Lecture 1: Review of holomorhic functions

September 3, 2014

The subject of this course is Riemann surfaces. A one sentence "definition" of a Riemann surface is that it is a geometric object obtained by patching together pieces of the complex plane along *biholomorphic maps*. Before studying Riemann surfaces, it's a good idea to understand the notion of a biholomorphic map, or more generally, just a holomorphic map. In this lecture we'll review the basic definitions and results on holomorphic maps, giving some intuitive explanations but no proofs. For the proofs you can see any basic text on complex analysis.

First, some notations. The set of complex numbers is

$$\mathbb{C} = \{a + b \cdot i \ \mid \ a, b \in \mathbb{R}\}.$$

For a complex number $z \in \mathbb{C}$ and a positive real number r, the *open disk* of radius r centered at z is

$$D(z, r) = \{w \in \mathbb{C} \ \mid \ |w - z| < r\},$$

and the analogous *closed disk* is

$$\overline{D}(z, r) = \{w \in \mathbb{C} \ \mid \ |w - z| \leq r\}.$$

Its *boundary* is the circle

$$\partial\overline{D}(z, r) = \{w \in \mathbb{C} \ \mid \ |w - z| = r\}.$$

Here is the first definition.

Definition 0.1. *A subset $U \subset \mathbb{C}$ is called* open *if for every $z \in U$, there exists a positive real number r such that $D(z, r) \subset U$.*

It is equivalent to ask that U be a union of open disks.

The reason we want to consider open subsets is that we want to differentiate. To differentiate, you need to know not just the values of your function at a point, but its values at all points sufficiently close to your point. An open subset is exactly a subset which, if it contains a point, contains all sufficiently close points as well.

Now we can say what is a holomorphic function.

Definition 0.2. *Let $U \subset \mathbb{C}$ be an open subset, and $f : U \to \mathbb{C}$ a function. We say that f is* holomorphic *if it is complex differentiable at z_0 for every $z_0 \in U$, i.e. for every $z_0 \in U$ the limit (of complex numbers)*

$$\lim_{z \to z_0} \frac{f(z) - f(z_0)}{z - z_0}$$

exists. (The limit value is denoted $f'(z_0)$, as usual.)

There are two ways of thinking about this definition. One is that it is the complex analog of the usual notion of a differentiable function of one real variable. So insofar as the complex numbers, with its algebraic operations and its notion of absolute value, is like the real numbers with its operations and absolute value, the notion of a holomorphic function is like the notion of a real differentiable function.

But this perspective hides the magic of holomorphic functions, which in many ways are very different from real differentiable functions. To bring out the magic, its helpful to think geometrically, viewing the complex numbers as the plane $\mathbb{R}^2$ and thus thinking of f as a function of two real variables.

Then we can ask, what is it that distinguishes a holomorphic function from an arbitrary real differentiable function of two variables? The answer can be phrased as follows. Suppose given a real differentiable function $f : U \to \mathbb{R}^2$ and a point $p \in U$. The derivative of f at p is by definition a certain linear map

$$df|_p : \mathbb{R}^2 \to \mathbb{R}^2,$$

where we think of the source (resp. target) as the vectors emanating from p (resp. $f(p)$). Then an immediate comparison of the definitions shows that f is holomorphic if and only if $df|_p$ is given by multiplication by some complex number (which will be $f'(p)$).

Now, it is only certain linear maps which are of that form. Multiplication by a complex number is either the zero map, or it is the composition of a rotation and a scaling. A particularly vivid consequence is that when f is holomorphic, $df|_p$ is either zero or it preserves oriented angles. Another consequence is a certain rigidity: the whole linear map $df|_p$ is determined by its value on any nonzero vector, since already that fixes the amount of rotation and scaling undergone by the whole map.

To bring this to life, recall the geometric interpretation of $df|_p$: if you leave the point p travelling with instantaneous velocity v, then your image under f will leave $f(p)$ travelling with instantaneous velocity $df|_p(v)$. Thus when f is holomorphic, it will preserve (oriented) angles to first order, since $df|_p$ preserves angles. In fact this is almost an equivalent condition: a real-differentiable function with nonvanishing derivative is holomorphic if and only if it is *conformal*, i.e. preserves angles, or even just if and only if it preserves right angles. That's because the only linear maps which do that are compositions of rotations and scaling, hence given by multiplication by a complex number. Furthermore there is the rigidity: the first-order movement under f in any one direction determines its first-order movement in every direction. E.g. $df|_p(i)$ is the 90-degree counterclockwise rotation of $df|_p(1)$. In coordinates, this is known as the Cauchy-Riemann equation.

Moreover, if you imagine fixing the length of v but sweeping out all possible directions evenly, your image under f will also have fixed length and sweep out all possible directions evenly. Thus a holomorphic map is balanced (the technical term is *harmonic*).

This first-order balance has a large scale manifestation, the *mean value property*:

Theorem 0.3. *Let $f : U \to \mathbb{C}$ be a holomorphic function on an open subset of $\mathbb{C}$, and let $\overline{D}(z, r)$ be a closed disc contained in U. Then the average value of f on the boundary $\partial\overline{D}(z, r)$ is equal to the value of f at the center, i.e.*

$$f(z) = \frac{1}{2\pi r} \int_{\partial\overline{D}(z,r)} f(w) \cdot |dw|.$$

Here we should be imagining w tracing around the circle $\partial\overline{D}(z, r)$, and $|dw|$ as picking up the infinitesimal arc length. The $2\pi r$ corresponds to the length of the circle, i.e. $\int_{\partial\overline{D}(z,r)} |dw|$.

There is a way to make this formula nicer. Namely, if we send w around at uniform speed counterclockwise, then we see the following formula:

$$\frac{dw}{w-z} = \frac{i \cdot |dw|}{r}.$$

The i comes from the fact that dw is perpendicular to $w-z$: the tangent to a circle is perpendicular to the line emanating from the center of the circle. Substituting this in to the integral, we find

$$f(z) = \frac{1}{2\pi i} \int_{\partial \overline{D}(z,r)} \frac{f(w)}{w-z} dw.$$

Here it should be implicit that we're supposed to move around $\partial \overline{D}(z,r)$ in a counterclockwise direction.

This formula is nicer: one reason is that in fact it is valid with z replaced by *any* point inside the disk $D(z;r)$. That's the *Cauchy integral formula*:

Theorem 0.4. *Let $f : U \to \mathbb{C}$ be holomorphic, with U open. If $\overline{D}$ is any closed disk inside U and z is any point in the interior D, then*

$$f(z) = \frac{1}{2\pi i} \int_{\partial D} \frac{f(w)}{w-z} dw,$$

if we go around the boundary once counterclockwise in performing the integration.

This formula is something of a key to the notion of holomorphic function. For example, it explains why holomorphic functions are so regular: the dependence on z in the right-hand side of the equation is easy to control; e.g. we can easily differentiate, or even expand into a power series. Here are some corollaries:

Theorem 0.5. *If f is holomorhpic, so is its derivative f'.*

Theorem 0.6. *If $f : U \to \mathbb{C}$ is holomorphic and $D(z,r)$ is any open disk contained in U, then the power series expansion for f cenetered at z converges in all of $D(z,r)$.*

Let me pause to give a cute example application of this theorem. Consider the power series

$$F(z) = \sum_{n \geq 0} F_n \cdot z^n,$$

where F_n is the n^{th} Fibonacci number. Then the recurrence relation $F_{n+2} = F_{n+1} + F_n$ and normalization $F_0 = F_1 = 1$ translate into the relation

$$F(z) = zF(z) + z^2 F(z) + 1.$$

Using this, we find that $F(z)$ is the power series expansion of

$$F(z) = \frac{1}{1 - z - z^2}$$

around 0. This function F is holomorphic everywhere except when $1 - z - z^2 = 0$, i.e. when $z = \frac{1 \pm \sqrt{5}}{2}$. From this we can conclude that the radius of convergence of our series is $\frac{1-\sqrt{5}}{2}$. Indeed, F is holomorphic in the disc with radius $\frac{\sqrt{5}-1}{2}$, so the above theorem tells us that the radius of convergence is at least $\frac{\sqrt{5}-1}{2}$; but on the other hand, the function F tends to ∞ as z approaches $\frac{1-\sqrt{5}}{2}$, and a power series can't tend to ∞ inside a region where it is defined since a power series is continuous; thus the radius of convergence is also at most $\frac{\sqrt{5}-1}{2}$, whence the claim. This tells you something about the order of growth of the F_n.

Another application of the Cauchy integral formula is the removable singularities theorem:

Theorem 0.7. *Let $U \subset \mathbb{C}$ be open, and $z \in U$. Suppose $f : U \setminus \{z\} \to \mathbb{C}$ is holomorphic, and is bounded in some disk around z. Then f extends uniquely to a holomorphic function on all of U.*

So merely the fact of f being bounded around z implies something much stronger.

Now we have reviewed the basics of holomoprhic functions. But there's one important question we should still discuss: where do holomorphic functions come from? Why do they exist?

The answer is that they come from a variety of sources. We'll talk about two of the more important ones in the next lecture: they come from solutions of polynomial (or more generally holomorphic) equations, and also from solutions of holomorphic differential equations.

But for now we'll give a different kind of answer. First of all, the identity function $z : \mathbb{C} \to \mathbb{C}$ is holomorphic, as is any constant function. Second, the class of holomoprhic functions is closed under addition and multiplication. Thus every polynomial is holomorphic. (The proof of these claims is formally just like the real variable case.) Then, we can access more holomorphic functions by taking limits of polynomials, thanks to the following theorem (also a consequence of the Cauchy formula!):

Theorem 0.8. *Suppose $f_1, f_2, \ldots$ is a sequence of holomorphic functions $U \to \mathbb{C}$ with pointwise limit $f : U \to \mathbb{C}$. Suppose also the following stronger property: for every closed ball $\overline{D} \subset U$, the value*

$$sup_{z \in \overline{D}} |f(z) - f_i(z)|$$

*tends to zero as i tends to ∞ (*uniform convergence on closed balls.*) Then f is holomorphic.*

So other examples of holomorphic functions come as uniform limits of polynomial functions. In fact, essentially all examples are of this form, since a power series is a uniform limit of its Taylor truncations.

Yet, there are many other ways to construct holomorphic functions than via the above path. One is particuarly vivid; it involves going back to the geometric understanding of holomorphic functions, and making it precise will be one of the goals of this course. Before stating a version of it, however, we should define the following important term:

Definition 0.9. *Let U and V be open subsets of $\mathbb{C}$. A* biholomorphism *between U and V is a holomorphic map $f : U \to V$ which is bijective and has holomorphic inverse.*

(The last condition in fact follows from the others, but we'll talk about that later.)

Now we can state Riemann's mapping theorem, which produces many interesting biholomorphisms:

Theorem 0.10. *Let $U \subset \mathbb{C}$ be an open subset which is simply-connected and not equal to $\mathbb{C}$. Then there exists a biholomorphism between U and the unit disk $D = D(0, 1)$.*

We will explain the term "simply-connected" later.

There is something of a physical "proof" of this theorem, which also gives some intuition for the notion of a biholomorphism. The idea is this: we can decorate the unit disk by imagining the radial lines joining the center 0 to the boundary, and also imagining all the circles of radius < 1 centered at 0. We want to find the corresponding decorations of U; then the biholomorphism will be constructed to match these up.

For this, we start by fixing a point $p \in U$ which will correspond to $0 \in D$. Then we place an electric charge at p, and consider the magnetic field it generates. But we do this subject to the boundary condition that the potential should be zero on the boundary of U. Then we can decorate U by the field lines on the one hand, and by the equipotential lines on the other. These will correspond to the radial lines and the circles in D.

The reason this all works, heuristically, is that field lines are always perpendicular to equipotential lines. That is why the resulting map $U \leftrightarrow D$ will be biholomorphic: right angles are preserved.

That was of course a non-rigorous argument, but one can make it rigorous. At the very least, there is a moral to the story: since nature decides what to do by solving real differential equations, we should try to construct holomorphic functions using solutions to (certain) real differential equations as well.

Two finish off this lecture, we'll review the geometry of some of the more basic holomorphic maps, namely the exponential function and the n^{th} power map for n a positive integer.

The exponential function $z \mapsto e^z$ is a holomorphic map $\mathbb{C} \to \mathbb{C}$. To picture it it's easiest to use Cartesian real coordinates on the source $\mathbb{C}$. Namely,

$$e^{a+bi} = e^a \cdot (cos(b) + i \cdot sin(b)).$$

In other words, e^{a+bi} maps to the complex number with polar coordinates (e^a, b) where the angle b is given in radians. What does this look like? Well, let's first imagine the field of vertical lines in the source $\mathbb{C}$. Each of these wraps around a circle centered at the origin, with periodicity of $2\pi i$. If our line is far to the left this circle is of small radius, and if our line is far to the right it is of large radius. On the other hand the horizonal lines simply map to rays emanating from the origin. Note, however, that the origin itself is not in the image of the map. As you move far to the left your values approach the origin, but they never get there.

All in all there is a periodicity in this map: adding a multiple of $2\pi i$ does not change the value under the exponential map. Thus it's convenient to imagine taking the horizontal strip ranging between the lines $b = 0$ and $b = 2\pi$, and making an infinite cylinder out of it by gluing these boundary lines together. Then we could say that the exponential map is giving a biholomorphism between this cylinder and the open subset $\mathbb{C} \setminus \{0\}$. Thus the exponential map "straightens out" $\mathbb{C} \setminus \{0\}$ into a more manifestly symmetric geometric model having the same conformal geometry (angles are preserved by the exponential function). E.g., now the operation of scaling by a non-zero complex number on $\mathbb{C} \setminus \{0\}$ turns into just translating along the cylinder.

What about the n^{th} power map? Since it's a question of multiplication, polar coordinates are the more useful: it sends 0 to 0, and it sends (r, θ) to $(r^n, n\theta)$. The picture is that on the unit circle, the n^{th} power map wraps around n times evenly. That is, if you travel around the unit circle once at constant speed s, then your image under the n^{th} power map will travel around the unit circle n times at constant speed ns. If you move away from the unit circle along a radial line, the argument of your image doesn't change, but your radius will change according to the real-valued n^{th} power map applied to your original radius. We'll have occasion to point out some features of this example later.