General Elliptic Curves:

General Definition: Let $\boldsymbol{F}$ be a field. A *general elliptic curve* is the solution set $V_P \subset P^2(\boldsymbol{F})$ to a degree 3 homogeneous polynomial P provided that.

1. V_P is nonsingular: If $v \in \boldsymbol{F}^3$ is a nonzero and $P(v) = 0$ then $\nabla P(v) \neq 0$.
2. P is irreducible: P does not factor into lower degree homogeneous polynomials. In other words, V_P is not the union of a line and a conic.

I have been a bit sloppy about not mentioning this second condition. Sometimes the first condition rules out the second condition and sometimes it doesn't. Here are some examples of polynomials which violate one or both conditions.

- If P is the product of 3 degree 1 polynomials then V_P is a union of 3 lines. V_P will be singular at the places where the lines intersect. This example works in any field.
- Suppose that $\boldsymbol{F} = \boldsymbol{R}$ and $P = (z)(x^2 + y^2 - z^2)$. In this case V_P is the union of the line at infinity and the unit circle, and V_P is non-singular.
- Consider the previous example with $\boldsymbol{F} = \boldsymbol{C}$. In this case, both P and ∇P vanish at $(1, i, 0)$. What is going on is that this point is a place where the conic and the line intersect.

Inflection Points: Let V_P be an elliptic curve. An *inflection point* on V_P is a point $p \in V_P$ which P contains with multiplicity 3. Geometrically this means that the tangent line to V_P at p only intersects P at p. There is also an algebraic characterization. (As in the previous notes) we can move the picture by a projective transformation (which doesn't effect the multiplicities) so that the point is $[0 : 1 : 0]$ and the tangent line is given by $Z = 0$. In this case the multiplicity of $[0 : 1 : 0]$ is defined to be the multiplicity of the root $x = 0$ of the equation $f(x) = P(x, 1, 0)$.

Here is an example. Consider the Weierstrass elliptic curve corresponding to

$$P(x, y, z) = y^2 z - x^3 - axz^2 - bz^3.$$

We have $P(0, 1, 0) = 0$ and

$$\nabla P(0, 1, 0) = (0, 0, 1).$$

So $[0:1:0] \in V_P$ and the tangent line to this point is given by $Z = 0$. This line only intersects V_P at $[0:1:0]$ so this point is an inflection point. We can also see this algebraically: $P(x, 1, 0) = x^3$.

When we work over $\boldsymbol{R}$ we can give another geometric characterizaton of inflection points. Call V_P *locally convex* at p if all points of V_P sufficiently close to p lie on the same side of the tangent line L_p. Figure 1 shows what local convexity looks like. V_P is locally convex at the black point but not at the red point.

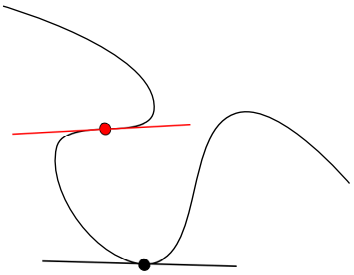

Figure 1: Local convexity and local non-convexity

Lemma 0.1 *If $p \in V_P$ is an inflection point if and only if V_P is not locally convex at p.*

Proof: For the purposes of working in the ordinary plane, we swap the roles of the y and z coordinates. We move by a projective transformation so that $p = [0:0:1]$ and L_p is given by $Y = 0$. That is, the tangent line is the x-axis.

Suppose first that p is not an inflection point. Then 0 is a double (but not triple) root of $f(x) = P(x, 0, 1)$. This means that $f''(0) \neq 0$. Let's consider the case when $f''(0) > 0$ and ∇P is a positive multiple of $(0, 1, 0)$. The other cases are similar. In this situation, $f(x) \geq 0$ for x near 0. Also, if all of x, y, y' are sufficiently near 0 and $y < y'$ then $P(x, y, 1) < P(x, y', 1)$. This comes from the fact that the directional derivative of P in the vertical direction is positive near $(0, 0)$.

Suppose some point of V_P very near $(0, 0)$ lies above the x-axis. Then $P(x, y, 1) = 0$ for some pair x, y very near 0 and $y > 0$. But $P(x, y, 1) > P(x, 0, 1)$, forcing $f(x) < 0$. This contradicts the fact that $f(x) > 0$ for x sufficiently near 0. So, near $(0, 0)$ all points of V_P lie below the x-axis.

When p is an inflection point, 0 is a $f(x)$ takes on both signs near 0. The same argument shows that V_P lies on both sides of the x-axis near $(0, 0)$. ♠

Existence of Inflection Points: The reason we want inflection points to exist on elliptic curves is that, in order to define the group law, we need to let 0 be an inflection point. We have already seen that $[0:1:0]$ is an inflection point for any Weierstrass elliptic curve with respect to any field. This is one reason these are so nice. In this section I'll sketch an argument that a general elliptic curve over $\boldsymbol{R}$ has an inflection point. This implies, of course, that a general elliptic curve over $\boldsymbol{C}$ also has an inflection point. A more sophisticated algebraic arguement would show that a general elliptic curve over $\boldsymbol{C}$ has exactly 9 inflection points.

Let V_P be a general elliptic curve over $\boldsymbol{R}$. Assume that V_P has no inflection points. The following lemma works even when V_P is assumed to have inflection points, but we'll use the extra hypothesis to make proof simpler. (All we need is one point which is not an inflection point for this to work.)

Lemma 0.2 *There exists a line which intersects V_P exactly once and is not tangent to V_P.*

Proof: Let L_p be any tangent line. Since p is not an inflection point, we know that locally V_P lies on one side of L_p. Let q be the other intersection point of $L_p \cap V_P$. If we consider lines L_p^* that are near L_p and contain q, then those on one side of L_p will intersect V_P in three points and those on the other will intersect V_P in one point. In particular, we can find a line which interects V_P once and is not tangent to V_P at the intersection point. Figure 2 shows the construction. ♠

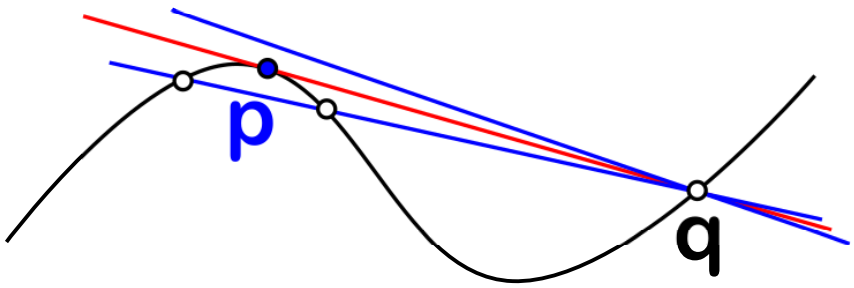

Figure 2: Wiggling the line.

We move the picture by a projective transformation so that the line from the lemma is the line at infinity. This implies that $V_P \cap \boldsymbol{R}^2$ has one unbounded connected component and a finite union of bounded components.

We'll suppose first that $V_P \cap \boldsymbol{R}^2$ has no bounded components. So, $V_P \cap \boldsymbol{R}^2$ is just one connected unbounded component. We orient the tangent lines to $V_P \cap \boldsymbol{R}^2$ so that the orientation varies continuously. Locally V_P either lies to the left or to the right of these tangent lines. By continuity, the answer (left/right) cannot switch. So, let's say that locally V_P always lies on the left. Suppose that some line intersects $V_P \cap \boldsymbol{R}^2$ in 3 points. Then we get the contradiction shown in Figure 3: Some tangent line locally contains V_P on the right and some other line locally contains V_P on the left. This is just the mean value theorem in action.

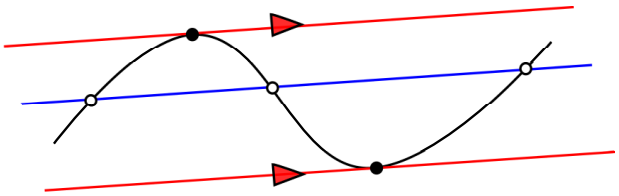

Figure 3: The mean value theorem in action

Since every line which intersects an elliptic curve twice also intersects it a third time, we see that any line containing two points of $V_P \cap \boldsymbol{R}^2$ which is not tangent to V_P (i.e. a secant line) must also intersect the line at infinity. Hence V_P contains infinitely many points at infinity. But then V_P contains the whole line at infinity, and is not irreducible. This is a contradiction.

It remains to consider the case when $V_P \cap \boldsymbol{R}$ has both an unbounded component and a bounded component. The argument above shows that there must exist a line which intersects the unbounded component in two points and the bounded component in at least one point. But any such line will either intersect the bounded component in a second point or will be tangent to it. In either case, we have produced a line which intersects V_P four times, counting multiplicity. This is a contradiction.

Equivalence with Weierstrass Curves: Suppose now that $\boldsymbol{F}$ is any field and V_P is an elliptic curve which has an inflection point. When $\boldsymbol{F} = \boldsymbol{R}$ or $\boldsymbol{F} = \boldsymbol{C}$ this is always the case, thanks to the argument in the previous section. We call V_P *nice* if $[0:1:0]$ is an inflection point and the line $Z = 0$ is the tangent line. So, any elliptic curve with an inflection point is projectively equivalent to a nice one.

Lemma 0.3 *A nice elliptic curve has the equation*

$$Ax^3 + Cz^3 + Fx^2z + Gxz^2 + Hy^2z + Iyz^2 + Jxyz$$

with $A \neq 0$.

Proof: The general equation is $P(x,y,z) =$

$$Ax^3 + By^3 + Cz^3 + Dx^2y + Exy^2 + Fx^2z + Gxz^2 + Hy^2z + Iyz^2 + Jxyz.$$

The conditions imply right away that 0 is a triple root of the polynomial

$$P(x,1,0) = Ax^3 + Dx^2 + Ex + B.$$

When 0 is a triple root, this forces $B = D = E = 0$ and $A \neq 0$. ♠

We call the nice elliptic curve *excellent* if $H \neq 0$. If we work over $\boldsymbol{R}$ or $\boldsymbol{C}$, almost every nice elliptic curve is excellent in the sense that if we pick the coefficients at random we will get $H \neq 0$. Put another, the set of nice elliptic curves forms a 7-dimensional vector space and the set of nice elliptic curves which are not excellent forms a 6-dimensional vector space. So, in any reasonable sense of the word, most of the nice ones are excellent.

Theorem 0.4 *Every excellent elliptic curve is projectively equivalent to a Weierstrass elliptic curve.*

Proof: All the transformations we make amount to composing with a projective transformation. The substitution $(x,y,z) \to (\alpha x, \beta y, z)$ for suitable α, β leads to $A = H = 1$. The transformation $(x,y,z) \to (x, y + \lambda x, z)$ leads to

$$x^3 + Cz^3 + F'x^2z + G'xz^2 + y^2z + Iyz^2 + (J + 2\lambda)xyz.$$

Here F' and G' are the new coefficients. (This has nothing to do with taking derivatives.) Taking $\lambda = -J/2$ leads to

$$x^3 + Cz^3 + F'x^2z + G'xz^2 + y^2z + Iyz^2.$$

A suitable transformation $(x,y,z) \to (x + \alpha z, y + \beta z, z)$ kills off F' and I. This gives

$$x^3 + C'z^3 + G''xz^2 + y^2z.$$

This is an equation for a Weierstrass elliptic curve. ♠

Remark: When $H = 0$ we can use projective transformations to get to the equation

$$yz^2 + cxyz - (x^3 + axz^2 + bz^3).$$

I think that this is not projectively equivalent to a Weierstrass elliptic curve, but it should be birationally equivalent to one.

A Mystery Explained: The results above do not seem to square with a dimension count. The space of general cubic curves is 9 dimensional and the (Lie) group of projective transformations is 8 dimensional. Above we showed that we can map almost every general elliptic curve to a Weierstrass elliptic curve, but this does not seem to square with the fact that there is a 2-parameter family of W.-elliptic curves. What is going on?

Let $E(a, b)$ be the W. elliptic curve given by the equation

$$y^2z - (x^3 + axz^2 + bz^3).$$

The solution to the mystery is that $E(a, b)$ and $E(\lambda^2 a, \lambda^3 b)$ are projectively equivalent. The idea is that the change of variables

$$(x, y, z) \to (x/\lambda, y/\lambda^{3/2}, z)$$

transforms the equation above to

$$\lambda^{-3}y^2z - \lambda^{-3}(x^3 + \lambda^2 xz^2 + \lambda^3 z^3).$$

Cancelling out the λ^{-3}, which does not change the solution set, gives the equation for $E(\lambda^2 a, \lambda^3 b)$. So, what is going on is that there is only a 1-dimensional set of Weierstrass elliptic curves up to projective transformations.