornaría–Voight–Yuen, Brumer–Kramer, Boxer–Calegari–Gee–Pilloni).
- There are scattered examples of K3 surfaces and Calabi-Yau threefolds for which the middle-degree *L*-functions have been matched with GL_2-modular forms (many authors).

Some of these statements depend on the *Faltings–Serre criterion* for certifying an equality of two *L*-functions based on comparison of Euler factors; this depends on the existence of Galois representations on both sides. Here computing the geometric *L*-function is the easy part...

Contents

Abelian varieties and the Honda-Tate theorem

For X an abelian variety of dimension g over $\mathbb{F}_q$, the possibilities for $Z(X, T)$ are dictated by the Honda-Tate theorem, and each one corresponds to a unique isogeny class. These have been tabulated in LMFDB for small g, q (Dupuy-K–Roe–Vincent).

This provides a natural intermediate step to tabulating abelian varieties over $\mathbb{F}_q$ up to isomorphism. There will be a Collaboration workshop at ICERM at the end of January to work on this.

Abelian varieties and the Honda-Tate theorem

For X an abelian variety of dimension g over $\mathbb{F}_q$, the possibilities for $Z(X, T)$ are dictated by the Honda-Tate theorem, and each one corresponds to a unique isogeny class. These have been tabulated in LMFDB for small g, q (Dupuy-K–Roe–Vincent).

This provides a natural intermediate step to tabulating abelian varieties over $\mathbb{F}_q$ up to isomorphism. There will be a Collaboration workshop at ICERM at the end of January to work on this.

Curves

By contrast, for X a curve of genus g over $\mathbb{F}_q$, the possibilities for $Z(X, T)$ are much harder to pin down (except for small g). For example, there are extra restrictions coming from positivity conditions:

$$\#X(\mathbb{F}_q) \geq 0$$
$$\#X(\mathbb{F}_{q^{mn}}) \geq \#X(\mathbb{F}_{q^n}).$$

For g large compare to q, these conditions cause the Weil upper bound $\#X(\mathbb{F}_q)$ to be suboptimal (Ihara, Drinfeld-Vlăduţ).

Much attention has gone into studying the maximum number of $\mathbb{F}_q$-points on a curve of genus g, as this a key question in algebraic coding theory; see https://manypoints.org.

Curves

By contrast, for X a curve of genus g over $\mathbb{F}_q$, the possibilities for $Z(X, T)$ are much harder to pin down (except for small g). For example, there are extra restrictions coming from positivity conditions:

$$\#X(\mathbb{F}_q) \geq 0$$
$$\#X(\mathbb{F}_{q^{mn}}) \geq \#X(\mathbb{F}_{q^n}).$$

For g large compare to q, these conditions cause the Weil upper bound $\#X(\mathbb{F}_q)$ to be suboptimal (Ihara, Drinfeld-Vlăduţ).

Much attention has gone into studying the maximum number of $\mathbb{F}_q$-points on a curve of genus g, as this a key question in algebraic coding theory; see https://manypoints.org.

Curves

By contrast, for X a curve of genus g over $\mathbb{F}_q$, the possibilities for $Z(X, T)$ are much harder to pin down (except for small g). For example, there are extra restrictions coming from positivity conditions:

$$\#X(\mathbb{F}_q) \geq 0$$
$$\#X(\mathbb{F}_{q^{mn}}) \geq \#X(\mathbb{F}_{q^n}).$$

For g large compare to q, these conditions cause the Weil upper bound $\#X(\mathbb{F}_q)$ to be suboptimal (Ihara, Drinfeld-Vlăduț).

Much attention has gone into studying the maximum number of $\mathbb{F}_q$-points on a curve of genus g, as this a key question in algebraic coding theory; see https://manypoints.org.

Curves

By contrast, for X a curve of genus g over $\mathbb{F}_q$, the possibilities for $Z(X, T)$ are much harder to pin down (except for small g). For example, there are extra restrictions coming from positivity conditions:

$$\#X(\mathbb{F}_q) \geq 0$$
$$\#X(\mathbb{F}_{q^{mn}}) \geq \#X(\mathbb{F}_{q^n}).$$

For g large compare to q, these conditions cause the Weil upper bound $\#X(\mathbb{F}_q)$ to be suboptimal (Ihara, Drinfeld-Vlăduț).

Much attention has gone into studying the maximum number of $\mathbb{F}_q$-points on a curve of genus g, as this a key question in algebraic coding theory; see https://manypoints.org.

Curves

By contrast, for X a curve of genus g over $\mathbb{F}_q$, the possibilities for $Z(X, T)$ are much harder to pin down (except for small g). For example, there are extra restrictions coming from positivity conditions:

$$\#X(\mathbb{F}_q) \geq 0$$
$$\#X(\mathbb{F}_{q^{mn}}) \geq \#X(\mathbb{F}_{q^n}).$$

For g large compare to q, these conditions cause the Weil upper bound $\#X(\mathbb{F}_q)$ to be suboptimal (Ihara, Drinfeld-Vlăduţ).

Much attention has gone into studying the maximum number of $\mathbb{F}_q$-points on a curve of genus g, as this a key question in algebraic coding theory; see https://manypoints.org.

K3 surfaces

For X a K3 surface over $\mathbb{F}_q$, the zeta function has the form

$$\frac{1}{(1-T)(1-qT)(1-q^2T)Q_{\mathrm{alg}}(T)Q_{\mathrm{trans}}(T)}$$

where $Q_{\mathrm{alg}}(T)$ has roots of the form q^{-1} times a root of unity, while $Q_{\mathrm{trans}}(T)$ has other roots of absolute value q^{-1}.

There is no conjecture for exactly what zeta functions should occur over a fixed $\mathbb{F}_q$. However, a theorem of Taelman–Ito shows that all feasible candidates for $Q_{\mathrm{trans}}(T)$ occur, except that one might have to pass from $\mathbb{F}_q$ to some extension to make things match up. (As in Honda-Tate, this is proved by making analytic constructions in characteristic 0 and then descending these to number fields.)

K3 surfaces

For X a K3 surface over $\mathbb{F}_q$, the zeta function has the form

$$\frac{1}{(1-T)(1-qT)(1-q^2T)Q_{\mathrm{alg}}(T)Q_{\mathrm{trans}}(T)}$$

where $Q_{\mathrm{alg}}(T)$ has roots of the form q^{-1} times a root of unity, while $Q_{\mathrm{trans}}(T)$ has other roots of absolute value q^{-1}.

There is no conjecture for exactly what zeta functions should occur over a fixed $\mathbb{F}_q$. However, a theorem of Taelman–Ito shows that all feasible candidates for $Q_{\mathrm{trans}}(T)$ occur, except that one might have to pass from $\mathbb{F}_q$ to some extension to make things match up. (As in Honda-Tate, this is proved by making analytic constructions in characteristic 0 and then descending these to number fields.)

Canvassing K3 surfaces

In order to probe the Taelman–Ito theorem, with Sutherland we found all smooth quartic K3 surfaces over $\mathbb{F}_2$ and their zeta functions. There are about 5×10^5 surfaces up to PGL_4-equivalence, and about 5×10^4 distinct zeta functions.

This is much smaller than the number of candidate zeta functions (about 1.5×10^6), but we identified about 2000 candidates which can *only* arise from smooth quartics (and not any other type of K3 surfaces), and these do all occur without any base extension required.

It may be feasible to do a similar census in some other cases:

- for other types of K3 surfaces over $\mathbb{F}_2$;
- for quartic K3 surfaces over $\mathbb{F}_3$ (Costa–Harvey–K, in progress);
- for cubic fourfolds over $\mathbb{F}_2$ (Auel–Costa–K–Sutherland, in progress).

Canvassing K3 surfaces

In order to probe the Taelman–Ito theorem, with Sutherland we found all smooth quartic K3 surfaces over $\mathbb{F}_2$ and their zeta functions. There are about 5×10^5 surfaces up to PGL_4-equivalence, and about 5×10^4 distinct zeta functions.

This is much smaller than the number of candidate zeta functions (about 1.5×10^6), but we identified about 2000 candidates which can *only* arise from smooth quartics (and not any other type of K3 surfaces), and these do all occur without any base extension required.

It may be feasible to do a similar census in some other cases:

- for other types of K3 surfaces over $\mathbb{F}_2$;
- for quartic K3 surfaces over $\mathbb{F}_3$ (Costa–Harvey–K, in progress);
- for cubic fourfolds over $\mathbb{F}_2$ (Auel–Costa–K–Sutherland, in progress).

Canvassing K3 surfaces

In order to probe the Taelman–Ito theorem, with Sutherland we found all smooth quartic K3 surfaces over $\mathbb{F}_2$ and their zeta functions. There are about 5×10^5 surfaces up to PGL_4-equivalence, and about 5×10^4 distinct zeta functions.

This is much smaller than the number of candidate zeta functions (about 1.5×10^6), but we identified about 2000 candidates which can *only* arise from smooth quartics (and not any other type of K3 surfaces), and these do all occur without any base extension required.

It may be feasible to do a similar census in some other cases:

- for other types of K3 surfaces over $\mathbb{F}_2$;
- for quartic K3 surfaces over $\mathbb{F}_3$ (Costa–Harvey–K, in progress);
- for cubic fourfolds over $\mathbb{F}_2$ (Auel–Costa–K–Sutherland, in progress).

Contents

The Hodge and Tate conjectures

For X smooth proper over $\mathbb{F}_q$, the pole order of $Z(X, T)$ at q^{-1} is at least the Néron-Severi rank (Picard number) of X. Equality is predicted by the *Tate conjecture*.

This is equivalent to predicting that the generalized eigenspace $F = q$ on étale H^2 is spanned by cycle classes (and in particular is semisimple). Compare with the Lefschetz (1,1) theorem: for X smooth proper over $\mathbb{C}$, $H^2(X, \mathbb{Z}) \cap H^{1,1}$ is spanned by cycle classes.

By analogy with the Hodge conjecture, one has a similar inequality and conjecture about the pole order of $Z(X, T)$ at q^{-i} for all positive integers i.

The Hodge and Tate conjectures

For X smooth proper over $\mathbb{F}_q$, the pole order of $Z(X, T)$ at q^{-1} is at least the Néron-Severi rank (Picard number) of X. Equality is predicted by the *Tate conjecture*.

This is equivalent to predicting that the generalized eigenspace $F = q$ on étale H^2 is spanned by cycle classes (and in particular is semisimple). Compare with the Lefschetz (1,1) theorem: for X smooth proper over $\mathbb{C}$, $H^2(X, \mathbb{Z}) \cap H^{1,1}$ is spanned by cycle classes.

By analogy with the Hodge conjecture, one has a similar inequality and conjecture about the pole order of $Z(X, T)$ at q^{-i} for all positive integers i.

The Hodge and Tate conjectures

For X smooth proper over $\mathbb{F}_q$, the pole order of $Z(X, T)$ at q^{-1} is at least the Néron-Severi rank (Picard number) of X. Equality is predicted by the *Tate conjecture*.

This is equivalent to predicting that the generalized eigenspace $F = q$ on étale H^2 is spanned by cycle classes (and in particular is semisimple). Compare with the Lefschetz (1,1) theorem: for X smooth proper over $\mathbb{C}$, $H^2(X, \mathbb{Z}) \cap H^{1,1}$ is spanned by cycle classes.

By analogy with the Hodge conjecture, one has a similar inequality and conjecture about the pole order of $Z(X, T)$ at q^{-i} for all positive integers i.

The (weak) conjecture of Birch and Swinnerton-Dyer

For X an elliptic curve over K, it is conjectured that the order of vanishing at $L_1(X, s)$ at $s = 1$ equals the rank of the finitely generated abelian group $X(K)$. For $K = \mathbb{Q}$, this is known when the order of vanishing is ≤ 1 (Gross–Zagier, Kolyvagin).

This has an awkward side effect for computations: it is not known how to exhibit an example where the order of vanishing at $s = 1$ is *provably* any larger than 3. Such an example would be useful for proving completeness of tables of imaginary quadratic fields of a given class number (Goldfeld–Watkins).

The analogous statement over a finite extension of $\mathbb{F}_q(t)$ is a case of the Tate conjecture; consequently, it is known that the order of vanishing is at least the algebraic rank.

The (weak) conjecture of Birch and Swinnerton-Dyer

For X an elliptic curve over K, it is conjectured that the order of vanishing at $L_1(X, s)$ at $s = 1$ equals the rank of the finitely generated abelian group $X(K)$. For $K = \mathbb{Q}$, this is known when the order of vanishing is ≤ 1 (Gross–Zagier, Kolyvagin).

This has an awkward side effect for computations: it is not known how to exhibit an example where the order of vanishing at $s = 1$ is *provably* any larger than 3. Such an example would be useful for proving completeness of tables of imaginary quadratic fields of a given class number (Goldfeld–Watkins).

The analogous statement over a finite extension of $\mathbb{F}_q(t)$ is a case of the Tate conjecture; consequently, it is known that the order of vanishing is at least the algebraic rank.

The (weak) conjecture of Birch and Swinnerton-Dyer

For X an elliptic curve over K, it is conjectured that the order of vanishing at $L_1(X, s)$ at $s = 1$ equals the rank of the finitely generated abelian group $X(K)$. For $K = \mathbb{Q}$, this is known when the order of vanishing is ≤ 1 (Gross–Zagier, Kolyvagin).

This has an awkward side effect for computations: it is not known how to exhibit an example where the order of vanishing at $s = 1$ is *provably* any larger than 3. Such an example would be useful for proving completeness of tables of imaginary quadratic fields of a given class number (Goldfeld–Watkins).

The analogous statement over a finite extension of $\mathbb{F}_q(t)$ is a case of the Tate conjecture; consequently, it is known that the order of vanishing is at least the algebraic rank.

Specialization and Picard numbers

The Picard number does not decrease upon specialization from K to $\mathbb{F}_q$. Consequently, one can use a zeta function computation to give an upper bound on the Picard number of a variety over either K or $\mathbb{F}_q$.

The same is true for the geometric Picard number, so in principle one can also use a zeta function computation to give an upper bound on the Picard number of a variety over either K or $\mathbb{F}_q$. There is a catch: the geometric Picard number over $\mathbb{F}_q$ always has the same parity as $\dim H^2$, whereas this is not true over K.

So for example, it is not clear how to use a zeta function computation to check that a K3 surface over K has geometric Picard number 1. This is possible by a method of van Luijk, more on which shortly.

Specialization and Picard numbers

The Picard number does not decrease upon specialization from K to $\mathbb{F}_q$. Consequently, one can use a zeta function computation to give an upper bound on the Picard number of a variety over either K or $\mathbb{F}_q$.

The same is true for the geometric Picard number, so in principle one can also use a zeta function computation to give an upper bound on the Picard number of a variety over either K or $\mathbb{F}_q$. There is a catch: the geometric Picard number over $\mathbb{F}_q$ always has the same parity as $\dim H^2$, whereas this is not true over K.

So for example, it is not clear how to use a zeta function computation to check that a K3 surface over K has geometric Picard number 1. This is possible by a method of van Luijk, more on which shortly.

Specialization and Picard numbers

The Picard number does not decrease upon specialization from K to $\mathbb{F}_q$. Consequently, one can use a zeta function computation to give an upper bound on the Picard number of a variety over either K or $\mathbb{F}_q$.

The same is true for the geometric Picard number, so in principle one can also use a zeta function computation to give an upper bound on the Picard number of a variety over either K or $\mathbb{F}_q$. There is a catch: the geometric Picard number over $\mathbb{F}_q$ always has the same parity as $\dim H^2$, whereas this is not true over K.

So for example, it is not clear how to use a zeta function computation to check that a K3 surface over K has geometric Picard number 1. This is possible by a method of van Luijk, more on which shortly.

Special cubic fourfolds

Consider the moduli space of smooth cubic fourfolds in $\mathbb{P}^5$. A cubic fourfold is *special* if it contains a surface not homologous to a complete intersection; the special cubic fourfolds form a countable union of divisors in the moduli space.

Ranestad–Voisin exhibited four divisors on the moduli space of smooth cubic fourfolds in $\mathbb{P}^5$ which did not appear to be of this type. Confirming this entails finding a nonspecial cubic fourfold of each of four prescribed shapes.

Ranestad–Voisin were able to treat one of the four shapes by hand. Addington–Auel handled two more by computing zeta functions over $\mathbb{F}_2$ by a modified[†] brute-force method. Costa–Harvey–K handled the fourth by computing with p-adic cohomology over $\mathbb{F}_p$ (with $p = 31$).

[†]A true brute-force calculation would require enumerating points on a fourfold over $\mathbb{F}_{2^{11}}$; instead, one considers a certain conic fibration in which one can read off the point counts on fibers without enumerating.

Special cubic fourfolds

Consider the moduli space of smooth cubic fourfolds in $\mathbb{P}^5$. A cubic fourfold is *special* if it contains a surface not homologous to a complete intersection; the special cubic fourfolds form a countable union of divisors in the moduli space.

Ranestad–Voisin exhibited four divisors on the moduli space of smooth cubic fourfolds in $\mathbb{P}^5$ which did not appear to be of this type. Confirming this entails finding a nonspecial cubic fourfold of each of four prescribed shapes.

Ranestad–Voisin were able to treat one of the four shapes by hand. Addington–Auel handled two more by computing zeta functions over $\mathbb{F}_2$ by a modified[†] brute-force method. Costa–Harvey–K handled the fourth by computing with p-adic cohomology over $\mathbb{F}_p$ (with $p = 31$).

[†]A true brute-force calculation would require enumerating points on a fourfold over $\mathbb{F}_{2^{11}}$; instead, one considers a certain conic fibration in which one can read off the point counts on fibers without enumerating.

Special cubic fourfolds

Consider the moduli space of smooth cubic fourfolds in $\mathbb{P}^5$. A cubic fourfold is *special* if it contains a surface not homologous to a complete intersection; the special cubic fourfolds form a countable union of divisors in the moduli space.

Ranestad–Voisin exhibited four divisors on the moduli space of smooth cubic fourfolds in $\mathbb{P}^5$ which did not appear to be of this type. Confirming this entails finding a nonspecial cubic fourfold of each of four prescribed shapes.

Ranestad–Voisin were able to treat one of the four shapes by hand. Addington–Auel handled two more by computing zeta functions over $\mathbb{F}_2$ by a modified[†] brute-force method. Costa–Harvey–K handled the fourth by computing with p-adic cohomology over $\mathbb{F}_p$ (with $p = 31$).

[†]A true brute-force calculation would require enumerating points on a fourfold over $\mathbb{F}_{2^{11}}$; instead, one considers a certain conic fibration in which one can read off the point counts on fibers without enumerating.

Contents

Special values in number theory

Starting from Dirichlet's class number formula, many conjectures in number theory have emerged relating special values of *L*-functions to arithmetic quantities. These include:

- Stark's conjectures on units in number fields;
- the strong form of the conjecture of Birch and Swinnerton-Dyer for elliptic curves;
- the Bloch–Kato Tamagawa number conjecture;
- Beilinson's conjecture on regulators in algebraic *K*-theory.

In some settings, there is room to make the conjectures more precise; experimental evidence, based on computations of *L*-functions, may play a role here (as it did for BSD).

Special values in number theory

Starting from Dirichlet's class number formula, many conjectures in number theory have emerged relating special values of *L*-functions to arithmetic quantities. These include:

- Stark's conjectures on units in number fields;
- the strong form of the conjecture of Birch and Swinnerton-Dyer for elliptic curves;
- the Bloch–Kato Tamagawa number conjecture;
- Beilinson's conjecture on regulators in algebraic *K*-theory.

In some settings, there is room to make the conjectures more precise; experimental evidence, based on computations of *L*-functions, may play a role here (as it did for BSD).

Special values in number theory

Starting from Dirichlet's class number formula, many conjectures in number theory have emerged relating special values of L-functions to arithmetic quantities. These include:

- Stark's conjectures on units in number fields;
- the strong form of the conjecture of Birch and Swinnerton-Dyer for elliptic curves;
- the Bloch–Kato Tamagawa number conjecture;
- Beilinson's conjecture on regulators in algebraic K-theory.

In some settings, there is room to make the conjectures more precise; experimental evidence, based on computations of L-functions, may play a role here (as it did for BSD).

Special values in number theory

Starting from Dirichlet's class number formula, many conjectures in number theory have emerged relating special values of *L*-functions to arithmetic quantities. These include:

- Stark's conjectures on units in number fields;
- the strong form of the conjecture of Birch and Swinnerton-Dyer for elliptic curves;
- the Bloch–Kato Tamagawa number conjecture;
- Beilinson's conjecture on regulators in algebraic *K*-theory.

In some settings, there is room to make the conjectures more precise; experimental evidence, based on computations of *L*-functions, may play a role here (as it did for BSD).

Special values in number theory

Starting from Dirichlet's class number formula, many conjectures in number theory have emerged relating special values of L-functions to arithmetic quantities. These include:

- Stark's conjectures on units in number fields;
- the strong form of the conjecture of Birch and Swinnerton-Dyer for elliptic curves;
- the Bloch–Kato Tamagawa number conjecture;
- Beilinson's conjecture on regulators in algebraic K-theory.

In some settings, there is room to make the conjectures more precise; experimental evidence, based on computations of L-functions, may play a role here (as it did for BSD).

Special values in number theory

Starting from Dirichlet's class number formula, many conjectures in number theory have emerged relating special values of L-functions to arithmetic quantities. These include:

- Stark's conjectures on units in number fields;
- the strong form of the conjecture of Birch and Swinnerton-Dyer for elliptic curves;
- the Bloch–Kato Tamagawa number conjecture;
- Beilinson's conjecture on regulators in algebraic K-theory.

In some settings, there is room to make the conjectures more precise; experimental evidence, based on computations of L-functions, may play a role here (as it did for BSD).

K3 surfaces

For X a K3 surface over $\mathbb{F}_q$, the Tate conjecture is known (many authors); hence $Z(X, T)$ has a pole at $T = q^{-1}$ of order equal to the Picard number. By the Artin-Tate formula, the leading coefficient is the discriminant of the Picard lattice times the order of the Brauer group; the latter is a perfect square.

Elsenhans–Jahnel observed that comparing the Artin-Tate formulas over $\mathbb{F}_q$ and $\mathbb{F}_{q^2}$ yields a previously unknown restriction on $Z(X, T)$.

van Luijk observed that if X is a K3 surface over K whose reduction to $\mathbb{F}_q$ has the same Picard number, then the discriminant does not change upon reduction. In particular, if X has two reductions with geometric Picard number 2 whose discriminants are not in the same square class, then X has geometric Picard number 1.

Question: can zeta/L-function computations be used to gain any effective control over the Brauer group of a K3 surface over K?

K3 surfaces

For X a K3 surface over $\mathbb{F}_q$, the Tate conjecture is known (many authors); hence $Z(X, T)$ has a pole at $T = q^{-1}$ of order equal to the Picard number. By the Artin-Tate formula, the leading coefficient is the discriminant of the Picard lattice times the order of the Brauer group; the latter is a perfect square.

Elsenhans–Jahnel observed that comparing the Artin-Tate formulas over $\mathbb{F}_q$ and $\mathbb{F}_{q^2}$ yields a previously unknown restriction on $Z(X, T)$.

van Luijk observed that if X is a K3 surface over K whose reduction to $\mathbb{F}_q$ has the same Picard number, then the discriminant does not change upon reduction. In particular, if X has two reductions with geometric Picard number 2 whose discriminants are not in the same square class, then X has geometric Picard number 1.

Question: can zeta/L-function computations be used to gain any effective control over the Brauer group of a K3 surface over K?

K3 surfaces

For X a K3 surface over $\mathbb{F}_q$, the Tate conjecture is known (many authors); hence $Z(X, T)$ has a pole at $T = q^{-1}$ of order equal to the Picard number. By the Artin-Tate formula, the leading coefficient is the discriminant of the Picard lattice times the order of the Brauer group; the latter is a perfect square.

Elsenhans–Jahnel observed that comparing the Artin-Tate formulas over $\mathbb{F}_q$ and $\mathbb{F}_{q^2}$ yields a previously unknown restriction on $Z(X, T)$.

van Luijk observed that if X is a K3 surface over K whose reduction to $\mathbb{F}_q$ has the same Picard number, then the discriminant does not change upon reduction. In particular, if X has two reductions with geometric Picard number 2 whose discriminants are not in the same square class, then X has geometric Picard number 1.

Question: can zeta/L-function computations be used to gain any effective control over the Brauer group of a K3 surface over K?

K3 surfaces

For X a K3 surface over $\mathbb{F}_q$, the Tate conjecture is known (many authors); hence $Z(X, T)$ has a pole at $T = q^{-1}$ of order equal to the Picard number. By the Artin-Tate formula, the leading coefficient is the discriminant of the Picard lattice times the order of the Brauer group; the latter is a perfect square.

Elsenhans–Jahnel observed that comparing the Artin-Tate formulas over $\mathbb{F}_q$ and $\mathbb{F}_{q^2}$ yields a previously unknown restriction on $Z(X, T)$.

van Luijk observed that if X is a K3 surface over K whose reduction to $\mathbb{F}_q$ has the same Picard number, then the discriminant does not change upon reduction. In particular, if X has two reductions with geometric Picard number 2 whose discriminants are not in the same square class, then X has geometric Picard number 1.

Question: can zeta/L-function computations be used to gain any effective control over the Brauer group of a K3 surface over K?

Contents

The Chebotarev density theorem

For L/K a Galois extension of number fields with group G, if one sorts the prime ideals of K according to their associated Frobenius conjugacy class in G (ignoring ramified primes), each class receives a density of primes equal to its size. That is, the Frobenius classes are equidistributed for the image of the Haar measure on G (a/k/a the uniform measure) on the set of conjugacy classes.

A concrete consequence of this is that if $f(x)$ is an irreducible polynomial over K, the distribution of shapes of the prime factorizations of its reductions at prime ideals of K corresponds to the distribution of cycle structures for random elements of the Galois group of f.

The Chebotarev density theorem

For L/K a Galois extension of number fields with group G, if one sorts the prime ideals of K according to their associated Frobenius conjugacy class in G (ignoring ramified primes), each class receives a density of primes equal to its size. That is, the Frobenius classes are equidistributed for the image of the Haar measure on G (a/k/a the uniform measure) on the set of conjugacy classes.

A concrete consequence of this is that if $f(x)$ is an irreducible polynomial over K, the distribution of shapes of the prime factorizations of its reductions at prime ideals of K corresponds to the distribution of cycle structures for random elements of the Galois group of f.

The Sato-Tate conjecture

For X an elliptic curve over K, write the factor of $L_1(X, s)$ for the prime ideal $\mathfrak{p}_K$ as

$$(1 - \alpha q^{1/2 - s})^{-1}(1 - \beta q^{1/2 - s})^{-1}, \qquad q = \mathrm{Norm}(\mathfrak{p}_K).$$

Then α, β lie on the unit circle and are complex conjugates of each other. As $\mathfrak{p}_K$ varies, one observes numerically that α is equidistributed for one of three measures:

- if E has CM defined over K, the uniform measure on the upper unit semicircle;
- if E has CM over a larger field, the average of the previous with a point measure at i;
- if E has no CM, the image of Haar measure on $SU(2)$.

The first two cases are due to Hecke. The third is the *Sato-Tate conjecture*, known if K is totally real or a CM field (Taylor et al.).

The Sato-Tate conjecture

For X an elliptic curve over K, write the factor of $L_1(X, s)$ for the prime ideal $\mathfrak{p}_K$ as

$$(1 - \alpha q^{1/2-s})^{-1}(1 - \beta q^{1/2-s})^{-1}, \qquad q = \text{Norm}(\mathfrak{p}_K).$$

Then α, β lie on the unit circle and are complex conjugates of each other. As $\mathfrak{p}_K$ varies, one observes numerically that α is equidistributed for one of three measures:

- if E has CM defined over K, the uniform measure on the upper unit semicircle;
- if E has CM over a larger field, the average of the previous with a point measure at i;
- if E has no CM, the image of Haar measure on $SU(2)$.

The first two cases are due to Hecke. The third is the *Sato-Tate conjecture*, known if K is totally real or a CM field (Taylor et al.).

The Sato-Tate conjecture

For X an elliptic curve over K, write the factor of $L_1(X, s)$ for the prime ideal $\mathfrak{p}_K$ as

$$(1 - \alpha q^{1/2-s})^{-1}(1 - \beta q^{1/2-s})^{-1}, \qquad q = \mathrm{Norm}(\mathfrak{p}_K).$$

Then α, β lie on the unit circle and are complex conjugates of each other. As $\mathfrak{p}_K$ varies, one observes numerically that α is equidistributed for one of three measures:

- if E has CM defined over K, the uniform measure on the upper unit semicircle;
- if E has CM over a larger field, the average of the previous with a point measure at i;
- if E has no CM, the image of Haar measure on $SU(2)$.

The first two cases are due to Hecke. The third is the *Sato-Tate conjecture*, known if K is totally real or a CM field (Taylor et al.).

The Sato-Tate conjecture

For X an elliptic curve over K, write the factor of $L_1(X, s)$ for the prime ideal $\mathfrak{p}_K$ as

$$(1 - \alpha q^{1/2-s})^{-1}(1 - \beta q^{1/2-s})^{-1}, \qquad q = \mathrm{Norm}(\mathfrak{p}_K).$$

Then α, β lie on the unit circle and are complex conjugates of each other. As $\mathfrak{p}_K$ varies, one observes numerically that α is equidistributed for one of three measures:

- if E has CM defined over K, the uniform measure on the upper unit semicircle;
- if E has CM over a larger field, the average of the previous with a point measure at i;
- if E has no CM, the image of Haar measure on $\mathrm{SU}(2)$.

The first two cases are due to Hecke. The third is the *Sato-Tate conjecture*, known if K is totally real or a CM field (Taylor et al.).

Sato-Tate groups

For any[‡] abelian variety X over K, one can define a compact Lie group, the *Sato-Tate group* of X, and a sequence of conjugacy classes associated to the factors of $L_1(X, s)$ in a similar fashion. These are expected to be equidistributed for the image of Haar measure.

For example, if X is an elliptic curve, then the Sato-Tate group is:

- $SU(2)$ if E has no CM;
- $SO(2)$ if E has CM defined over K;
- the normalizer of $SO(2)$ in $SU(2)$ if E has CM over a larger field. Note that this group has 2 connected components; E has ordinary reduction at p if and only if p corresponds to a class contained in the identity connected component.

[‡]Some complications ensue if the Mumford-Tate conjecture is not known for X, but this never happens in dimensions ≤ 3.

Sato-Tate groups

For any[‡] abelian variety X over K, one can define a compact Lie group, the *Sato-Tate group* of X, and a sequence of conjugacy classes associated to the factors of $L_1(X, s)$ in a similar fashion. These are expected to be equidistributed for the image of Haar measure.

For example, if X is an elliptic curve, then the Sato-Tate group is:

- $SU(2)$ if E has no CM;
- $SO(2)$ if E has CM defined over K;
- the normalizer of $SO(2)$ in $SU(2)$ if E has CM over a larger field. Note that this group has 2 connected components; E has ordinary reduction at p if and only if p corresponds to a class contained in the identity connected component.

[‡]Some complications ensue if the Mumford-Tate conjecture is not known for X, but this never happens in dimensions ≤ 3.

Sato-Tate groups

For any[‡] abelian variety X over K, one can define a compact Lie group, the *Sato-Tate group* of X, and a sequence of conjugacy classes associated to the factors of $L_1(X, s)$ in a similar fashion. These are expected to be equidistributed for the image of Haar measure.

For example, if X is an elliptic curve, then the Sato-Tate group is:

- $\mathrm{SU}(2)$ if E has no CM;
- $\mathrm{SO}(2)$ if E has CM defined over K;
- the normalizer of $\mathrm{SO}(2)$ in $\mathrm{SU}(2)$ if E has CM over a larger field. Note that this group has 2 connected components; E has ordinary reduction at $\mathfrak{p}$ if and only if $\mathfrak{p}$ corresponds to a class contained in the identity connected component.

[‡]Some complications ensue if the Mumford-Tate conjecture is not known for X, but this never happens in dimensions ≤ 3.

Sato-Tate groups

For any[‡] abelian variety X over K, one can define a compact Lie group, the *Sato-Tate group* of X, and a sequence of conjugacy classes associated to the factors of $L_1(X, s)$ in a similar fashion. These are expected to be equidistributed for the image of Haar measure.

For example, if X is an elliptic curve, then the Sato-Tate group is:

- $\mathrm{SU}(2)$ if E has no CM;
- $\mathrm{SO}(2)$ if E has CM defined over K;
- the normalizer of $\mathrm{SO}(2)$ in $\mathrm{SU}(2)$ if E has CM over a larger field. Note that this group has 2 connected components; E has ordinary reduction at p if and only if p corresponds to a class contained in the identity connected component.

[‡]Some complications ensue if the Mumford-Tate conjecture is not known for X, but this never happens in dimensions ≤ 3.

Sato-Tate groups

For any[‡] abelian variety X over K, one can define a compact Lie group, the *Sato-Tate group* of X, and a sequence of conjugacy classes associated to the factors of $L_1(X, s)$ in a similar fashion. These are expected to be equidistributed for the image of Haar measure.

For example, if X is an elliptic curve, then the Sato-Tate group is:

- $\mathrm{SU}(2)$ if E has no CM;
- $\mathrm{SO}(2)$ if E has CM defined over K;
- the normalizer of $\mathrm{SO}(2)$ in $\mathrm{SU}(2)$ if E has CM over a larger field. Note that this group has 2 connected components; E has ordinary reduction at $\mathfrak{p}$ if and only if $\mathfrak{p}$ corresponds to a class contained in the identity connected component.

[‡]Some complications ensue if the Mumford-Tate conjecture is not known for X, but this never happens in dimensions ≤ 3.

Sato-Tate groups of abelian surfaces

Fité–K–Rotger–Sutherland classified Sato-Tate groups for abelian surfaces; there are 52 of them, of which 34 can occur over $\mathbb{Q}$. See https://math.mit.edu/~drew/ for animations illustrating these.

The groups that do not occur over $\mathbb{Q}$ all occur over quadratic or biquadratic fields. There even exists a choice of K for which all 52 groups occur over K (Fité–Guitart).

The generic Sato-Tate group for an abelian surfaces is $\mathrm{USp}(4)$; for this group, equidistribution is not known in any case (and seems out of reach of current modularity lifting techniques). By contrast, one can prove equidistribution in most of the exceptional cases (Johansson).

Sato-Tate groups of abelian surfaces

Fité–K–Rotger–Sutherland classified Sato-Tate groups for abelian surfaces; there are 52 of them, of which 34 can occur over $\mathbb{Q}$. See https://math.mit.edu/~drew/ for animations illustrating these.

The groups that do not occur over $\mathbb{Q}$ all occur over quadratic or biquadratic fields. There even exists a choice of K for which all 52 groups occur over K (Fité–Guitart).

The generic Sato-Tate group for an abelian surfaces is $USp(4)$; for this group, equidistribution is not known in any case (and seems out of reach of current modularity lifting techniques). By contrast, one can prove equidistribution in most of the exceptional cases (Johansson).

Sato-Tate groups of abelian surfaces

Fité–K–Rotger–Sutherland classified Sato-Tate groups for abelian surfaces; there are 52 of them, of which 34 can occur over $\mathbb{Q}$. See `https://math.mit.edu/~drew/` for animations illustrating these.

The groups that do not occur over $\mathbb{Q}$ all occur over quadratic or biquadratic fields. There even exists a choice of K for which all 52 groups occur over K (Fité–Guitart).

The generic Sato-Tate group for an abelian surfaces is $\mathrm{USp}(4)$; for this group, equidistribution is not known in any case (and seems out of reach of current modularity lifting techniques). By contrast, one can prove equidistribution in most of the exceptional cases (Johansson).

Abelian threefolds

Fité–K– Sutherland are working on the corresponding result for abelian threefolds. There are approximately[§] 400 possible Sato-Tate groups.

We have no idea how many of these groups can occur over $\mathbb{Q}$. However, Sutherland has a large database of genus 3 curves over $\mathbb{Q}$ of small discriminant, for which it should be feasible to compute L-functions and match these against the classification.

[§]As of this writing, the list stands at 413 groups, but this analysis is not complete.

Abelian threefolds

Fité–K– Sutherland are working on the corresponding result for abelian threefolds. There are approximately[§] 400 possible Sato-Tate groups.

We have no idea how many of these groups can occur over $\mathbb{Q}$. However, Sutherland has a large database of genus 3 curves over $\mathbb{Q}$ of small discriminant, for which it should be feasible to compute L-functions and match these against the classification.

[§]As of this writing, the list stands at 413 groups, but this analysis is not complete.

K3 surfaces

It would be natural to classify Sato-Tate groups for K3 surfaces, in order to compare against numerical data coming from their *L*-functions (this is feasible to gather for certain low-degree shapes of K3 surfaces). For Picard rank ≥ 17, this basically reduces to the case of abelian surfaces.

This is closely related to classifying finite subgroups of $SO_n(\mathbb{Z})$ for $n \leq 21$. This is principle tractable (the maximal finite subgroups of $SL_n(\mathbb{Z})$ are known in this range) but is computationally intensive; moreover, mapping the results back to Sato-Tate groups will require some additional (human) effort.

K3 surfaces

It would be natural to classify Sato-Tate groups for K3 surfaces, in order to compare against numerical data coming from their L-functions (this is feasible to gather for certain low-degree shapes of K3 surfaces). For Picard rank ≥ 17, this basically reduces to the case of abelian surfaces.

This is closely related to classifying finite subgroups of $\mathrm{SO}_n(\mathbb{Z})$ for $n \leq 21$. This is principle tractable (the maximal finite subgroups of $\mathrm{SL}_n(\mathbb{Z})$ are known in this range) but is computationally intensive; moreover, mapping the results back to Sato-Tate groups will require some additional (human) effort.

Contents

Hypergeometric motives

One can describe many one-parameter families of *L*-functions associated to generalized hypergeometric differential equations. These are sometimes (but usually not) related to familiar classes of varieties.

These give essentially the only examples of *L*-functions in dimension ≥ 5 that are feasible to compute at present. This depends heavily on specialized *p*-adic methods for these varieties (and on work in progress).

Hypergeometric motives

One can describe many one-parameter families of L-functions associated to generalized hypergeometric differential equations. These are sometimes (but usually not) related to familiar classes of varieties.

These give essentially the only examples of L-functions in dimension ≥ 5 that are feasible to compute at present. This depends heavily on specialized p-adic methods for these varieties (and on work in progress).