A Whirlwind Tour of Elliptic Curves

- In this talk I aim to discuss as many interesting aspects of elliptic curves as possible. Note that "interesting" means (of course) "what I, the speaker, think is interesting", which likely will not – and perhaps, should not – align precisely with "what you, the audience, think is interesting".
- I will assume that you already have been told what an elliptic curve is, and have some sort of vague idea of what the group law, isogenies, torsion points, and the like, are. The goal is for the talk to be mostly self-contained but (unfortunately) the different parts of the story are sufficiently interconnected that it is not really possible to establish a satisfactory total ordering on the topics.

1 What is an "elliptic curve"?

- There are several different ways we typically think of elliptic curves.
- The very first time we mention an elliptic curve, we usually start by saying that an elliptic curve is a curve with a Weierstrass equation $y^2 = x^3 + a_4x + a_6$.
  - Approximately the next sentence after this will say something like "and one reason elliptic curves are interesting is that there is a group law, which we define geometrically by drawing a line through P and Q, finding the third intersection point, and then reflecting that point about the x-axis".
  - And perhaps a paragraph or two later someone defines the discriminant $\Delta = -16(4A^3 + 27B^2)$ and says that that cubic curve is a nonsingular elliptic curve if $\Delta \neq 0$.
  - And then a bit later someone else defines the j-invariant $j = 1728 \cdot \dfrac{(4A)^3}{\Delta}$, and says that it characterizes an elliptic curve up to isomorphism over an algebraic closure.
  - And then someone else comes along and points out that the differential $\omega = \dfrac{dx}{y}$ is a holomorphic differential on the curve. (In fact, it is even translation-invariant, meaning if that one applies the pushforward of the "add Q to every point on the curve" map to ω, one obtains just ω again.)
- The Weierstrass equation is an acceptable way to think of an elliptic curve most of the time (e.g., for computations), but it is not really the "right" way.
  - Issue #1: This is not the right form if the characteristic is 2 or 3. The most general form should properly be $y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6$. If the characteristic is not 2 we can complete the squares on the left, and if the characteristic is not 3 we can depress the cube on the right.
  - Issue #2: An elliptic curve has a point at infinity which is extremely relevant for the group law (being the identity and all), and so we should probably be thinking about an elliptic curve as a projective curve. So we should actually be writing it with projective coordinates, as something like $Y^2Z = X^3 + a_4XZ^2 + a_6Z^3$.
  - Issue #3: What does it mean to say that two elliptic curves are "the same"? Do we take that to mean that their Weierstrass equations are the same (i.e., identical), or maybe that there is some change of coordinates taking one to the other, and if the latter, where does the change of coordinates need to be defined? (And so on.)
  - Issue #4: Proving things about elliptic curves using Weierstrass equations generally requires a ton of algebra. Canonical example: show that the group law is associative using the Weierstrass equation. (Another example: try showing that the multiplication-by-m map has degree m^2.) There also frequently arise annoying special cases one must deal with when using the group law, such as using a tangent line to double a point.
- One can work out most everything one would ever want to say about elliptic curves just by talking about Weierstrass forms, but this rapidly becomes annoying for a variety of reasons.

2 The Right Definition

- It turns out that the "right" way to define an elliptic curve is the following:
- An elliptic curve is a pair (E, O), where E is a smooth projective curve of genus 1 and $O \in E$. (The point O will be the identity under the group law. Usually we will be lazy and not specify what O is.)
  - We say E is defined over K if $E(K)$, the K-points of E, is a curve, and contains O.
  - This solves (at least) the issue of saying when two elliptic curves are "the same": namely, if there is a birational equivalence of smooth projective curves sending the marked point to the marked point.
- Using Riemann-Roch we can, rather magically, get a Weierstrass equation out of this definition.

3 A Brief Review of Divisors and Riemann-Roch

- If C is a smooth curve, the divisor group $\mathrm{Div}(C)$ is the free abelian group generated by the points of C. Elements look like $\sum_{P \in C} n_p(P)$ where the n_p are integers all but finitely many of which are zero.
- The degree of $\sum_{P \in C} n_p(P)$ is defined to be $\sum_{P \in C} n_p$, the sum of the coefficients. The kernel of the degree map is the set of degree-zero divisors $\mathrm{Div}^0(C)$.
- If f is a nonzero function on C, then we define $\mathrm{div}(f) = \sum_{P \in C} \mathrm{ord}_P(f)\,(P)$, where $\mathrm{ord}_P(f)$ denotes the degree of zeroness or poleness of f at P (if f has a zero at P the order is degree of the zero, and if f has a pole at P is -1 times the degree of the pole). This is a divisor, in fact a degree-zero divisor, because nonzero functions on smooth projective curves have equal and finite numbers of zeroes and poles.
- We call a divisor principal if it is the divisor of some nonzero function. Principal divisors form a subgroup of $\mathrm{Div}(C)$ (since $\mathrm{div}(1) = 0$ and $\mathrm{div}(f) - \mathrm{div}(g) = \mathrm{div}(f/g)$), and so we can form the quotient of divisors modulo principal divisors, called the divisor class group or the Picard group, denoted $\mathrm{Pic}(C)$. We call two divisors equivalent if they are the same in the Picard group (i.e., if they differ by a principal divisor).
  - The degree-zero part of the Picard group is denoted $\mathrm{Pic}^0(C)$. We have an exact sequence $1 \to \bar{K}^* \to \bar{K}(C)^* \to \mathrm{Div}^0(C) \to \mathrm{Pic}^0(C) \to 1$, which is the function-field analogue of the exact sequence $1 \to \{\text{units}\} \to K^* \to \{\text{fractional ideals}\} \to \mathrm{Cl}(K) \to 1$.
- If $\omega$is a nonzero differential on C, we can likewise associate a divisor to it by computing what amounts to its order of vanishing / poleness, but which is annoying to define precisely. It turns out that up to equivalence (in the Picard group) all nonzero differentials have the same divisor, which we will call the canonical divisor K_C.
- We say a divisor is effective, writing $D \geq 0$, if all coefficients of D are nonnegative, and we extend this to a partial ordering of all divisors.
- For C defined over K, for a divisor D, we define the space $\mathcal{L}(D)$ to be finite-dimensional $\bar{K}$-vector space containing 0 and the nonzero functions f in $\bar{K}(C)$ with $\mathrm{div}(f) \geq -D$, and we set $l(D)$ to be its dimension. [This is secretly a line bundle but I don't want to get too much into it.]
  - Note that if $\deg(D) < 0$, then $l(D) = 0$, because if $f \in L(D)$ is nonzero then $0 = \deg(\mathrm{div}(f)) \geq \deg(-D) = -\deg(D)$.
- Riemann-Roch then says the following: If C is a smooth curve of genus g, and K_C is the canonical divisor class associated to any nonzero differential, and D is any divisor, then $l(D) - l(K_C - D) = \deg(D) - g + 1$.
  - By setting $D = 0$ one sees that $l(K_C) = g$, and then by setting $D = K_C$ one sees $\deg(K_C) = 2g-2$. Then, for any other divisor of degree larger than $2g-2$, one has that $\deg(K_C - D) < 0$, hence $l(K_C - D) = 0$. Thus if D has degree $> 2g-2$, then $l(D) = \deg(D) - g + 1$.

4 A Smooth Genus-1 Curve Has a Weierstrass Equation (or, Riemann-Roch Rocks)

- Theorem: If (E, O) is an elliptic curve defined over K, then there exist functions $x, y \in K(E)$ such that the map $\phi : E \to \mathbb{P}^2$ sending P to $[x(P) : y(P) : 1]$ is an isomorphism of E/K onto a Weierstrass curve with coefficients in K which sends O to the point $[0 : 1 : 0]$. And conversely, any smooth cubic curve given by a Weierstrass equation is an elliptic curve. In other words, "elliptic curves are given by Weierstrass equations".
- Proof: Consider the vector spaces $L(n(O))$ for $n = 1, 2, \cdots$. By the remarks I made just after Riemann-Roch, we know that $l(n(O)) = n$ for all $n \geq 1$, because $n(O)$ is a divisor of degree larger than $2g - 2 = 0$.
  - $L(2(O))$ is 2-dimensional. The function 1 is in here, so extend to a basis by choosing another function x. By construction, x must have a pole of exact order 2 at O, else x would be in $L(1(O))$.
  - $L(3(O))$ is 3-dimensional. The functions 1 and x are in here, so extend to a basis by choosing another function y. By the same logic, y must have a pole of exact order 3 at O.
  - Now $L(6(O))$ is 6-dimensional. But here are seven functions that are in it: $1, x, x^2, x^3, y, xy, y^2$. They must therefore be linearly dependent. In this dependence, neither of the coefficients of x^3 and y^2 can be zero, as otherwise every other term in the dependence would have a different order of pole at O. Now rescale, and check this is an isomorphism.
- Now, we can also use Riemann-Roch to prove that the points on an elliptic curve satisfy a group law. This method has the advantage of being far less of a pain than doing the group law directly.
- Theorem: For every degree-zero divisor D there exists a unique point $P \in E$ such that $D \sim (P) - (O)$. If $\sigma : \mathrm{Div}^0(E) \to E$ denotes this map, then σ is a bijection between $\mathrm{Pic}^0(E)$ and E, and the group law induced from $\mathrm{Pic}^0(E)$ via σ is the same as the geometric group law. (In other words, if we think of E as a group with the geometric law, then $E \cong \mathrm{Pic}^0(E)$ via σ.)
  - For the existence of such a P, Riemann-Roch says that $\dim(L(D + (O))) = 1$. Take f a generator; then $\mathrm{div}(f) \geq -D - (O)$ and $\deg(\mathrm{div}(f)) = 0$, so $\mathrm{div}(f) = -D - (O) + (P)$ for some P, whence $D \sim (P) - (O)$.
  - Uniqueness follows from the fact that if C has genus 1, then (as divisors) $(P) \sim (Q)$ iff $P = Q$: this holds because if $\mathrm{div}(f) = (P) - (Q)$ then $f \in L((Q))$. But Riemann-Roch says that $\dim(L((Q))) = 1$, so f must be constant so that $P = Q$.
  - $\sigma((P) - (O)) = P$ so σ is certainly surjective on Div^0. By the definition of σ we have $(\sigma(D_1)) - (\sigma(D_2)) \sim D_1 - D_2$, so $D_1 \sim D_2$ if and only if $\sigma(D_1) = \sigma(D_2)$, which shows that σ is a bijection from $\mathrm{Pic}^0(E)$ to E.
  - On $\mathrm{Pic}^0(E)$, the inverse map of σ is $\kappa : P \to (P) - (O)$. We want to see that $\kappa(P + Q) = \kappa(P) + \kappa(Q)$, where the addition on the left is the geometric group law, and the addition on the right is the addition of divisor classes in the Picard group. Equivalently, we want to see that $(P + Q) - (P) - (Q) + (O) \sim 0$.
  - Let f be the line through P and Q, let R be the third intersection point of E with this line, and let f' be the line through R and O. Then since the line $Z = 0$ intersects E at O with multiplicity 3, we have $\mathrm{div}(f/Z) = (P) + (Q) + (R) - 3(O)$ and $\mathrm{div}(f'/Z) = (R) + (P + Q) - 2(O)$.
  - Therefore, $(P + Q) - (P) - (Q) + (O) = \mathrm{div}(f/f') \sim 0$. So we are done.
- Theorem: The group law defines morphisms $+ : E \times E \to E$ and $- : E \to E$.
  - Proof: It is enough to show that the maps are rational, since rational maps from a smooth curve to a variety are morphisms. Then just check, using the formulas for the maps.

5 Complex Elliptic Curves as Lattices

- Let us now think about elliptic curves over the complex numbers. Let me give some extremely sketchy motivation:

- Over $\mathbb{C}$, an elliptic curve is a smooth curve of genus 1. Geometrically, if we think of $\mathbb{C}$ as being 2-dimensional over $\mathbb{R}$, this more or less says an elliptic curve is an orientable 2-dimensional real manifold of genus 1, which, by the classification of surfaces, means it is a torus.
- So how do we get a torus from $\mathbb{C}$? Answer: a torus is $\mathbb{C}/\Lambda$, where Λ is a discrete rank-2 lattice in $\mathbb{C}$. So this is what we should be getting.

- Now imagine we have an elliptic curve with a Weierstrass equation $y^2 = (x - r_1)(x - r_2)(x - r_3)$ over the complex numbers. We know that $\omega = \dfrac{dx}{y} = \dfrac{dx}{\sqrt{(x - r_1)(x - r_2)(x - r_3)}}$ is a holomorphic differential on E, and so we can try to integrate this differential in the complex plane to get a map from E to $\mathbb{C}$, by sending a point P to the integral $\int_0^P \omega$.
  - Historical note: The integral $\int \dfrac{dx}{\sqrt{(x - r_1)(x - r_2)(x - r_3)}}$ is, after some changes in variable, what one obtains after trying to calculate the arclength of an ellipse. This is the reason for the terms "elliptic integral", "elliptic function", and "elliptic curve".
- The problem is that this integral is not well-defined since this function needs branch cuts, and since we really want to do things projectively we should do this on the Riemann sphere. So make one branch cut from r_1 to r_2 and another from r_3 to ∞; topologically, this turns the Riemann sphere into a torus.
- Let α be a path looping around the r_1-r_2 branch cut once, and let β be a path looping around the $r_3 - \infty$ branch cut once. Since α and β generate $H^1(T)$, the difference between any two paths between 0 and P on our branch-cut Riemann sphere is homotopic to a linear combination of α and β.
- So the integral $\int_0^P \omega$ is well-defined up to adding a linear combination of $\omega_1 = \int_\alpha \omega$ and $\omega_2 = \int_\beta \omega$. Ergo we have a well-defined map $E(\mathbb{C}) \to \mathbb{C}/\Lambda$ via $P \to \int_0^P \omega \pmod{\Lambda}$, where $\Lambda = \langle \omega_1, \omega_2 \rangle$. As we should expect (and is indeed true) the two period elements are $\mathbb{R}$-linearly independent, so this really is a discrete lattice.
- The fact that ω is translation-invariant implies that this map is an isomorphism of groups: $\int_0^{P+Q} \omega = \int_0^P \omega + \int_P^{P+Q} \omega = \int_0^P \omega + \int_0^Q \omega$.
- One can go off and show, in fact, that this map from $E(\mathbb{C})$ to $\mathbb{C}/\Lambda$ is a complex-analytic isomorphism.
- So now let us discuss functions on lattices.
- An elliptic function (relative to a lattice Λ) is a meromorphic function on $\mathbb{C}$ which satisfies $f(z + \omega) = f(z)$ for all $\omega \in \Lambda$ and $z \in \mathbb{C}$.
  - There are many nice things that one can say about elliptic functions: for example, an elliptic function has the same number of zeroes as poles, and if it has no poles (and thus no zeroes) it must be constant. There is also no elliptic function that has a single simple pole.
- Here is the standard example of an elliptic function: the Weierstrass $\wp$-function is defined to be the series $\wp(z; \Lambda) = z^{-2} + \sum_{\omega \in \Lambda *} \left[\dfrac{1}{(z - \omega)^2} - \dfrac{1}{\omega^2} \right]$.
- It is a (not too hard) theorem that, in fact, every elliptic function is a rational function in $\wp$ and $\wp'$.
- The Laurent series for $\wp$ around $z = 0$ is given by $\wp(z) = z^{-2} + \sum_{k=1}^{\infty} (2k + 1) G_{2k+2} z^{2k}$, where $G_{2k}(\Lambda) = \sum_{\omega \in \Lambda *} \omega^{-2k}$ is the so-called Eisenstein series of weight $2k$ (for Λ).
- By comparing Laurent expansions, one can show that $\wp'(z)^2 = 4\wp(z)^3 - g_2 \wp(z) - g_3$, where $g_2 = 60 G_4$ and $g_3 = 140 G_6$.
  - One needs only to compute the terms with a nonpositive power of z and check that they are equal on both sides, for then the difference between left and right is an elliptic function with no poles which is zero at $z = 0$.

- That expression is very (!!) suggestive. In fact, if Λ is any lattice, then the map $\mathbb{C}/\Lambda \to E$ sending $z \to [\wp(z) : \wp'(z) : 1]$, where E is the elliptic curve with affine equation $y^2 = 4x^3 - g_2x - g_3$, is a complex-analytic isomorphism of complex Lie groups.
- After checking a few additional details about the maps involved, one obtains the following equivalence of categories: [Elliptic curves / $\mathbb{C}$, Isogenies] $\leftrightarrows$ [Lattices up to homothety, maps from Λ_1 to Λ_2 with $\alpha\Lambda_1 \subset \Lambda_2$].
- Now let us use this very nice characterization of elliptic curves as lattices to say some very general things:
    - The endomorphism ring of E is either $\mathbb{Z}$ or an order in an imaginary quadratic extension of $\mathbb{Q}$ (and in the latter case, if ω_1 and ω_2 generate the lattice, then $\mathbb{Q}(\omega_2/\omega_1)$ is imaginary quadratic and the order lives in there): follows by looking at what kinds of maps send a lattice into itself.
        * In the latter case, we say that the elliptic curve has <u>complex multiplication</u> (the name coming from the fact that, if we view E as a lattice, the endomorphisms are literally multiplication by a complex number). There is an extremely rich theory of such CM curves: for example, if $K = \mathbb{Q}(\sqrt{-D})$ then there are exactly $h(K)$ curves up to isomorphism whose endomorphism ring is the maximal order $\mathcal{O}_K$. Furthermore, $K(j(E))$ is the Hilbert class field of K.

6 Maps Between Elliptic Curves

- Let's now return to the general setting of elliptic curves, keeping in mind this lattice structure that morally ought to be lurking in the background.
- Whenever we have objects, we should study maps between them.
- If E_1 and E_2 are elliptic curves, an <u>isogeny</u> from E_1 to E_2 is a morphism $\phi : E_1 \to E_2$ satisfying $\phi(O) = O$. By morphism-ness of curves, if an isogeny is not the zero map, then it is surjective.
    - We say that E_1 and E_2 are isogeneous if there is an isogeny between them.
    - The terminology suggests that "isogenous" should be an equivalence relation, which is in fact true, but not at all obvious from the definition.
- A natural family of isogenies is the multiplication-by-m maps: for $m > 0$, $[m]P = P + P + \cdots + P$, where we add m of them, and for $m < 0$, $[m]P = [-m](-P)$.
    - The multiplication-by-m map, being a morphism, has some finite degree. Using a thoroughly atrocious mess of algebra, or cleverness, one can prove that the degree of the multiplication-by-m map has degree m^2.
    - The points P of E for which $[m]P = O$ are called the <u>m-torsion</u> points of E, and denoted $E[m]$.
- If we are over a finite field $\mathbb{F}_q$, the Frobenius map (taking everything to its qth power) is another example of an isogeny.
    - The degree of Frobenius is q. (This is fairly clear.)
- Isogenies are very nice. Here are some properties:
    - Every isogeny is a group homomorphism.
    - The kernel of a nonzero isogeny is finite.
    - The size of the kernel of a separable (nonzero) isogeny is equal to its degree.
    - If E is an elliptic curve and Φ is a finite subgroup, then there is a unique E' and separable isogeny $\phi : E \to E'$ has kernel Φ. (If E is defined over K and Φ is absolute-Galois-invariant, then the isogeny and target curve can be defined over K too.)
- What we would like is for "isogeny" to be an equivalence relation. We can show this by constructing what is called the dual isogeny:

- For any nonconstant isogeny $\phi : E_1 \to E_2$ of degree m, there exists a unique isogeny $\hat{\phi} : E_2 \to E_1$, called the dual isogeny: satisfying $\hat{\phi} \circ \phi = [m]$.
    - The exact construction is natural (one uses the action of ϕ^* on the divisor groups) but not extremely enlightening so I won't actually write it down.
    - The composition of ϕ and $\hat{\phi}$ in either order is multiplication by m (on the appropriate curve).
    - The composition of dual isogenies is the dual of the composition in reverse order.
    - The sum of duals is the dual of the sum.
    - The double dual is the original isogeny.
    - The degree of the dual and of the original isogeny are equal.
    - The multiplication-by-m map is its own dual, and thus $\deg[m] = m^2$.
- Using the above one can show that if E/K is an elliptic curve, then if $\mathrm{char}(K) = 0$ or m is prime to $\mathrm{char}(K)$ then $E[m] \cong (\mathbb{Z}/m\mathbb{Z})^2$, and if $\mathrm{char}(K) = p$ then $E[p^e]$ is either trivial or $\mathbb{Z}/p^e\mathbb{Z}$ (depending on whether the curve is supersingular or ordinary, respectively).
    - In the complex setting, if we think of E as $\mathbb{C}/\Lambda$, then the m-torsion of E is very clearly isomorphic to $(\mathbb{Z}/m\mathbb{Z})^2$, because the m-torsion of Λ is just $\frac{1}{m}\Lambda$.
    - It is also very clear that the multiplication by m map has degree m^2, because the degree of $[m]$ is just the number of points in $[m]^{-1}\{O\} = E[m]$, which is m^2.
- Now, the m-torsion points of E as a finite group are isomorphic to $(\mathbb{Z}/m\mathbb{Z})^2$. But there is far more to say about them: if E is defined over a field K, then the absolute Galois group $\mathrm{Gal}(\bar{K}/K)$ also acts on the m-torsion points, since if $[m]P = O$ then $[m](P^\sigma) = ([m]P)^\sigma = O$ too.
    - Therefore, we obtain a mod-m representation $\bar{\rho}_{E,m} : \mathrm{Gal}(\bar{K}/K) \to \mathrm{Aut}(E[m])$ of the absolute Galois group into $\mathrm{Aut}(E[m]) \cong GL_2(\mathbb{Z}/m\mathbb{Z})$, after choosing a basis.
    - We could study these representations individually. But this involves doing things in finite characteristic, which (for representation theory) can be really annoying.
- So let us instead define the Tate module $T_l(E) = \lim_{\leftarrow n} E[l^n]$. Each of the pieces has a $\mathbb{Z}/l^n\mathbb{Z}$-structure, so the Tate module has a natural $\mathbb{Z}_l$-structure.
    - By what we just saw about torsion elements, we see that $T_l(E) \cong \mathbb{Z}_l \times \mathbb{Z}_l$ if $l \neq \mathrm{char}(K)$, and is either $\mathbb{Z}_p$ or 0 otherwise.
    - The action of Galois commutes with the multiplication-by-l maps used for the inverse limit so Galois also acts on T_l.
    - So we get an l-adic representation $\rho_{E,l} : \mathrm{Gal}(\bar{K}/K) \to \mathrm{Aut}(T_l(E))$. By thinking of $\mathbb{Z}_l$ inside $\mathbb{Q}_l$, we get a representation of the Galois group over a field of characteristic 0.
- If ϕ is an isogeny, then since isogenies are group homomorphisms we see that ϕ sends m-torsion to m-torsion, and so ϕ induces a map on the Tate modules. So we obtain a homomorphism $\mathrm{Hom}(E_1, E_2) \to \mathrm{Hom}(T_l(E_1), T_l(E_2))$.
    - This map is injective: in fact, if we tensor up to $\mathbb{Z}_l$, this map remains injective. If K is a number field or a finite field, the map $\mathrm{Hom}(E_1, E_2) \otimes \mathbb{Z}_l \to \mathrm{Hom}(T_l(E_1), T_l(E_2))$ is actually an isomorphism, although this fact is very hard. Morally, to understand this, one should view the Tate module as $H^1(E; \mathbb{Z}_l)$; then the theorem is a statement about when a map on homology groups actually comes from an honest geometric map.
- One natural question might ask is: what does the image of $\rho_l(G)$ look like in $\mathrm{Aut}(T_l(E))$?
    - A theorem of Serre says that, if E/K is a non-CM curve and K is a number field, then the image is of finite index for all primes, and that for all but finitely many primes, the image is all of $\mathrm{Aut}(T_l(E))$.
    - If the curve has CM, the extra automorphisms force the action of Galois on the Tate module to be abelian, so a finite-index result cannot hold for CM curves.

7 The Hasse Bound and Weil Conjectures

- Now let us very briefly mention some things about elliptic curves over finite fields: let $E/\mathbb{F}_q$.
- Theorem (Hasse bound): We have the inequality $|\#E(\mathbb{F}_q) - (q+1)| \leq 2\sqrt{q}$.
  - Proof: Because the absolute Galois group $\bar{\mathbb{F}}_q/\mathbb{F}_q$ is topologically generated by Frobenius, a point $P \in E(\bar{\mathbb{F}}_q)$ is an $\mathbb{F}_q$-point if and only if it is fixed by the Frobenius map Frob_q: that is, if $\mathrm{Frob}_q(P) = P$.
  - Therefore, $E(\mathbb{F}_q) = \ker(1 - \mathrm{Frob}_p)$. So we want to know the size of the kernel of $1 - \mathrm{Frob}_p$.
  - We have $(1 - \mathrm{Frob}_p)(1 - \hat{\mathrm{Frob}}_p) = 1 - (\mathrm{Frob}_p + \hat{\mathrm{Frob}}_p) + \mathrm{Frob}_p \hat{\mathrm{Frob}}_p = 1 - \mathrm{tr}(\mathrm{Frob}_p) + q$, since we know that the degree of Frobenius is q.
  - Then we apply the Cauchy-Schwarz inequality to see that $|\mathrm{tr}(\mathrm{Frob}_p)| \leq \sqrt{q}$.
- More generally, one could count the $\mathbb{F}_{q^n}$ points of E: that would be the number of elements in the kernel of $1 - \mathrm{Frob}_p^n$, which, by the same observation, is $1 - \mathrm{tr}(\mathrm{Frob}_p^n) + q^n$.
  - For an elliptic curve over $\mathbb{F}_q$, Frobenius acts on the Tate module T_l for any l as a 2×2 matrix.
  - It turns out that, in fact, if we compute the determinant and trace of Frobenius on T_l for any l, both are actually in $\mathbb{Z}$, and independent of l. One can also verify that the eigenvalues of Frobenius are complex conjugates α and β whose product is q.
  - Therefore, $\#\ker(1 - \mathrm{Frob}_p^n) = 1 + q^n - \mathrm{tr}(\mathrm{Frob}_p^n) = 1 + q^n - \mathrm{tr}(\mathrm{Frob}_p^n \,|T_l) = 1 - \alpha^n - \beta^n + q^n$.
- This essentially proves the Weil conjectures for elliptic curves. Here is the statement:
  - If V/K is a projective variety where $K = \mathbb{F}_q$ and $K_n = \mathbb{F}_{q^n}$, then the (congruence) zeta function of V is defined as $Z(V/K;T) = \exp\left[\sum_{n=1}^{\infty} (\#(V(K_n)) \cdot \frac{T^n}{n}\right]$.
  - The Weil conjectures say that, for any smooth projective variety (i) the zeta function is a rational function in T, (ii) there is a functional equation, (iii) and that the zeta function factors as a product of terms $\frac{P_1(T)\cdots P_{2n-1}(T)}{P_0(T)\cdots P_{2n}(T)}$ where each P_i has integral coefficients and factors as a product $\prod(1 - \alpha_{i,j}T)$ where $|\alpha_{i,j}| = q^{i/2}$.
- From the calculation above, we can plug in and show that $Z(E/K;T) = \frac{1 - \mathrm{tr}(\mathrm{Frob}_p)T + qT^2}{(1-T)(1-qT)}$.

8 Other Things To Know

- There are many, many, many other interesting things about elliptic curves. Here are some of them:
  - The relation between elliptic curves and modular forms, modular functions, modular curves, etc.
  - The theory of elliptic curves with complex multiplication.
  - More about elliptic curves over finite fields: endomorphism rings and quaternion algebras, supersingular curves
  - Integral points on elliptic curves: Siegel's theorem, other theorems.
  - Elliptic curves over local fields, good and bad reduction, Neron-Ogg-Shafarevich.
  - Elliptic curves over global fields: Mordell-Weil, descent, naive heights and the canonical height, rank of elliptic curves, the Birch and Swinnerton-Dyer conjecture.
  - Mazur's theorem, for torsion points of elliptic curves defined over $\mathbb{Q}$.
  - Elliptic curve cryptography and Lenstra's factorization algorithm.