18.112 — Functions of a Complex Variable

Class by Andrew Lawrie

Fall 2023

Notes for the MIT class **18.112** (Functions of a Complex Variable), taught by Andrew Lawrie. All errors are my responsibility.

Contents

§1 September 7, 2023

§1.1 Introduction

(Here not everything will make sense; Prof. Lawrie wants to show us a glimpse of what's special about complex analysis.)

Complex analysis connects to basically every area of math (analysis, number theory, geometry, topology); but it's also essential in physics and engineering. So it's an example of a subject that's ubiqituous (like linear algebra).

First, what is complex analysis? At its heart, it's the study of functions $f\colon \mathbb{C} \to \mathbb{C}$ (where $\mathbb{C}$ is the complex numbers). We write complex numbers as $z = x + iy$ for real numbers x and y, where $i^2 = -1$. More generally, we can consider functions $f\colon \Omega \to \mathbb{C}$ where Ω is an open region in $\mathbb{C}$. We'll use z for our complex numbers, so we can write our functions as $z \mapsto f(z)$.

In 18.100 (real analysis) we studied functions of real variables, taking real values; here our domain and target space are both the complex numbers instead.

It turns out that this is somehow a completely different subject than real analysis, which is incredibly surprising. We mainly study functions that are *complex differentiable* (these are called *holomorphic functions*).

Definition 1.1. A function $f\colon \Omega \to \mathbb{C}$ is *differentiable at z* if the limit

$$\lim_{h \to 0} \frac{f(z+h) - f(z)}{h}$$

exists; if so, we call this $f'(z)$.

This is similar to the definition for real numbers; note that here h is a complex number.

It turns out that the theory of real differentiable functions is quite messy — you can have a function that's differentiable but not twice-differentiable. But it turns out that 'complex differentiable' is incredibly restrictive — the class of all complex differentiable functions is a very rigid set. Being a holomorphic function means that several incredible properties hold (and the first part of the class will be about understanding these properties):

1. Let γ be a loop (closed curve) in the complex plane, and f a holomorphic function on a set Ω with $\gamma \subseteq \Omega$. Then $\int_\gamma f(z)\, dz = 0$.
2. Any holomorphic f is infinitely differentiable (in other words, once the first derivative exists, the derivatives of all other orders also exist). (This is called the *regularity* or *analyticity* of holomorphic functions; this is very different from the real case.)
3. Analytic continuation — suppose we have some open $\Omega \subseteq \mathbb{C}$, and $f, g\colon \Omega \to \mathbb{C}$ are both holomorphic on Ω. If $f = g$ on any disc $\mathcal{D} \subseteq \Omega$, then we must actually have $f = g$ on all of Ω. (This requires Ω to be 'connected.')

These mean holomorphic functions are extremely rigid objects. We'll explore implications of all three of these facts throughout the course.

These are very much not true for functions of a real variable — for the third, if we have a nice smooth function, ther are many ways to continue it smoothly past a given point. So knowing two functions are equal on an open subset of $\mathbb{R}$ tells you nothing about what they do elsewhere, which is in stark difference.

We will see the second point through power series — we'll see that every holomorphic function has a power series expansion, which is going to tell us why this property is true.

The plan for the class is we'll first spend some time developing the general theory of holomorphic functions and understanding where these facts come from and their consequences. Later on, we'll study some specific functions. Another cool part of complex analysis is that since high school we've been developing a repertoire of functions, e.g. polynomials, rational functions (quotients of polynomials), sin and cos, exp, log, and so on. But complex analysis is our first chance to get our hands on some other functions that are just as natural as these, but for which we didn't have the background to understand. For example, sin and cos are periodic functions. In the complex plane, you can ask about *doubly* periodic functions. So we'll learn about *elliptic functions*, for example in particular the Weirstrass p-function. There's also the *theta function*, and the *Riemann zeta function* $\zeta(s)$ (one of the most famous unsolved problems in math, the Riemann hypothesis, is about this function), and the *gamma function* $\Gamma(z)$. Once we start studying these functions, complex analysis starts spilling over into other fields — the Riemann zeta function is closely related to the distribution of primes (so connects to number theory), and elliptic functions have deep connections to algebraic geometry. It also connects to physics (Schrodinger's equation) and therefore engineering, as well as probability.

Now we'll come back down to Earth and start from the beginning.

§1.2 Complex numbers

Definition 1.2. A *complex number* can be written as $z = x + iy$, where $x, y \in \mathbb{R}$ and $i^2 = -1$.

This gives us a system of numbers, together with an algebra:

(A) We have $(x_1 + iy_1) + (x_2 + iy_2) = (x_1 + x_2) + i(y_1 + y_2)$.

(M) We have $(x_1 + iy_1)(x_2 + iy_2) = (x_1x_2 - y_1y_2) + i(x_1y_2 + y_1x_2)$, using the distributive property together with the fact that $i^2 = -1$.

Alternatively, we can think of complex numbers geometrically, as points in the plane (with a real and imaginary axis); a complex number is a vector in $\mathbb{R}^2$. We can think of these vectors in polar form; for today, we'll use the notation $z = [r, \theta]$ to denote polar coordinates (where $r \geq 0$). (We'll forget this notation in half an hour, but it makes certain things easier to state.)

We can describe addition and multiplication here as well:

(A) Addition is as vectors — $(x_1, y_1) + (x_2, y_2) = (x_1 + x_2, y_1 + y_2)$.

($\widetilde{\text{M}}$) There's no way to multiply vectors together (and get another vector), but we can multiply complex numbers using polar form — $[r_1, \theta_1] \cdot [r_2, \theta_2] = [r_1r_2, \theta_1 + \theta_2]$.

Exercise 1.3. These two notions of multiplication are the same.

The usual caveats about polar coordinates apply — θ is well-defined only if $z \neq 0$ (meaning $r > 0$), and it's an angle, so it's well-defined only up to addition by 2π, i.e., $[r, \theta] = [r, \theta + 2k\pi]$ for all $k \in \mathbb{Z}$.

With this notion of addition and multiplication, the complex numbesr $\mathbb{C}$ become a *field*, i.e, the field axioms are satisfied.

Theorem 1.4

$\mathbb{C}$ with (A) and (M) (or (A) and ($\widetilde{\text{M}}$)) are a field.

This means the field axioms are satisfied:

1_a: $z_1 + z_2 = z_2 + z_1$ (addition is commutative);

2_a: $(z_1 + z_2) + z_3 = z_1 + (z_2 + z_3)$ (addition is associative);

$z + 0 = z$ (0 is an additive identity);

$z + (-z) = 0$ (all elements have additive inverses).

1_m: $z_1 z_2 = z_2 z_1$.

2_m: $(z_1 z_2)z_3 = z_1(z_2 z_3)$.

3_m: $z \cdot 1 = z$.

4_m: $zz^{-1} = 1$.

5. $z(z_1 + z_2) = zz_1 + zz_2$.

Here $0 = (0, 0)$ and $1 = (1, 0)$ in the geometric interpretation; $-z$ is $(-x, -y) = [r, \theta + \pi]$. We can define

$$z^{-1} = \left(\frac{x}{x^2 + y^2}, \frac{y}{x^2 + y^2}\right) = \left[\frac{1}{r}, -\theta\right].$$

You can start with the algebraic notion where we assume $i^2 = -1$, or we can use the geometric notion and prove it where $i = (0, 1) = [1, \frac{\pi}{2}]$ — we have $i^2 = [1 \cdot 1, \frac{\pi}{2} + \frac{\pi}{2}] = [1, \pi] = (-1, 0) = -1$.

We'll now dispense with this funny notation for polar coordinates.

Definition 1.5. For a complex number $z = x + iy$, we call x the *real part* of z and y the *imaginary part* of z, denoted $\operatorname{Re}(z)$ and $\operatorname{Im}(z)$. We call r the *length* of z and θ its *argument*, written $r = |z|$ and $\theta = \arg(z)$.

These have the algebraic properties that

$$\operatorname{Re}(z_1 + z_2) = \operatorname{Re}(z_1) + \operatorname{Re}(z_2)$$

and similarly $\operatorname{Im}(z_1 + z_2) = \operatorname{Im}(z_1) + \operatorname{Im}(z_2)$; we have $|z_1 z_2| = |z_1| |z_2|$ and $\arg(z_1 z_2) = \arg(z_1) + \arg(z_2)$.

Definition 1.6. The *complex conjugate* of z is the reflection of z over the real axis, denoted $\overline{z}$. If $z = x + iy$ then $\overline{z} = x - iy$.

We'll soon see why this definition is useful. Note first that $|\overline{z}| = |z|$, while $\arg(\overline{z}) = -\arg(z)$.

Theorem 1.7

Let z_1, z_2, z_3 be complex numbers.

1. z is real if and only if $z = \overline{z}$.
2. $\overline{\overline{z}} = z$.
3. $\overline{z_1 + z_2} = \overline{z_1} + \overline{z_2}$. (This is easy to see geometrically — we can add first and thren reflect, or reflect first vie versa.)
4. $\overline{z_1 z_2} = \overline{z_1 z_2}$.
5. $-\overline{z} = \overline{-z}$.
6. $\overline{z^{-1}} = \overline{z}^{-1}$.
7. $|z|^2 = z\overline{z}$.
8. $|z_1 + z_2|^2 = |z_1|^2 + 2\operatorname{Re}(z_1\overline{z_2}) + |z_2|^2$.

Proof of (7). We have $|z|^2 = |x + iy|^2 = x^2 + y^2$, from the algebraic definitions . We can write this as

$$x^2 - iy^2 = (x + iy)(x - iy) = z\overline{z}.$$

□

Proof of (8). First note that $\operatorname{Re}(z) = \frac{z+\overline{z}}{2}$, and $\operatorname{Im}(z) = \frac{z-\overline{z}}{2i}$. So then if we multiply out, we have

$$|z_1 + z_2|^2 = (z_1 + z_2)(\overline{z_1} + \overline{z_2}) = z_1\overline{z_1} + z_1\overline{z_2} + \overline{z_1}z_2 + z_2\overline{z_2} = |z_1|^2 + (z_1\overline{z_2} + \overline{z_1}z_2) + |z_2|^2.$$

Note that $z_1\overline{z_2} = \overline{\overline{z_1}z_2}$, using the previous properties. So the two middle terms are conjugates of each other; then using the fact that $\operatorname{Re}(z) = \frac{z+\overline{z}}{2}$, we get that this is $|z_1|^2 + 2\operatorname{Re}(z_1\overline{z_2}) + |z_2|^2$. □

A cool consequence of (8) is the 'parallelogram law,' where we have two complex numbers z_1 and z_2 and we consider the parallelogram they generate (with vertices z_1, z_2, 0, and $z_1 + z_2$, whose diagonal vectors are $z_1 + z_2$ and $z_1 - z_2$).

Corollary 1.8

We have $|z_1 + z_2|^2 + |z_1 - z_2|^2 = 2|z_1|^2 + 2|z_2|^2$.

(This is because the extra terms we have are $+2\operatorname{Re}(z_1\overline{z_2})$ and $-2\operatorname{Re}(z_1\overline{z_2})$, which cancel.) So the sum of squares of the lengths of the two diagonals of a parallelogram is equal to twice the sum of squares of the lengths of the sides.

§2 September 12, 2023

The first pset will be assigned this Friday and due next Friday.

§2.1 Polynomials

We're going to start developing the theory of complex functions. We'll start with the simplest example, namely polynomials:

Definition 2.1. A *polynomial* $p(z)$ is a $\mathbb{C}$-valued function of a complex variable z of the form

$$p(z) = c_0 z^n + c_1 z^{n-1} + \cdots + c_n,$$

where $c_k \in \mathbb{C}$ and $c_0 \neq 0$. We define n (i.e., the highest exponent with nonzero coefficient) as the *degree* of the polynomial, denoted $\deg(p)$. We cal lthe c_k *coefficients*, and c_0 the *leading coefficient*.

Definition 2.2. A *root* z_0 of a polynomial p is a complex number that satisfies $p(z_0) = 0$.

Example 2.3

Some examples of polynomials are $p(z) = z$, $p(z) = z^2 + 1$, $p(z) = z^2 - 1$. Note that $z^2 + 1$ has no real roots, but it does have two complex roots, namely $\pm i$. The roots of $z^2 - 1$ are ± 1, and the only root of z is 0.

Polynomials will come up a lot; we'll really develop the theory of polynomials and rational functions.

Theorem 2.4

If a polynomial $f(z)$ with $\deg(f) = n$ and leading coefficient c_0 has a root $z_0 \in \mathbb{C}$, then there is some polynomial $f_1(z)$ with $\deg(f_1) = n - 1$ and leading coefficient also c_0 such that for all z we have

$$f(z) = (z - z_0)f_1(z).$$

Proof. We have
$$0 = f(z_0) = c_0 z_0^n + c_1 z_0^{n-1} + \cdots + c_n.$$

Then we can write
$$f(z) = f(z) - f(z_0) = c_0(z^n - z_0^n) + c_1(z^{n-1} - z_0^{n-1}) + \cdots + c_{n-1}(z - z_0).$$

For each of these factors, we can pull out $z - z_0$ — we have the identity that for all $k \in \mathbb{N}$,
$$a^k - b^k = (a-b)(a^{k-1} + a^{k-2}b + \cdots + ab^{k-2} + b^{k-1}),$$

so from each of these factors we can pull out a $z - z_0$ factor (using this identity for $k = n, n-1, \ldots, 1$). Then we get
$$f(z) = (z - z_0)f_1(z),$$

where
$$f_1(z) = c_0(z^{n-1} + z^{n-2}z_0 + \cdots + z_0^{n-1}) + c_1(z^{n-2} + \cdots + z_0^{n-2}) + \cdots.$$

All terms except hte first have degree less than $n-1$, so
$$f_1(z) = c_0 z^{n-1} + \text{terms of degree} < n-1.$$

So f_1 has leading coefficient c_0 and degree $n-1$. □

In other words, this says we can always factor out a root from any polynomial.

Theorem 2.5 (Fundamental Theorem of Algebra, Gauss 1800)
Every polynomial with deg ≥ 1 has at least one root.

Proof. There are many proofs; we will give an analysis one. Let $f(z)$ be a polynomial of degree at least 1, and assume for contradiction that $f(z)$ has no roots; this means $f(z) \neq 0$ for all z.

Now note that $f(0) = c_n$, so $|f(0)| = |c_n| > 0$.

Since f has nonzero degree, now let's look at the function
$$|f(z)| = |c_0 z^n + \cdots + c_n|,$$

and let's think about the asymptotic behavior of $|f|$ for large z. When $|z|$ is really large, this behaves like $|z|^n$ — we have
$$|f(z)| = |c_0 z^n + \cdots + c_n| \asymp |z|^n.$$

(This is because $|z|^n$ is much bigger than any of the other factors.) In particular, if we take $|z|$ large enough, then this polynomial is in absolte value always bigger than $|c_n|$.

So we can find $R > 0$ such that $|f(z)| \geq 2|c_n|$ for all $|z| > R$ — we find a large disk outside of which the polynomial has huge absolute values.

Imagine drawing our disk of radius R centered at 0. The closed disk $\overline{\mathcal{D}(0,R)} = \{z \in \mathbb{C} \mid |z| \leq R\}$ is *compact* (closed and bounded), and $|f|$ is continuous on $\overline{\mathcal{D}(0,R)}$, so $|f|$ has a minimum on $\overline{\mathcal{D}(0,R)}$. ($|f|$ is a real-valued function, the type we studied in real analysis.)

We assumed f has no roots, so that minimum has to be strictly positive. This means:

- The minimum of $|f(z)|$ on this disk, which we denote $|f(z_0)|$, is in fact a global minimum — this is because the minimum is at most $|c_n|$ (since it's at most $f(0)$), and $|f|$ is always at least $2|c_n| > |c_n|$.
- We have $|f(z_0)| > 0$ by our contradiction assumption.

Now we normalize (this is a bit of a trick) — we have $f(z_0) \neq 0$, so we can define

$$g(z) = \frac{f(z_0 + z)}{f(z_0)}.$$

Then $g(0) = 1$, and $|g(z)| \geq 1$ for all z (because z_0 is the global minimum of $|f|$). This is also a polynomial of degree n, since we just divided by a complex number and translated.

Now we have a new polynomial whose absolute value is always at least 1, and is 1 at 0.

Since $\deg(f) = n \geq 1$, there is at least one term with nonzero coefficient; let k be the exponent of the smallest such term, so that

$$g(z) = 1 + cz^k + h(z)$$

where $h(z)$ is a polynomial with no terms of degree at most k. (Here k is the minimal order of a term with nonzero coefficient in g.)

We're now going to find a contradiction; what we'll do is pick a special z that makes this have absolute value strictly less than 1 (which is a contradiction).

Let's choose a special z. Before, we exploited the fact that if we take z with large absolute value, then the leading-order term dominates. But if we now take z's that are super small, then the terms with high degree become much smaller.

We choose z such that:

(1) We choose $\arg(z)$ such that $cz^k = -\left|cz^k\right|$. (This corresponds to choosing the right argument for z such that its kth power times c lies on the negative real axis.)

(2) We choose $|z|$ sufficiently small so that $\left|cz^k\right| < 1$ and $|h(z)| < \left|cz^k\right|$. We can do this because if we take z to be within a sufficiently small ball about 0, then $\left|z^k\right| < \frac{1}{|c|}$; and by taking Z sufficiently small, since all the terms in h have degrees greater than k, we can make the second inequality hold.

Now when we plug in this z, we can get something smaller than 1 in absolute value. For this we need the triangle inequality:

Theorem 2.6

If we have two complex numbers w_1 and w_2, then $|w_1 + w_2| \leq |w_1| + |w_2|$.

So then for this special z, we have

$$|g(z)| \leq \left|1 + cz^k\right| + |h(z)|.$$

Now since we rigged z such that cz^k is a pure real and negative real number, this is equal to

$$\left|1 - \left|cz^k\right|\right| + |h(z)|.$$

$1 - \left|cz^k\right|$ is positive, so this is equal to

$$1 - \left|cz^k\right| + |h(z)|.$$

And because of the last condition, we have $|h(z)| < cz^k$; this means we get a number that's strictly less than 1. This is a contradiction, since we assumed 1 is the global minimum of $|g|$. □

(Later in the class we will be more systematic about introducing open/closed sets, continuous functions, and so on.)

This proves the fundamental theorem of algebra. A corollary of this and the previous result is that we can *completely* factor any polynomial.

Corollary 2.7

Every polynomial of degree at least 1 is a product of linear factors

$$f(z) = c_0(z - z_1)(z - z_2)\cdots(z - z_n).$$

Proof. Use induction and the first two results — factor out the first root to get f_1, a polynomial of degree $n-1$. If $n-1=0$ then you're done; otherwise factor it to get $(z-z_0)(z-z_1)f_2$. Then do the same on f_2, and so on. ☐

Let's look at this formula

$$f(z) = c_0(z - z_1)\cdots(z - z_n). \tag{*}$$

There are exactly n roots of f, namely $z_1, \ldots, z_n$ (where $n = \deg(f)$). There might be repeated roots (for exam.e if $f(z) = (z-i)^2(z+i)^4$). We can relabel these roots — let $z_1, \ldots, z_r$ be the *distinct* roots of f, and let m_k be the number of times z_k appears in the list of roots. Then we can write

$$f(z) = c_0(z - z_1)^{m_1}\cdots(z - z_r)^{m_r}.$$

The LHS has degree n: the degree of the polynomial on the right-hand side is $m_1 + \cdots + m_r$. (The m_k's are called the *multiplicities* of the roots — the multiplicity of a root is the number of times it apppears in the factorization.)

Definition 2.8. For each positive integer n, the *nth roots of unity* are the n roots of the polynomial $z^n - 1$ (i.e., solutions of the equation $z^n = 1$).

By the thery we've just developed, there hvae to be n roots. They admit a very nice description:

Theorem 2.9

For every $n \in \mathbb{N}$, there are exactly n distinct nth roots of unity,

Example 2.10

For $n = 2$, the roots are ± 1.

For $n \geq 3$, they are the vertices of the regular n-gon inscribed in the unit circle with one vertex at $z = 1$.

Proof. We can just write down a formula for the roots: suppose that $n \geq 1$. Then ζ is an nth root of unity if and only if $\zeta^n = 1$. This means $\left|\zeta\right|^n = 1$, so $|\zeta| = 1$ (i.e., the nth roots of unity lie on the unit circle); and $\arg(\zeta^n) = \arg(1) = 0$. Since angles are only defined up to 2π, this means $\arg(\zeta^n) = 2\pi k$ for any k; so either $\arg(\zeta) = 0$ or $\arg(\zeta) = \frac{2\pi k}{n}$ (since $\arg(\zeta^n) = n\arg(\zeta)$ — when you multiply complex numbers, you add their angles).

We can make a list of these, and they'll start repeating — our list of $|\zeta| = 1$ and

$$\arg(\zeta) = 0, \frac{2\pi}{n}, \frac{2\pi\cdot 2}{n}, \ldots, \frac{2\pi(n-1)}{n},$$

After this they start repeating — the angle $\frac{2\pi n}{n}$ is the same as 0. So these are the different roots; and these give rise to the vertices of the regular n-gon. ☐

Using polar to rectangular coordinates, the point z with $|z| = r$ and $\arg(z) = \theta$ is

$$z = r\cos\theta + ir\sin\theta = re^{i\theta}$$

(we haven't defined $e^{i\theta}$, but in a few classes we'll define the complex exponential using power series; or you can define it as the unique solution to the ODE $x'(\theta) = ix(\theta)$ with $x(0) = 1$). Analysts sometimes like defining it using the differential equation, but actually proving that such an equation has solutions involves power series; so it's often cleaner to use the power series definition.

With this notation, a nice observation is that $\zeta = e^{2\pi i/n}$ are the roots of unity, and so you can write the whole set as $1, \zeta, \zeta^2, \ldots, \zeta^{n-1}$.

There's a bunch of fun exercises about nth roots of unity in our homework.

In the proof of the fundamental theorem of algebra, we were talking about continuous functions on closed disks; we'll now systematically develop some language for talking about such things.

Remark 2.11. When you want to study a complex surface (in algebraic geometry or advanced complex analysis), you can try to find the analytic or holomorphic functions on it. Oftentimes polynomials or rational functions are the only ones you get. So polynomials are fundamental objects.

§2.2 Open and Closed Sets

We're going to study complex functions of a complex variable; so we'll study functions $f\colon \mathcal{D} \to \mathcal{C}$ where $\mathcal{D} \subseteq \mathbb{C}$. (The domain $\mathcal{D}$ is the subset of $\mathbb{C}$ on which the function is defined.)

Example 2.12

For polynomials, we can think of them as functions $f\colon \mathbb{C} \to \mathbb{C}$ (you can define a polynomial on all complex numbers).

But for rational functions $f(z) = p(z)/q(z)$ where p and q are polynomials, we can only define $f\colon \mathcal{D} \to \mathbb{C}$ where $\mathcal{D} = \mathbb{C} \setminus \{z_1, \ldots, z_m\}$, where $z_1, \ldots, z_m$ are the roots of q. If we remove the roots of q from the complex plane, then we can define our rational function on the punctured complex plane.

For example, $\frac{1}{z^2+1}$ can be defined on $\mathbb{C} \setminus \{\pm i\}$.

Definition 2.13. A set $\mathcal{D} \subseteq \mathbb{C}$ is *open* if for every $z_0 \in \mathcal{D}$, there exists $r > 0$ such that the open disc $\{z \in \mathbb{C} \mid |z - z_0| < r\}$ is contained in $\mathcal{D}$.

You can draw this as a picture — you need to be able to find a sufficiently small disk centered at z_0 that's entirely contained in $\mathcal{D}$.

Example 2.14

1. $\mathbb{C}$ itself is open (you can take any radius r for any point).
2. The punctured plane $\mathbb{C} \setminus \{z_1, \ldots, z_m\}$ is open.
3. Open disks are open.
4. Closed disks are *not* open. (A closed disk is one including the boundary — if you take any point on the boundary, you cannot find any disk around that point strictly contained inside.)
5. The *slit plane* $\mathbb{C} \setminus (-\infty, 0]$ (where we take the complex plane and remove the negative real axis, including the point 0) is open.

On the first day, we mentioned that we'll often be considering paths, so we need to define that.

Definition 2.15. Let $p, q \in \mathbb{C}$. A *smooth path* γ from p to q is any function $\gamma\colon [a, b] \to \mathbb{C}$ of a real variable t (defined on a closed interval $[a, b]$) such that $\gamma(a) = p$ and $\gamma(b) = q$, and such that γ is continuously differentiable.

(This means γ exists and γ' exists, where the derivative is as we saw in calculus — note that γ is a function of a real variable.)

Example 2.16

1. The line segmemt between p and q, denoted $[p, q]$ and defined as
$$\gamma(t) = tq = (1 - t)p$$
(where $\gamma\colon [0, 1] \to \mathbb{C}$), is a smooth path. (We use $[p, q]$ to denote the image of this path.)
2. The circle $\{z \in \mathbb{C} \mid |z - z_0| = r\}$ is a smooth path — we can take $\gamma(t) = z_0 + r\cos(t) + ri\sin(t)$ for $\gamma\colon [0, 2\pi] \to \mathbb{C}$. This is a path from $z+r$ to itself. (In this example, we have $\gamma'(t) = -r\sin t + ir\cos t$, which is a continuous function of t.)
3. We can take any smooth path, and consider its reverse path — given γ, we define its reverse path as $\widetilde{\gamma}(t) = \gamma(-t)$. If $\gamma\colon [a, b] \to \mathbb{C}$ then $\gamma'\colon [-b, -a] \to \mathbb{C}$.

So smooth paths are not just the picture on the board, but the curve and a parametrization of the curve. When we talk about a smooth path, we're fixing a parametrization; different parametrizations can give the same image. For example, if we put $2t$ in our definition of the circle, we'd have the same image but we'd traverse it twice.

The derivative is defined as in calculus:
$$\gamma'(t) = \lim_{t_1 \to t, t_1 \neq t} \frac{\gamma(t_1) - \gamma(t)}{t_1 - t}$$
provided that this limit exists. Continuous is also defined as in calculus: φ is continuous at t_0 if $\varphi(t) \to \varphi(t_0)$ as $t \to t_0$.

We have smooth paths, but we'll often want to have a generalized version of this called piecewise smooth paths.

Definition 2.17. A *piecewise smooth path* γ from p to q is a finite sequence $\gamma_1, \ldots, \gamma_m$ of smooth paths from p to z_1, z_1 to z_2, ..., z_{m-1} to q. For notation, you can write $\gamma = \gamma_1 + \cdots + \gamma_m$ (this does not mean we're adding anything).

This thing might have corners (corresponding to the points z_i); there the limits for the derivatives might be different from the two sides.

Definition 2.18. A *region* $\mathcal{D}$ is an open *connected* subset of $\mathbb{C}$, where *connected* means that for all $p, q \in \mathcal{D}$, there exists a path γ in $\mathcal{D}$ from p to q.

All our examples of open sets are regions, but if we removed the whole real axis then we wouldn't get a region (because it won't be connected).

Definition 2.19. A subset $\mathcal{D} \subseteq \mathbb{C}$ is *convex* if for all $p, q \in \mathcal{D}$, the line segment $[p, q] \subseteq \mathcal{D}$.

Example 2.20

A circle (or oval) is a convex set; a curvy bean is not convex.

Every convex open set is a region — convex sets are in particular connected, since you can take these line segments as your path. But not every region is convex.

We'll now get to the main definition from which everything follows in this class.

Definition 2.21. Let $\mathcal{D} \subseteq \mathbb{C}$ be a region. A function $f\colon \mathcal{D} \to \mathbb{C}$ is called *holomorphic* (or *analytic*) if $f'(z)$ exists at each $z \in \mathcal{D}$ (i.e., f is complex differentiable at each $z \in \mathcal{D}$).

We'll soon define what $f'(z)$ means. For that, we need to understand what limits mean in the context of complex functions; many of the theorems we saw in real analysis about limits are still true, with the same proofs.

Definition 2.22. We define

$$f'(z) = \lim_{z_1 \to z, z_1 \neq z} \frac{f(z_1) - f(z)}{z_1 - z} = \lim_{h \to 0, h \neq 0} \frac{f(z+h) - f(z)}{h}.$$

We'll derive a slew of consequences the next weeks.

(In an open set if you can find a continuous path you can adjust it to make it a smooth path; so you can take either smooth or not in your definition of connected. For us, we'll only deal with smooth or piecewise smooth paths.)

Remark 2.23. In the plane, path-connectedness and connectedness are the same; so for connected just think about path-connected (there's also an analysis definition about overlapping sets, but we shouldn't worry about that — path-connected is generally weaker, but here it's the same).

§3 September 14, 2023

The first problem set will go out tomorrow and will be due the next Friday; that will be the typical rhythm throughout the semester (except for the midterm week). We will mostly assign the exercises from the notes. We'll assign all the exercises, so you can go through and do them whenever you want.

At the end of last class, we defined:

Definition 3.1. For a region $D \subseteq \mathbb{C}$, a function $f\colon D \to \mathbb{C}$ is *holomorphic* (or *analytic*) at a point $z \in D$ if

$$\lim_{z_1 \to z, z_1 \neq z} \frac{f(z_1) - f(z)}{z_1 - z}$$

exists; if so, the limit is called $f'(z)$.

Definition 3.2. We say f is holomorphic on D if $f'(z)$ exists for every $z \in D$.

The notion of a limit is the same as in 18.100 but specialized to $\mathbb{C}$:

Definition 3.3. Let $\varphi\colon D - \{z_0\} \to \mathbb{C}$ be a complex-valued function (φ doesn't have to be defined at z_0, though it may be), and let $L \in \mathbb{C}$. Then we say $\varphi(z) \to L$ as $z \to z_0$ (with $z \neq z_0$), or equivalently that

$$\lim_{z\to z_0, z\neq z_0} \varphi(z) = L,$$

if for all $\varepsilon > 0$ there exists $\delta > 0$ such that $|\varphi(z) - L| < \varepsilon$ for all $z \in D \setminus \{z_0\}$ with $|z - z_0| < \delta$.

In other words, whenever z is close to z_0 (but not actually equal to z_0), $\varphi(z_0)$ should be close to L. So if you give me an ε-ball around L, I can always find a δ-ball around z_0 such that every point inside the δ-ball gets mapped to the ε-ball.

The usual limit rules from analysis class apply, with the same proofs: If $\varphi(z) \to L$ and $\psi(z) \to M$ as $z \to z_0$, then:

(1) $\varphi(z) + \psi(z) \to L + M$ as $z \to z_0$;

(2) $\varphi(z)\psi(z) \to LM$ as $z \to z_0$;

(3) If $M \neq 0$ and $\psi(z) \neq 0$ on $D \setminus \{z_0\}$, then $\varphi(z)/\psi(z) \to L/M$.

These have the same proofs as in real analysis. To remind us of how (2) works, we can write

$$\varphi(z)\psi(z) - LM = (\varphi(z) - L)M + L(\psi(z) - M) + (\psi(z) - L)(\psi(z) - M).$$

We want to check that the LHS goes to 0; but each of these factors goes to 0, so this is true. (You can use this with the triangle inequality and ε-δ proofs to get an actual proof.)

Because of these limit rules, the derivative behaves nicely in the usual way: if f and g are holomorphic in $D \subseteq \mathbb{C}$, then

(1) $f + g$ is holomorphic in D and $(f + g)' = f' + g'$.

(2) (product rule) fg is holomorphic in D and $(fg)' = f'g + fg'$.

(3) (quotient rule) if $g(z) \neq 0$ in D, then f/g is holomorphic in D and

$$\left(\frac{f}{g}\right)' = \frac{f'g - fg'}{g^2}.$$

(We have the exact same proofs as in an analysis class — since the proofs just rely on the limit rules.)

Corollary 3.4

$f(z) = z^n$ for $n \in \mathbb{N}$ is holomorphic in $\mathbb{C}$, and $f'(z) = nz^{n-1}$.

To prove this, it's enough to prove that $f(z) = z$ is holomorphic (with $f'(z) = 1$); then we use the product rule.

Lemma 3.5

- If f is holomorphic in D, then for all $n \in \mathbb{N}$, f^n is holomorphic in D, and $(f^n)' = nf^{n-1}f'$.
- More generally, we have $(f_1 \cdots f_n)' = f_1' f_2 \cdots f_n + f_1 f_2' f_3 \cdots f_n + \cdots + f_1 \cdots f_{n-1} f_n'$.

(This denotes taking powers, not function composition. The first rule is sort of the chain rule for powers.)

So we have n factors, with one f differentiated in each. (This follows from repeated power rule.)

Using this, we can differentiate the only polynomials we've introduced so far, which are polynomials.

§3.1 Differentiating polynomials

Theorem 3.6

Every polynomial $f(z) = c_0z^n + \cdots + c_n$ is holomorphic in $\mathbb{C}$, and its derivative is the polynomial

$$f'(z) = nc_0z^{n-1} + (n-1)c_1z^{n-2} + \cdots + c_{n-1}.$$

This is a direct consequence of the above rules.

Now let's take this a bit further, and start thinking about the roots.

Theorem 3.7

Suppose that f is a nonzero polynomial, and let $z_0 \in \mathbb{C}$ be a root of f with multiplicity m. If $m = 1$, then $f'(z_0) \neq 0$ (i.e., z_0 is not a root of f'); if $m > 1$, then z_0 is a root of f' of multiplicity $m - 1$.

Proof. We proved last time that if z_0 is a multiplicity-m root, we can factor out all m copies of it to get $f(z) = (z - z_0)^m g(z)$ for some nonzero polynomial g with $g(z_0) \neq 0$.

Now using the product and chain rules, we have

$$f'(z) = m(z - z_0)^{m-1}g(z) + (z - z_0)^m g'(z) = (z - z_0)^{m-1}(mg(z) + (z - z_0)g'(z)),$$

which we can write as $(z - z_0)^{m-1}h(z)$ where $h(z) = mg(z) + (z - z_0)g'(z)$.

If $m = 1$, then the first factor is not present, and $f'(z) = h(z)$; and we know $g(z_0) \neq 0$, while $(z_0 - z_0)g'(z_0) = 0$. So $f'(z_0) = h(z_0) \neq 0$.

If $m > 1$, then z_0 is a root of multiplicity $m - 1$, because we have this factorization and $h(z_0) \neq 0$. □

Corollary 3.8

If f is a polynomial of degree $n \geq 1$, then there are fewer than n different complex numbers w for which the function $g_w(z) = f(z) - w$ has a multiple root.

A root of $f(z) - w$ is a point in the domain of f that is sent to w (i.e., a z such that $f(z) = w$).

Proof. If $w_1, \ldots, w_n$ are distinct, and $f(z_1) = w_1, \ldots, f(z_n) = w_n$, then $z_1, \ldots, z_n$ are distinct (since f is a function). If all of the z_j's are multiple roots of their corresponding equation $f(z) - w_j$ (with multiplicities $m_j > 1$), then they are also roots of $(f(z) - w_j)' = f'(z)$; so they are distinct roots of f'. (The point is that $g_w'(z) = f'(z)$.) This is too many roots, since f' is a nonzero polynomial of degree $n - 1$ (and therefore can only have at most $n - 1$ roots). □

Remark 3.9. There exist functions that are real differentiable but not holomorphic (i.e., differentiable as functions $\mathbb{R}^2 \to \mathbb{R}^2$ but not $\mathbb{C} \to \mathbb{C}$) — for example, $f(z) = \overline{z}$. If we think of this as a function $\mathbb{R}^2 \to \mathbb{R}^2$, then it sends $(x, y) \mapsto (x, -y)$, and is therefore differentiable. But it is not holomorphic; we will see that soon.

This function 'looks' super nice, so it's called *anti-holomorphic*. Later we'll see holomorphicity does something nice with the way things are oriented — holomorphic functions are supposed to preserve angles in some sense.

Being holomorphic is actually extremely rigid.

Just the derivative existing — it's really odd that this simple definition is going to mean so many things. But it'll be a very rigid class that reminds you a lot of polynomials.

Remark 3.10. For the 'hit multiple times' intuition which I didn't write down, consider $f(z) = z^2$, and then $g_w(z) = z^2 - w$. The only case this has a double root is when $w = 0$. For example, the roots of $z^2 - 4$ are ± 2.

In particular, this means 'There are finitely many points where the function is not *locally* injective.'

§3.2 Rational functions

Theorem 3.11

If p and q are polynomials, then the function $f(z) = p(z)/q(z)$ (called a *rational function*) is holomorphic in $\mathbb{C} - \{z_1, \ldots, z_m\}$ where $z_1, \ldots, z_m$ are the roots of $q(z)$.

So if we puncture $\mathbb{C}$ at the roots of q, then our function is holomorphic.

Example 3.12

$(z - z_0)^{-k}$ is holomorphic in $\mathbb{C} \setminus \{z_0\}$, and its derivative is $-k(z - z_0)^{-k-1}$.

This follows from the quotient rule.

§3.3 Convex hull

Definition 3.13. Given complex numbers $z_1, \ldots, z_n$, the *convex hull* $H = H(z_1, \ldots, z_n)$ is the set of all points $z \in \mathbb{C}$ which can be written as

$$z = \lambda_1 z_1 + \cdots + \lambda_n z_n$$

where $\lambda_k \in \mathbb{R}$, $\lambda_k \geq 0$, and $\lambda_1 + \cdots + \lambda_n = 1$.

This generalizes the line segment between two points.

Example 3.14

- If $n = 2$, the convex hull of z_1 and z_2 is $[z_1, z_2]$ (the line segment between z_1 and z_2, which is defined as $\{(1 - \lambda)z_1 + \lambda z_2 \mid \lambda \in [0, 1]\}$).
- If $n = 3$, the convex hull of three points is the triangle formed by those three points. How do we see this? Well, if we have $\lambda_2 + \lambda_3 > 0$ and $z \in H(z_1, z_2, z_3)$, then $z = \lambda_1 z_1 + \lambda_2 z_2 + \lambda_3 z_3 = \lambda_1 z_1 + (\lambda_2 + \lambda_3)(\lambda_2/(\lambda_2 + \lambda_3) z_2 + \lambda_3/(\lambda_2 + \lambda_3) z_3)$, which is $\lambda_1 z_1 + (\lambda_2 + \lambda_3)\tilde{z}$ for some $\tilde{z}$ as defined as above. This $\tilde{z}$ is on the line segment between z_2 and z_3 (since these two coefficients in it sum to 1). Then z is on the line segment between z_1 and $\tilde{z}$. So the convex hull sweeps out this triangle, and you can build up this intuition for larger values of n as well.

We can give ourselves some other ways to think of the convex hull:

Theorem 3.15

Let $H = H(z_1, \ldots, z_n)$ be the convex hull of $z_1, \ldots, z_n$. Then H is the smallest convex set containing $z_1, \ldots, z_n$, meaning that:

(1) H is convex.

(2) H is contained in every convex set K containing $z_1, \ldots, z_n$.

Proof. To prove (1), we have to show that if we take any two points in the convex hull, then the line segment between them is as well. So let $p, q \in H$, so that

$$p = \sum \lambda_j z_j \text{ and } q = \sum \mu_j z_j$$

where $\sum \lambda_j = \sum \mu_j = 1$. Now we want to look at the line segment between p and q, meaning that we want to consider $(1-t)p + tq$ for $t \in [0, 1]$; this is equal to

$$\sum_{j=1}^{n} ((1-t)\lambda_j + t\mu_j) z_j.$$

So it is of the form required, where these are our new λ's; it suffices to check that $\sum((1-t)\lambda_j + t\mu_j) = 1$, which is true because it is $(1-t)\sum \lambda_j + t \sum \mu_j = (1-t) + t = 1$.

Next, we want to prove that if we take any convex set, then H is contained in it. We'll think about the case $n = 3$ — suppose we have any convex set, and we want to show H is inside K, so take z to be a point in the convex hull. Assume that $\lambda_2 + \lambda_3 > 0$ (otherwise our point is just z_1). Then the line segment $[z_2, z_3]$ has to be in K, which means our point $\tilde{z}$ described earlier is in k; then since z_1 is in k, so is $[z_1, \tilde{z}]$, which contains z.

(The proof for general n is similar.) □

§3.4 The Gauss–Lucas Theorem

Theorem 3.16 (Gauss–Lucas)

Let f be a polynomial with $\deg(f) = n \geq 1$. Then each root of f' is contained in the convex hull of the roots of f.

So knowing the roots of f doesn't easily tell you the roots of f', but it does let you trap them.

Proof. Let $z_1, \ldots, z_n$ be the roots of f (not necessarily distinct); then we can write

$$f(z) = c(z - z_1) \cdots (z - z_n),$$

where $c \neq 0$ is the leading coefficient of f. Then applying the product rule, we have

$$f'(z) = c \sum_k (z - z_1) \cdots \widehat{(z - z_k)} \cdots (z - z_n)$$

where this notation means the kth term is removed (since the derivative of $z - z_k$ is just 1).

We can then consider

$$\frac{f'(z)}{f(z)} = \sum_{k=1}^{n} \frac{1}{z - z_k}.$$

(this quotient makes sense for all $z \neq z_1, \ldots, z_n$).

Now let z^* be a root of f'. If z^* is a root of f, then it is definitely in the convex hull. If not, then this quotient makes sense ($f(z^*) \neq 0$) and the numerator is 0, so this sum is 0 — i.e.,

$$\sum_{k=1}^{n} \frac{1}{z^* - z_k} = 0.$$

If a complex number is 0, then so is its conjugate; so we can write

$$0 = \sum_{k=1}^{n} \frac{1}{\overline{z^* - z_k}} = \sum_{k=1}^{n} \frac{z^* - z_k}{|z^* - z_k|^2}.$$

This means we can write

$$z^* \sum_{k=1}^{n} \frac{1}{|z^* - z_k|^2} = \sum \frac{z_k}{|z^* - z_k|^2}.$$

In other words, we have

$$z^* = \sum \frac{1}{\sum \frac{1}{|z^* - z_j|^2}} \cdot \frac{1}{|z - z_k|^2} \cdot z_k.$$

And we have now given an expression of z^* as an element of the convex hull — $z^* = \sum \lambda_k z_k$ where

$$\lambda_k = \frac{\frac{1}{|z^* - z_k|^2}}{\sum_{j=1}^{n} \frac{1}{|z^* - z_j|^2}}$$

(these sum to 1 because the denominator is the sum of all the numerators). So we're done. □

§3.5 Path Integrals

Let $\varphi\colon [a, b] \to \mathbb{C}$ be a complex-valued function of a real variable. Then the integral

$$\int_a^b \varphi = \int_a^b \varphi(t)\, dt$$

is defined in the same way as in calculus (it's a complex-valued function, so you can split it into real and imaginary parts). More precisely, if $\varphi(t) = \alpha(t) + i\beta(t)$, then we define

$$\int_a^b \varphi = \int_a^b \alpha(t)\, dt + i \int_a^b \beta(t)\, dt,$$

where these are the Riemann integrals, defined in the usual way as a limit of Riemann sums —

$$\int_a^b \varphi = \lim \sum_{k=1}^{n} \varphi(t_k^*)(t_k - t_{k-1})$$

where we have a partition $a = t_0 \leq t_1 \leq \cdots \leq t_n = b$, and we choose points in the partition as $t_{k-1} \leq t_k^* < t_k$.

In calculus class, you proved that this limit exists and is unique; we'll assume it.

We have the following facts (which you can prove at the level of Riemann sums):

- $\int_a^b (f + g) = \int_a^b f + \int_a^b g$.
- $\int_a^b cf = c \int_a^b f$.
- $\int_a^b f = \int_a^c f + \int_c^b f$ if $a \leq c \leq b$.
- $\left| \int_a^b f \right| \leq \int_a^b |f|$.
- $\int_a^b 1 = b - a$.

§4 September 19, 2023

§4.1 Path integrals

Last time, for a function $f\colon \mathbb{R} \to \mathbb{C}$, we defined its integral $\int_a^b f(t)\,dt$ as a limit of Riemann sums. We then started listing various properties of the integral:

Proposition 4.1

If f and g are two Riemann integrable functions, then:

(1) $\int_a^b (f+g) = \int_a^b f + \int_a^b g$.

(2) For any $c \in \mathbb{C}$, $\int_a^b cf = c\int_a^b$.

(3) For any $c \in [a,b]$, $\int_a^b f = \int_a^c f + \int_c^b f$.

(4) The triangle inequality — $|\int_a^b f| \leq \int_a^b |f|$.

(5) For f real-valued, if $f \geq 0$ on $[a,b]$, then $\int_a^b f \geq 0$. In particular, if $f \leq g$ on $[a,b]$, then $\int_a^b f \leq \int_a^b g$.

(6) $\int_a^b 1 = b - a$.

These follow directly by viewing the Riemann integral as a limit of Riemann sums (these properties all hold for the corresponding sums, so they still hold when we pass to the limit).

Last time we started talking about smooth paths; we'll now define a *path integral*. First, we need to define the *length* of a path:

Definition 4.2. Let $z\colon [a,b] \to \mathbb{C}$ be a smooth path in $\mathbb{C}$. Its *length* L is defined as

$$L = \int_a^b |z'(t)|\,dt.$$

Example 4.3

(1) The line segment $[p,q]$ has length $|q-p|$.

(2) The circle $|z - z_0| = r$ has length $2\pi r$.

(3) If we have a path z, then its reverse path $\widetilde{z}(t) = z(-t)$ has the same length. (So length is independent of the direction in which we traverse the path.)

(These are left as exercises — you parametrize the paths, and then compute the corresponding integrals.)

We've defined the length of a *smooth* path; if we have a piecewise smooth path (where we concatenate a bunch of smooth paths), we simply add up the lengths of the components.

Definition 4.4. The length L of a piecewise smooth path $z = z_1 + \cdots + z_n$ (where each of the z_j are smooth paths) is the sum of the lengths of $z_1, \ldots, z_n$.

For a while, we'll be trying to compute line integrals of holomorphic functions; so let's now define those.

Definition 4.5. Let $D \subseteq \mathbb{C}$ be a region, and $f\colon D \to \mathbb{C}$ a continuous function. Let $\gamma\colon [a,b] \to D$ be a smooth path in D. Then the *path integral* of f along γ is defined as

$$\int_\gamma f = \int_a^b f(\gamma(t))\gamma'(t)\,dt.$$

The integrand is a continuous function of t (since f is continuous, γ is continuous, and we assumed that γ' is continuous), so this integral is defined.

Remark 4.6. We may use the notation $f \in \mathcal{C}(D;\mathbb{C})$ to mean that f is a continuous function $D \to \mathbb{C}$.

Definition 4.7. If $\gamma = \gamma_1 + \cdots + \gamma_n$ is a piecewise smooth path, then $\int_\gamma f = \sum_j \int_{\gamma_j} f$.

So to integrate along a piecewise smooth path, you integrate along all its smooth components and add them up.

From this definition and our properties of the integral, we get a bunch of properties of the path integral:

Proposition 4.8

Let $f, g \in \mathcal{C}(D;\mathbb{C})$ and $\gamma\colon [a,b] \to D$ a smooth path.

(1) $\int_\gamma (f+g) = \int_\gamma f + \int_\gamma g$.

(2) $\int_\gamma cf = c\int_\gamma f$ for any $c \in \mathbb{C}$.

(3) For any smooth paths α and β such that $\gamma = \alpha + \beta$, we have $\int_\gamma f = \int_\alpha f + \int_\beta f$.

(4) $\left|\int_\gamma f\right| \leq LM$ where L is the length of γ and $M = \max_{z\in\gamma} |f(z)|$.

(For (4), note that a path is a compact subset of $\mathbb{C}$, so f achieves a maximum on it.)

We'll prove (4); the rest are straightforward from the definition of the path integral and the corresponding properties of the integral.

Proof of (4). We have

$$\left|\int_\gamma f\right| = \left|\int_a^b f(\gamma(t))\gamma'(t)\,dt\right| \leq \int_a^b |f(\gamma(t))|\,|\gamma'(t)|\,dt \leq M\int_a^b |\gamma'(t)|\,dt = ML.$$ □

Here's an easy but fundamental example of the subject:

Theorem 4.9

For any circle γ centered at any $z_0 \in \mathbb{C}$, oriented counterclockwise, we have

$$\int_\gamma \frac{dz}{z - z_0} = 2\pi i.$$

(As with normal integrals, we may write either $\int_\gamma f$ or $\int_\gamma f(z)\,dz$.)

Interestingly, the radius of the circle doesn't appear here. This is a hint at some of the rigidity of complex function theory.

Proof. Let $\gamma(\theta) = z_0 + r\cos\theta + ir\sin\theta$ be the parametrization of a circle of radius r centered at z_0, so $\gamma\colon [0, 2\pi] \to \mathbb{C}$. Then

$$\gamma'(\theta) = -r\sin\theta + ir\cos\theta = i(\gamma(\theta) - z_0).$$

So then

$$\int_\gamma \frac{dz}{z - z_0} = \int_0^{2\pi} f(\gamma(\theta))\gamma'(\theta)\,d\theta = \int_0^{2\pi} \frac{1}{\gamma(\theta) - z_0} \cdot i(\gamma(\theta) - z_0)\,d\theta = \int_0^{2\pi} i\,d\theta = 2\pi i. \qquad \square$$

Remark 4.10. What if the circle is in the opposite orientation? Then we get $-2\pi i$.

A few more facts about path integrals:

Theorem 4.11 (FTC for path integrals)
Let $D \subseteq \mathbb{C}$ be a region, and $f \in \mathcal{C}(D; \mathbb{C})$. Suppose that F is an anti-derivative of f on D, i.e., $F' = f$ on D (in particular F is holomorphic on D). Let $\gamma\colon [a, b] \to D$ be a smooth path from p to q. Then

$$\int_\gamma f(z)\,dz = F(q) - F(p).$$

This is just like the fundamental theorem of calculus.

Proof. We will essentially just apply the fundamental theorem of calculus. Since γ is smooth, by the fundamental theorem of calculus for the Riemann integral, we have

$$\int_\gamma f = \int_a^b f(\gamma(t))\gamma'(t)\,dt = \int_a^b F'(\gamma(t))\gamma'(t)\,dt = \int_a^b (F(\gamma(t)))'\,dt$$

by the chain rule. Now by the usual fundamental theorem of calculus, this is equal to $F(\gamma(b)) - F(\gamma(a)) = F(q) - F(p)$. $\square$

Remark 4.12. This generalizes to piecewise smooth paths — you use the same theorem on each smooth component of the piecewise smooth path, and the intermediate endpoints will cancel out. Explicitly, $\int_\gamma f = \sum_j \int_{\gamma_j} f = F(\gamma_1(b_1)) - F(\gamma_1(a_1)) + \cdots = F(q) - F(p)$, since $\gamma_1(b_1) = \gamma_2(a_2)$ and so on.

This gives us a way of computing integrals, if we can find antiderivatives.

Corollary 4.13
We have $\int_\gamma 1\,dz = q - p$.

The proof is to take $f = 1$ in the theorem, and $F(z) = z$.

Corollary 4.14
For each polynomial $f(z) = c_0 z^n + \cdots + c_n$, the FTC applies with

$$F(z) = \frac{c_0}{n+1} z^{n+1} + \cdots + c_n z.$$

(Every polynomial has an antiderivative on the complex plane, as given above.)

Corollary 4.15

If γ is a closed path (i.e., $p = q$) and f is continuous and has an antiderivative on D, then $\int_\gamma f = 0$.

This is because if F is the antiderivative, then we'd get $F(p) - F(p) = 0$.

In particular, if we integrate any polynomial over any closed loop in $\mathbb{C}$, then we get 0 (because all polynomials have a globally defined antiderivative).

However, the function $\frac{1}{z - z_0}$ does not have an antiderivative (or else the integral would be 0). We can see that there's a problem caused by z_0; even though the integral doesn't know about this point (the path doesn't contain this point), it feels it somehow. We'll unpack this later.

Remark 4.16. What if we chose D to contain the circle but exclude the point z_0, in that example? The point is that we can't find an antiderivative (even on the punctured disk). This tells us that we can't define log just by removing a point — to define log you need to remove a whole line segment. (log is a super fascinating complex function, and the fact that we have to remove this slit means it's a wild thing.)

One of the fundamental theorems of the subject is:

Theorem 4.17 (Cauchy's integral theorem, 1820)

Let D be a convex region and $f\colon D \to \mathbb{C}$ be holomorphic in D. Let γ be a closed smooth pathin D. Then

$$\int_\gamma f(z)\, dz = 0.$$

(These ideas had been circulating for a while; Cauchy put everything on firm analysis foundations in the early 1800s.)

(For now we need convex; this holds in greater generality, but that involves complicated things we'll see later.)

We just saw that this holds for polynomials, and fails for the creature $1/(z - z_0)$. This theorem says that it holds for *any* holomorphic function (on a convex region).

We'll prove this using a slightly more modern proof given by Goursat around 1900. It's based on two lemmas. The core one is the following (which proves Cauchy's theorem restricting to the case of triangles):

Lemma 4.18 (Goursat's Lemma, 1900)

Let $D \subseteq \mathbb{C}$ be a convex region, $f \in \mathcal{H}(D; \mathbb{C})$, and Δ a triangle in D (oriented counterclockwise). Then

$$\int_\Delta f(z)\, dz = 0.$$

Remark 4.19. We use $\mathcal{H}(D; \mathbb{C})$ to denote the class of holomorphic functions on D.

The theorem will follow from that lemma together with another (which is where we will use convexity):

Lemma 4.20 (Antiderivative Lemma)

Let D be a convex region, $f \in \mathcal{C}(D; \mathbb{C})$, and assume that $\int_\Delta f = 0$ for all triangles Δ in D. Then f has a holomorphic antiderivative in D.

Note that here we only require f to be a continuous function.

These two lemmas combined, together with the fundamental theorem of calculus, give us Cauchy's theorem — take a holomorphic function f. Then by the first lemma, the path integral of f along any triangle is 0; by the second lemma, this means that f has an antiderivative. So since f has an antiderivative, its integral around any loop is 0.

> **Remark 4.21.** The Goursat lemma came later, when modern analysis had been better developed; it's incredibly slick.

> **Remark 4.22.** We can instead allow piecewise smooth paths; and we can also sort of relax convexity.

Proof of antiderivative lemma. Choose any $z_0 \in D$. Then for every $z \in D$, set $F(z) = \int_{[z_0,z]} f$. We will show that F is holomorphic and an antiderivative of f. (This is a pretty natural candidate antiderivative if we think about FTC.) Note that this is where we're using convexity — by convexity, the line segment $[z_0, z]$ is contained in D, so F is well-defined.

To do this, we'll use our assumption that the integral of f on a triangle is 0. Take any $z_1 \in D$, and consider the triangle $\Delta = [z_0, z_1] + [z_1, z] + [z, z_0]$. Then we have

$$0 = \int_\Delta f = \int_{[z_0,z_1]} f + \int_{[z_1,z]} f + \int_{[z,z_0]} f.$$

The first term is $F(z_1)$ and the last is $-F(z)$ (since we have the reverse path), so this means we have

$$F(z_1) - F(z) = \int_{[z,z_1]} f(\zeta)\, d\zeta = \int_{[z,z_1]} f(z)\, d\zeta + \int_{[z,z_1]} (f(\zeta) - f(z))\, dz.$$

In other words, this is equal to

$$F(z_1) - F(z) = f(z)(z_1 - z) + \int_{[z,z_1]} (f(\zeta) - f(z)) d\zeta.$$

We're trying to show that the derivative of F is equal to f, so we want to consider the difference quotient

$$\frac{F(z_1) - F(z)}{z_1 - z}.$$

Call the last term $\varepsilon(z, z_1) = \int_{[z,z_1]} (f(\zeta) - f(z))d\zeta$; we need to show that it goes to 0 as $z_1 \to z$, faster than $z_1 - z$. To do this, note that

$$|\varepsilon(z, z_1)| \le |z_1 - z| \cdot \max_{\zeta \in [z,z_1]} |f(\zeta) - f(z)|\,.$$

Now rearranging the top line, we have

$$\left| \frac{F(z_1) - F(z)}{z_1 - z} - f(z) \right| \le \max_{\zeta \in [z,z_1]} |f(\zeta) - f(z)|\,.$$

Letting $z_1 \to z$, since f is continuous the right-hand side goes to 0, which means $F' = f$. □

> **Remark 4.23.** You can generalize 'convex' but it involves topological arguments.

We'll do the Goursat lemma next time.

§5 September 21, 2023

§5.1 Cauchy's integral theorem

We're starting to get to exciting material; last class we saw the following major result.

Theorem 5.1 (Cauchy's integral theorem)
Let $D \subseteq \mathbb{C}$ be a convex region, $f: D \to \mathbb{C}$ a holomorphic function, and γ a closed path in D. Then

$$\int_\gamma f(z)\,dz = 0.$$

(We assume all paths are piecewise smooth.)

Last time, we wrote down two lemmas:

Lemma 5.2 (Goursat lemma)
Let $D \subseteq \mathbb{C}$ be a convex region, $f: D \to \mathbb{C}$ a holomorphic function, and Δ a triangle in D. Then $\int_\Delta f(z)\,dz = 0$.

This is a special case of Cauchy's theorem, when the closed path γ is a triangle.

Lemma 5.3 (Antiderivative lemma)
Let $D \subseteq \mathbb{C}$ be a convex region and $f: D \to \mathbb{C}$ a *continuous* function. Assume that $\int_\Delta f(z)\,dz = 0$ on every triangle. Then f has a holomorphic antiderivative.

The third ingredient in the proof is the fundamentla theorem of calculus for line integrals:

Theorem 5.4 (FTC)
If $D \subseteq \mathbb{C}$ is any region and f is continuous on D and has a holomorphic antiderivative F (i.e., $F' = f$), then $\int_\gamma f = 0$ on every closed loop γ.

(This is a special case of FTC for loops — in general FTC says that for a path γ from p to q, $\int_\gamma f = F(q) - F(p)$.)

Remark 5.5. We're seeing Goursat's proof, which was from much later than Cauchy's. When Cauchy proved the theorem, complex function theory wasn't quite rigorous because we didn't fully know what a function was.

We'll first prove the Goursat lemma, and then the Cauchy theorem.

Proof of Goursat lemma. Consider some triangle. (All our paths are oriented counterclockwise.)

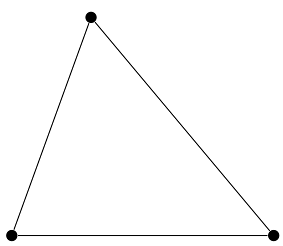

We want to compute $\int_\Delta f$, so we take this triangle and chop it into four smaller triangles.

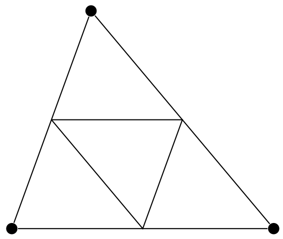

Then we can write

$$\int_\Delta f = \sum \int_{\Delta^{(j)}} f,$$

since every path that we added in gets traversed in both directions (and therefore cancels out).

We want to show that this integral equals 0; one way to do this is to estimate its size. First we'll try to estimate $\left|\int_\Delta f\right|$; note that the integral of f around one of these smaller triangles (in absolute value) is at least $\frac{1}{4}$ of the integral of f around the larger triangle. Call that triangle $\Delta_1 \in \{\Delta^{(1)}, \ldots, \Delta^{(4)}\}$, so that

$$\left|\int_\Delta f\right| \leq 4\left|\int_{\Delta_1} f\right|.$$

Let's suppose that this is the triangle on the bottom-left. Then we chop up this triangle into four pieces in the same way, and select Δ_2 to be one of those four pieces in the same way:

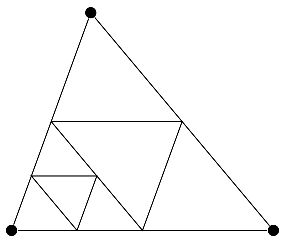

Then we have

$$\left|\int_\Delta f\right| \leq 4\left|\int_{\Delta_1} f\right| \leq 4^2\left|\int_{\Delta_2} f\right| \leq \cdots.$$

Proceeding in this fashion, we get a sequence of triangles Δ_1, Δ_2, ..., such that

$$\left|\int_\Delta f\right| \leq 4^n\left|\int_{\Delta_n} f\right|.$$

Now let's think about the lengths of these triangles — if $L(\Delta) = L$, then $L_n = L(\Delta_n) = \frac{1}{2^n}L$ (since every time we split the triangles up, each of the smaller triangles has half the perimeter of the original).

Now note that each of the triangles we pick is inside the previous one, so $\Delta \supseteq \Delta_1 \supseteq \Delta_2 \subseteq \cdots$ (thinking about the area inside them, including the boundary). Each of these triangles (with their boundary) is compact, so we have a nested sequence of compact sets; this means their intersection is nonempty, so there exists some z^* that is inside Δ_n for each n (i.e., there's a point in every single one of these triangles).

We'll now use the fact that f is holomorphic — f is holomorphic in D, so $f'(z^*)$ exists. This means

$$\lim_{z \to z^*} \frac{f(z) - f(z^*)}{z - z^*} = f'(z^*)$$

exists. Unpacking this, for z near z^* (but not equal to z^*), we can write

$$f(z) = f(z^*) + f'(z^*)(z - z^*) + r(z)(z - z^*),$$

where $r(z) \to 0$ as $z \to z^*$. (This is essentially the first-order Taylor expansion, with $r(z)$ as the remainder.) Note that we can simply define $r(z^*) = 0$, and then this expression holds for all z, including z^*.

But now we have

$$\int_{\Delta_n} f(z)\,dz = \int (f(z^*) + f'(z^*)(z - z^*))\,dz + \int_{\Delta_n} r(z)(z - z^*)\,dz$$

(the nth triangle is pretty small, so the points z on this triangle will be close to z^*). The first function $f(z^*) + f'(z^*)(z - z^*)$ is just a polynomial, so it has an antiderivative — and therefore its integral is 0 (as we've seen before).

Now we'll need to estimate the second integral. Note that we have a 4^n and a $\frac{1}{2^n}$, so when we estimate this term we have to gain another $\frac{1}{2^n}$ from it. To do so, we have

$$\left|\int_{\Delta_n} f(z)\,dz\right| \le \left|\int_{\Delta_n} r(z)(z - z^*)\,dz\right| \le \frac{1}{2^n} L \max_{z\in\Delta_n} |r(z)|\,|z - z^*|$$

(using the fact that our length is $L/2^n$). But $|z - z^*|$ is at most the perimeter of the triangle as well, so this is at most

$$\frac{L^2}{4^n} \max_{z\in\Delta_n} |r(z)|\,.$$

Plugging back into our estimate, we have

$$\left|\int_{\Delta} f\right| \le L^2 \max_{z\in\Delta_n} |r(z)|$$

for all n. Now taking $n \to \infty$, the right-hand side goes to 0 (as z is on the boundary of a smaller and smaller triangle containing z^*); so since the left-hand side doesn't depend on n, it has to equal 0. □

Remark 5.6. You can see that this proof took some techniques from real analysis — Cauchy didn't have these (he was developing the subject), but we do.

Remark 5.7. We really do need f to be holomorphic, not just differentiable as a function $\mathbb{R}^2 \to \mathbb{R}^2$ — we used this to write $f(z) = f(z^*) + f'(z^*)(z - z^*) + \cdots$. For functions $\mathbb{R}^2 \to \mathbb{R}^2$, there's Green's theorem that we learned in calculus (which is the proof that Cauchy gave) — you can evaluate this type of integral using Green's theorem, and you need the right things to cancel for the answer to be 0; those end up being essentially the Cauchy–Riemann equations.

So a lot of this is special to complex function theory, but there do exist generalizations to *harmonic functions* (which are functions $\mathbb{R}^2 \to \mathbb{R}$ that satisfy certain properties — the real and imaginary parts of holomorphic functions turn out to be harmonic functions, and several things like this hold for harmonic functions).

But for general differentiable functions $\mathbb{R}^2 \to \mathbb{R}^2$, this theorem is not true.

§5.2 Cauchy integral formula

Now we get to our next big theorem, also from Cauchy (who was trying to compute integrals, and realized this).

Theorem 5.8

Let $D \subseteq \mathbb{C}$ be a convex region and $f \colon D \to \mathbb{C}$ holomorphic. Then for all $z_0 \in D$ and every circle γ containing z_0, we have

$$f(z_0) = \frac{1}{2\pi i} \int_\gamma \frac{f(z)\,dz}{z - z_0}. \tag{1}$$

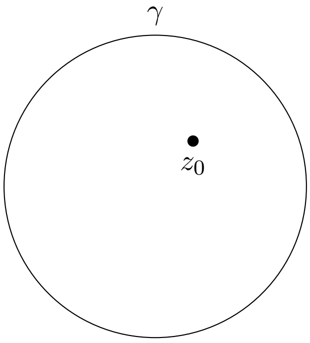

This means if we have a circle γ and a point z_0 inside it, we can figure out the value of $f(z_0)$ just by looking at the values of f on the boundary — so we can recover the entire function inside the disk just by looking at its values on the boundary. This is really remarkable.

Remark 5.9. In the case of real-valued functions, there's an analogous formula for harmonic functions — and more generally, this holds for functions solving certain PDEs. But the proof in complex analysis is more beautiful.

(Circles are always oriented counterclockwise.)

Proof. We'll first reduce the problem of computing the integral around γ to computing the integral around a small circle γ_r (of radius r) centered at z_0.

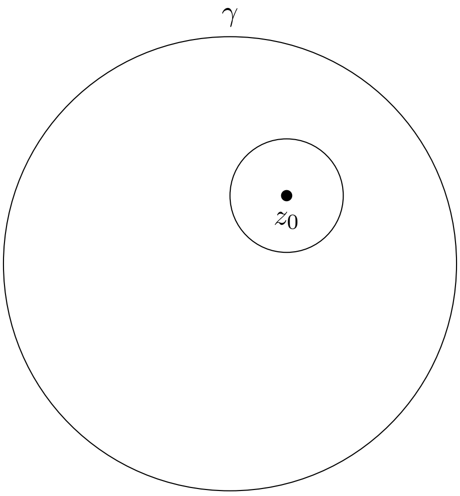

To do so, we create four small pieces — each of γ_1, γ_2, γ_3, and γ_4 is a loop (consisting of the two straight segments, and the pieces of γ_r and γ).

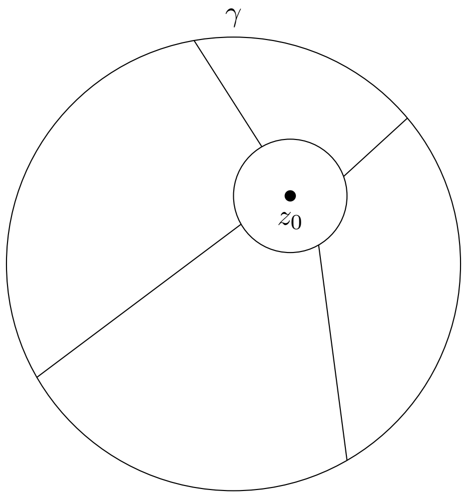

When we sum up the integrals along these paths, the straight segments cancel out — so we have

$$\int_\gamma = \int_{\gamma_r} + \int_{\gamma_1} + \int_{\gamma_2} + \int_{\gamma_3} + \int_{\gamma_4}.$$

But now let's look at the function we're trying to integrate. None of the regions bounded by γ_i have z_0 inside them, and f and $\frac{1}{z-z_0}$ are holomorphic in the complex plane exluding z_0, so in particular they're holomorphic inside these regions.

We want to claim that these integrals are all 0, using the Cauchy integral theorem. But we have a slight problem that these regions are not convex, but we can just find a nice convex region that contains γ_3 but not z_0 (by taking a half-plane, for example — we can assume that each of our arcs has length less than 180°, so that we can find a line that is strictly above the arc and cuts between z_0 and this arc, and we can declare D' to be the intersection of D with this half-plane and use it to apply the Cauchy theorem — it's straightforward to show that the intersection of two convex sets is convex, since for any two points the line segment between them is in both sets).

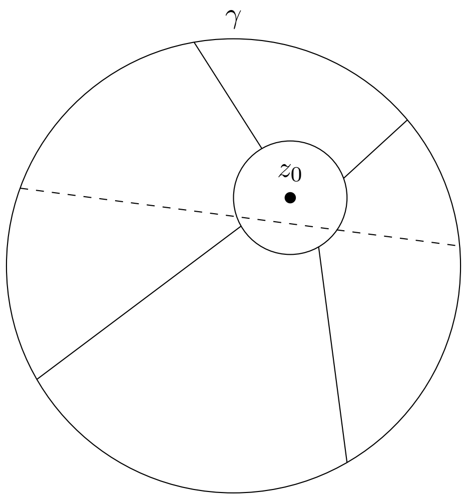

Remark 5.10. It's possible to generalize Cauchy's theorem to more general domains, but even for simply connected ones you can have issues defining the interior and exterior of a curve. So it's simpler to just work with convex ones — in fact Cauchy only worked with rectangles.

So then each of these integrals is 0 by the Cauchy theorem — we have

$$\int_{\gamma_j} \frac{f(z)}{z - z_0}\, dz = 0$$

for each j. So this gives us that

$$\int_{\gamma} \frac{f(z)}{z - z_0}\, dz = \int_{\gamma_r} \frac{f(z)}{z - z_0}\, dz$$

for all $r > 0$ (small enough that the small circle is contained in the big disk).

So now we've reduced the problem to the following claim (we can also take the limit as $r \to 0$, since we can choose r):

> **Claim 5.11** — We have $\int_{\gamma_r} \frac{f(z)}{z-z_0}\, dz \to 2\pi i f(z_0)$ as $r \to 0$.

To prove this, recall that for all r we have

$$\int_{\gamma_r} \frac{dz}{z - z_0} = 2\pi i$$

(we saw this last time). So the thing we want to prove is equivalent to showing that

$$\int_{\gamma_r} \frac{f(z) - f(z_0)}{z - z_0}\, dz \to 0$$

as $r \to 0$ (by subtracting the two sides). But the integrand converges to $f'(z_0)$, and is therefore bounded as $z \to z_0$; this means

$$\left| \int_{\gamma_r} \frac{f(z) - f(z_0)}{z - z_0}\, dz \right| \leq L(\gamma_r) \cdot 2\left|f'(z_0)\right| = 4\pi r \left|f'(z_0)\right|$$

for small r. (Here the 2 in $2\left|f'(z_0)\right|$ is because the integrand converges to $f'(z_0)$, and so its absolute value is bounded by e.g. $2\left|f'(z_0)\right|$ for z close to z_0, which occurs if we take r small enough.) This goes to 0 as $r \to 0$, and we're done. □

In fact, not only are the values that f takes in the interior of a disk determined by the values it takes at the boundary, but all its derivatives are *also* bounded.

> **Theorem 5.12** (Cauchy derivative theorem)
> Let D be a convex region and $f\colon D \to \mathbb{C}$ holomorphic. Then f', f'', f''', … all exist and are holomorphic in D, and for all $z_0 \in D$ and circles γ containing z_0, we have
>
> $$f^{(n)}(z_0) = \frac{n!}{2\pi i} \int_{\gamma} \frac{f(z)\, dz}{(z - z_0)^{n+1}}.$$

So in particular, once f has *one* complex derivative it has all of them. And not only do these derivatives exist, but we can compute them using only the values that f takes on the boundary of some circle.

This is an astounding theorem.

Proof. We'll prove this by induction, with (1) as our base case. Now suppose that $n \geq 1$, and that we have proven the statement for $n - 1$ — so $f^{(n-1)}$ is holomorphic in D, and for all $z_0 \in D$ and circles γ about z_0, we have

$$f^{(n-1)}(z_0) = \frac{(n-1)!}{2\pi i} \int_{\gamma} \frac{f(z)}{(z - z_0)^n}\, dz.$$

This is our induction hypothesis; now we have to prove the corresponding statement for n.

We want to show that the nth derivative exists, so to do that, we should look at the difference quotient

$$\frac{f^{(n-1)}(z_1) - f^{(n-1)}(z_0)}{z_1 - z_0}.$$

We want to take $z_1 \to z_0$, so we can think of z_1 as being very close to z_0 — in particular, we can think of it as being inside γ.

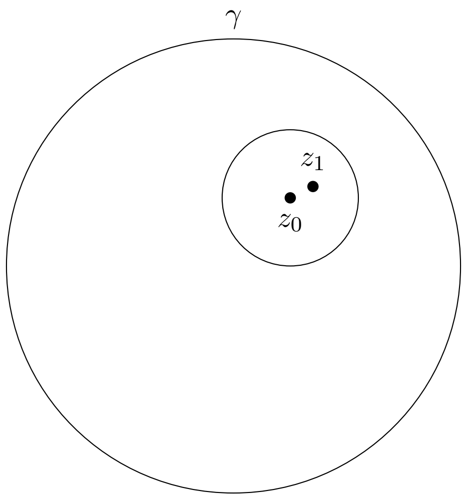

Now we have

$$\frac{f^{(n-1)}(z_1) - f^{(n-1)}(z_0)}{z_1 - z_0} = \frac{(n-1)!}{2\pi i} \int_\gamma \frac{\frac{1}{(z-z_1)^n} - \frac{1}{(z-z_0)^n}}{z_1 - z_0} f(z)\, dz.$$

Note that for each fixed z, the large fraction goes to $\frac{n}{(z-z_0)^{n+1}}$ as $z_1 \to z_0$ — imagine that we fix z, and think of this as a function of z_0. In other words, we think of $\varphi_z(\zeta) = \frac{1}{(z-\zeta)^n}$, so that $\varphi_z'(\zeta) = \frac{n}{(z-\zeta)^{n+1}}$.

Unfortunately, we have an integral and a limit, and we can't interchange them without thinking more carefully about how the function approaches its limit (in order to pass the limit through the integral — we're trying to take a limit as $z_1 \to z_0$, but we need to put that through the integral, and to do that we need to prove some sort of uniformity in z). (As we've seen in real analysis, there are examples where we can't exchange the order of integration; here we need to show that we can.)

In order to do this, we need an estimate.

Lemma 5.13

If φ, φ', and φ'' are holomorphic in a region E (not necessarily convex), and $[p, q]$ a line segment inside E, then

$$\left|\frac{\varphi(q) - \varphi(p)}{q - p} - \varphi'(p)\right| \le |q - p| \max_{[p,q]} |\varphi''|.$$

(This looks like what you'd get from the second-order Taylor estimate, so it's not super surprising. But we need to prove it in the setting of complex functions.)

Proof. We still use the fundamental theorem of calculus; this tells us that

$$\varphi(q) - \varphi(p) = \int_{[p,q]} \varphi'(\xi) d\xi.$$

But then we have

$$\varphi'(\xi) - \varphi'(p) = \int_{[p,\xi]} \varphi''(\xi_1)\, d\xi_1.$$

Plugging in this formula for $\varphi'(\xi)$ into our original integral, we get

$$\varphi(q) - \varphi(p) = \int_{[p,q]} \varphi'(p)\, d\xi + \int_{[p,q]} \int_{[p,\xi]} \varphi''(\xi_1)\, d\xi_1\, d\xi.$$

But the first term has nothing to do with ξ, so it's just $\varphi'(p)(q-p)$, and so

$$\varphi(q) - \varphi(p) = (q-p)\varphi'(p) + \int_{[p,q]} \int_{[p,\xi]} \varphi''(\xi_1)\, d\xi_1\, d\xi,$$

and we have

$$\left| \frac{\varphi(q) - \varphi(p)}{q-p} - \varphi'(p) \right| \le \frac{1}{q-p} \int_{[p,q]} \int_{[p,\xi]} \max_{[p,q]} |\varphi''(\xi_1)| = \frac{1}{q-p}(q-p)^2 \max |\varphi''| . \qquad \square$$

We'll finish the proof next time; the point is that z_0 and z_1 are both inside γ_r, so we can bound them away from the boundary (we need to do this because we have division by $z_i - z$). □

§6 September 26, 2023

Last time, we stated but didn't finish proving the Cauchy derivative formula:

Theorem 6.1 (Cauchy derivative formula)
Let $f\colon D \to \mathbb{C}$ be holomorphic in a convex region $D \subseteq \mathbb{C}$. Then $f^{(n)}$ is holomorphic in D for all $n \in \mathbb{N}$, and for all $z_0 \in D$ and all positively oriented circles $\gamma \subseteq D$ containing z_0, we have

$$f^{(n)}(z_0) = \frac{n!}{2\pi i} \int_\gamma \frac{f(z)\, dz}{(z-z_0)^{n+1}}.$$

So the value that any derivative of f takes at z_0 is determined by the values f takes on the boundary of the circle.

Last time, we proved the following lemma about estimates for the difference quotient of a function based on its derivative.

Lemma 6.2
Let E be a region, and suppose that φ, φ', and φ'' are holomorphic in E and $[p,q] \in E$. Then

$$\left| \frac{\varphi(q) - \varphi(p)}{q-p} - \varphi'(p) \right| \le |p-q| \max_{[p,q]} |\varphi''| .$$

(We proved this last time; it follows from the definition of the derivative and the fundamental theorem of calculus.)

Last time, we started the proof of the Cauchy derivative formula; now we'll finish it.

Proof. We use induction on n. The base case $n = 0$ is the Cauchy integral formula (which we proved previously). For the induction step, assume that the Cauchy derivative formula is true for $n-1$; we want to prove it for n.

To consider the nth derivative at z_0, we consider the difference quotient

$$\frac{f^{(n-1)}(z_1) - f^{(n-1)}(z_0)}{z_1 - z_0}.$$

We take the limit as $z_1 \to z_0$, so we can assume that z_1 also lies inside γ.

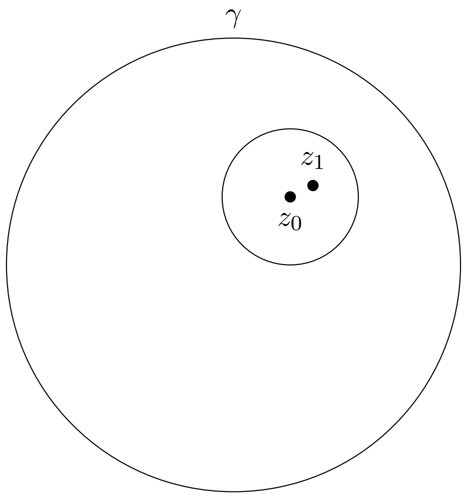

Then using the inductive hypothesis, we can rewrite this as

$$\frac{f^{(n-1)}(z_1) - f^{(n-1)}(z_0)}{z_1 - z_0} = \frac{(n-1)!}{2\pi i} \int_\gamma \left(\frac{\frac{1}{(z-z_1)^n} - \frac{1}{(z-z_0)^n}}{z_1 - z_0} \right) f(z)\, dz.$$

Last time, we noticed that for each *fixed* z, as we let $z_1 \to z_0$, the quantity in parentheses approaches the derivative of $\frac{1}{(z-\zeta)^n}$ at z_0, which is $\frac{n}{(z-z_0)^{n+1}}$.

We're now going to apply the lemma to this function — define

$$\varphi_z(\zeta) = \frac{1}{(z-\zeta)^n},$$

so that

$$\varphi_z'(\zeta) = \frac{n}{(z-\zeta)^{n+1}} \text{ and } \varphi_z''(\zeta) = \frac{n(n+1)}{(z-\zeta)^{n+2}}.$$

We want to apply this lemma, so we need to choose a region. For each fixed z this function has a singularity at z, which is on γ, so we want to exclude γ; so we define our region E as the open disk bounded by γ (so that all of these functions φ_z — for each $z \in \gamma$ — are holomorphic in ζ).

Now we can use this lemma to compare the term we have inside our integral to $\varphi'(z_0)$ — by the lemma, we have

$$\left| \frac{\frac{1}{(z-z_1)^n} - \frac{1}{(z-z_0)}^n}{z_1 - z_0} - \varphi_z'(z_0) \right| \le |z - z_0| \max_{\zeta \in [z_0, z_1]} \frac{n(n+1)}{|z-\zeta|^{n+2}},$$

where $\varphi_z'(z_0) = \frac{n}{(z-z_0)^{n+1}}$. We want to show that the right-hand side is small. To do this, we have z_0 inside an open disk bounded by γ, and z_1 close to z_0. We can just choose a small disk centered around z_0, where if z_0 is 2δ from the boundary of the disk, then we take the disk centered at z_0 with radius δ and only consider z_1 in here.

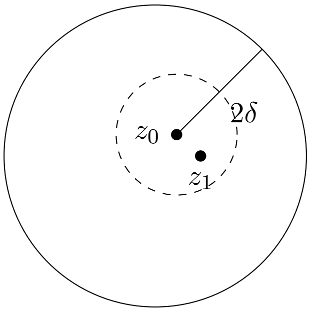

Then we're considering ζ on the interval $[z_0, z_1]$ and z on the boundary of γ; the closest they can be is δ. Then the right-hand side is at most $|z_1 - z_0| \cdot n(n+1)/\delta^{n+2}$.

So then we can write our original integral as

$$\left| \frac{f^{(n-1)}(z_1) - f^{(n-1)}(z_0)}{z_1 - z_0} - \frac{n!}{2\pi i} \int_\gamma \frac{1}{(z-z_0)^{n+1}} \, dz \right| \le \frac{1}{2\pi} \cdot 2\pi r \, |z_1 - z_0| \cdot \frac{(n+1)!}{\delta^{n+2}} \max_\gamma |f| \, ,$$

which goes to 0 as $z_1 \to z_0$. (Here r is the radius of γ, so that $2\pi r = L(\gamma)$.)

To elaborate, the expression we want is

$$\frac{(n-1)!}{2\pi i} \int_\gamma \frac{n}{(z-z_0)^{n+1}} f(z) \, dz + \frac{(n-1)!}{2\pi i} \int_\gamma \left(\cdots - \frac{n}{(z-z_0)^{n+1}} \right) f(z) \, dz,$$

where $\cdots$ is the thing we were originally trying to integrate. The first term is the thing we have on the left (the second term); then we use the bound from our lemma to bound the absolute value of the quantity *inside* the integral (for all $z \in \gamma$), and then bound by the length of the integral to finish. □

Remark 6.3. Do we really need them to be circles? Well, you can generalize this by taking any other nice-looking object and using the Cauchy theorem with a circle inside it(that the integral of a holomorphic function on a closed loop is 0). But it's easier to write a clean statement with circles.

§6.1 The Cauchy–Riemann equations

Suppose that f is holomorphic in some region D, and take some $z_0 \in D$. Then

$$f'(z_0) = \lim_{h \to 0} \frac{f(z_0 + h) - f(z_0)}{h}$$

by the definition of the complex derivative. In other words,

$$f(z_0 + h) = f(z_0) + f'(z_0)h + h \cdot \varepsilon(h),$$

where $|\varepsilon(h)| \to 0$ as $h \to 0$. This definition says essentially that $f(z_0) + f'(z_0)h$ is the approximation of f as a linear function.

Let's compare this with what we get when we think of f as a function $\mathbb{R}^2 \to \mathbb{R}^2$, forgetting the complex structure. If we identify $z = x + iy$ with the point $(x, y) \in \mathbb{R}^2$, and we define $F \colon D \to \mathbb{R}^2$ where $D \subseteq \mathbb{R}^2$ and $z_0 \in D$, then F is differentiable at p_0 if there exists a linear transformation $J(z_0) \colon \mathbb{R}^2 \to \mathbb{R}^2$ such that

$$F(p_0 + h) = F(p_0) + J(p_0)h + \varepsilon(h) \cdot h,$$

where p_0 and h are now vectors in $\mathbb{R}^2$, and $|\varepsilon(h)| \to 0$ as $|h| \to 0$. This is similar to before, but if we write $F(x, y) = (u(x, y), v(x, y))$, then we can identify J with the matrix

$$J = \begin{bmatrix} \partial_x u & \partial_y u \\ \partial_x v & \partial_y v \end{bmatrix},$$

called the *Jacobian matrix*. A function is differentiable if this holds. This is a linear transformation; if we plug in a point we get a 2×2 matrix.

(We are being vague about what $\varepsilon(h)$ is; it may be better to write this in terms of the difference quotient, with absolute values everywhere.)

Since we can view complex differentiable functions as functions $\mathbb{R}^2 \to \mathbb{R}^2$, we can ask, what does being complex differentiable have to say about the Jacobian?

(For a complex-differentiable function we can write it like this; the converse is not true, meaning that being differentiable as a function on $\mathbb{R}^2$ doesn't make you complex differentiable.)

If f is holomorphic at z_0, then we have to consider

$$\frac{f(z_0 + h) - f(z_0)}{h}.$$

We're taking the limit as $h \to 0$, and h is a complex number; this means we need to take the limit along any path tending to 0 (e.g. a straight line, winding around, and so on), and all these limits have to agree. In particular, if I take the limit along the x-axis, i.e., I restrict $h = h_1 + i0$, then in this difference quotient we get

$$f'(z_0) = \lim_{h_1 \to 0} \frac{f(x + h_1 + iy) - f(x + iy)}{h_1}.$$

This limit is the same as $\partial_x f(z_0)$ by the definition of the partial (I just vary the x-component of f, and not the y-component).

Similarly, I should be able to recover the partial with respect to y by taking the limit along the purely imaginary axis — taking $h = 0 + ih_2$¡ this should equal

$$f'(z_0) = \lim_{h_2 \to 0} \frac{f(x + i(y + h_2)) - f(x + iy)}{ih_2} = \frac{1}{i} \partial_y f(z_0).$$

What you recover from this is that if f is holomorphic at z_0, then

$$f'(z_0) = \partial_x f(z_0) = \frac{1}{i} \partial_y f(z_0).$$

If we unwind these by writing $f(x, y) = (u(x, y), v(x, y))$ by splitting into real and imaginary parts, then unwinding this gives the *Cauchy–Riemann equations*: we get

$$\partial_x u = \partial_y v \text{ and } \partial_x v = -\partial_y u.$$

These are a requirement on the partial derivatives of f with respect to x and y (where we're identifying a function of a complex variable with a function in $\mathbb{R}^2$). So if you're holomorphic at z_0, these have to hold.

Then for a holomorphic function, if we try to look at the Jacobian matrix, these equations give some conditions on the Jacobian matrix — the Jacobian has to be of the form

$$J(z_0) = \begin{bmatrix} \partial_x u & -\partial_x v \\ \partial_x v & \partial_x u \end{bmatrix}.$$

In particular, we can rewrite this as

$$\sqrt{(\partial_x u)^2 + (\partial_x v)^2} \cdot \begin{bmatrix} \partial_x u / \sqrt{(\partial_x u)^2 + (\partial_x v)^2} & -\partial_x v / \sqrt{(\partial_x u)^2 + (\partial_x v)^2} \\ \partial_x v / \sqrt{(\partial_x u)^2 + (\partial_x v)^2} & \partial_x u / \sqrt{(\partial_x u)^2 + (\partial_x v)^2} \end{bmatrix},$$

where everything is evaulated at z_0. This is given by $\rho \cdot A$ where $\rho \in \mathbb{R}_{>0}$ and A is a matrix of the form

$$A = \begin{bmatrix} a & -b \\ b & a \end{bmatrix}$$

where $\det(A) = 1$, i.e., $A \in \mathrm{SO}_2(\mathbb{R})$. More generally, what happens with $\mathrm{SO}_2(\mathbb{R})$ is that angles are preserved. So this means that the linear trnasformation defined by the complex derivative of a holomorphic function has this special structure, that it preserves angles (it's called *conformal*). We'll get back to this later in the class.

We've seen a couple of amazing theorems and now we're going to start reaping consequences of them. A question back in our mind should be where this theory is coming from. We can compare that for a arbitrary differentiable function $\mathbb{R}^2 \to \mathbb{R}^2$ the Jacobian matrix can be anything; whereas here the Cauchy–Riemann equations have to hold, which means the Jacobian has to specifically be a matrix of this form.

Some other neat things that hold — these Cauchy–Riemann equations are pretty, but may be a pain to remember. So people often write them more succintly by introducing notation. One bit of notation you might see a lot is the notation ∂_z and $\partial_{\overline{z}}$ — we define

$$\partial_z = \frac{1}{2}(\partial_x - i\partial_y) \text{ and } \partial_{\overline{z}} = \frac{1}{2}(\partial_x + i\partial_y)$$

(this involves taking linear combinations of partial derivatives). f is holomorphic at z_0 if and only if the Cauchy–Riemann equations hold, which can be encapsulated as saying that $\partial_{\overline{z}} f = 0$ (sinc ethis says $\partial_x f + i\partial_y f = 0$, and $1/i = -i$). Then you just have to remember what ∂_z and $\partial_{\overline{z}}$ mean, and you can remember these equations.

Another way to derive this is noting that if f is complex differentiable then it should be real differentiable, and noting what linear transformations can be expressed as multiplication by a complex number — these two expressions match up, but in the top we have multiplication $f'(z_0) \cdot h$. We can rewrite this viewing f as a function $\mathbb{R}^2 \to \mathbb{R}^2$ and rewriting multiplication by a complex number as a linear operator on $\mathbb{R}^2$; the type of things you get this way are exactly what we described.

§6.2 Back to theory

This answers a question from a few weeks ago about a converse to the Cauchy theorem.

Theorem 6.4 (Morera's theorem)
Let f be continuous in a region D. If $\int_\Delta f = 0$ for every triangle Δ in D, then f is holomorphic in D.

This is essentially a converse to the Cauchy theorem, which says that if you're holomorphic then your integral is 0 on any closed loop. This is even stronger than the converse — we only have to check for triangles. In fact, we've done most of the proof already.

Proof. If D is convex, then f has an antiderivative F in D by the antiderivative lemma, meaning F is holomorphic and $F' = f$. But we just proved that when a function is holomorphic, its derivative is also holomorphic; so f is also holomorphic by the Cauchy derivative formula.

If D is not convex, it doesn't matter much — we're just trying to prove f is holomorphic in D, i.e., that f is complex-differentiable at every point $z_0 \in D$ (since being holomorphic at a point is a local condition). And we can just find a tiny disk containing z_0 in D (since D is an open set) and apply the previous argument, since disks are convex. ☐

The next result, which we won't prove yet, is kind of shocking.

Definition 6.5. A function $f\colon \mathbb{C} \to \mathbb{C}$ is called *entire* if f is holomorphic on $\mathbb{C}$.

In particular, polynomials are entire functions.

Definition 6.6. A function $f\colon S \to \mathbb{C}$ is *bounded* on S if there exists $B > 0$ such that $|f(z)| \leq B$ for all $z \in S$.

Theorem 6.7 (Liouville's theorem)

Every bounded entire function is constant.

So if you're a holomorphic function on all the complex plane, there has to be some direction where your absolute value becomes infinite along. This is in stark contrast to functions on $\mathbb{R}$ — for example, $f(x) = \cos(x)$ is bounded on $\mathbb{R}$, even though cos has infinitely many derivatives. So there's tons of bounded smooth functions on $\mathbb{R}$, but there are no such (interesting) ones in $\mathbb{C}$.

Remark 6.8. $\cos(\operatorname{Re}(z))$ is not holomorphic. But we will construct a function $\cos(z)$ to be an entire function, which will be unbounded by this theorem. What'll happen is that it's bounded on the real axis, but it becomes unbounded on the imaginary axes (and along any other axis) — then you get the hyperbolic cosine, where $\cos(iy) = \cosh(y)$.

There's also a generalized Liouville theorem — if your function grows slower than z^n, then it has to be polynomial. So any function that's not a polynomial has to have a direction where it diverges faster than any polynomial. You can classify all polynomials like that — is this function controlled by $(1+z)^n$ for some n? If yes, then it's polynomial. This is an example of the rigidity of complex functions.

Thursday's lecture will be recorded and put on Canvas; there will be a new Canvas section called Zoom/-Panopto and there will be a recording of Section 8. This might happen a couple of times this semester.

§7 October 3, 2023

The first thing we'll do is start expanding the class of functions that we have been working with.

§7.1 The exponential function

Definition 7.1. The *exponential* function $\exp\colon \mathbb{C} \to \mathbb{C}$ (written as $z \mapsto e^z$) is defined in the following way: for $z = x + iy$, we define

$$e^z = e^x \cos y + ie^x \sin y.$$

You can also define e^z in other ways (as a power series, or a solution to a differential equation).

Theorem 7.2

The function e^z is entire (i.e., is holomorphic on $\mathbb{C}$), and $(e^z)' = e^z$.

Remark 7.3. You can also *define* e^z in this way, as the unique solution to $(e^z)' = e^z$ with $e^0 = 1$.

We proved the fact that e^z is entire on the previous problem set. Here's a proof of the differential equation:

Proof. Once we know e^z is holomorphic, we can use the Cauchy–Riemann equations — we know

$$(e^z)' = \frac{d}{dx}e^z = \frac{d}{dx}(e^x \cos y + ie^x \sin y) = e^x \cos y + ie^x \sin y = e^z,$$

using the fact that $\frac{d}{dx}e^x = e^x$ (i.e., that we already know this fact about the real function). □

Theorem 7.4

For all $z_1, z_2 \in \mathbb{C}$, we have

$$e^{z_1+z_2} = e^{z_1}e^{z_2}.$$

Proof. We can again use the definition, the fact that $e^{x_1+x_2} = e^{x_1}e^{x_2}$ for real x_1, x_2, and the formulas

$$\cos(y_1 + y_2) = \cos y_1 \cos y_2 - \sin y_1 \sin y_2, \quad (2)$$
$$\sin(y_1 + y_2) = \sin y_1 \cos y_2 + \sin y_2 \cos y_1.$$

□

Theorem 7.5

We have $e^z = 1$ if and only if $z = 2\pi in$ for some $n \in \mathbb{Z}$.

Proof. We have $e^z = 1$ if and only if $e^x \cos y = 1$ and $e^x \sin y = 0$; since e^x is never 0, the second equation implies that $\sin y = 0$, and therefore $y = m\pi$ for some $m \in \mathbb{Z}$; now $e^x \cos y = e^x \cos m\pi = e^x(-1)^m$. This is equal to 1 if and only if $x = 0$ and m is even (since e^x is always positive), i.e., $z = 2\pi in$ for integer n. □

Corollary 7.6

There does not exist $z_0 \in \mathbb{C}$ such that $e^{z_0} = 0$.

Note that e^{z_0} might be negative, or might not be real; but it's never equal to 0.

Proof. If $e^{z_0} = 0$, then we would have

$$1 = e^0 = e^{z_0 - z_0} = e^{z_0}e^{-z_0} = 0e^{-z_0} = 0,$$

contradiction. □

Theorem 7.7

If $z = x + iy$, then $|e^z| = e^x$ and $\arg(e^z) = y$.

(This is from the homework.)

Corollary 7.8

The nth roots of unity are precisely $\zeta = e^{2\pi ik/n}$ for $k \in \{0, \dots, n-1\}$.

(We defined the nth roots of unity in the first class.)

Theorem 7.9

The function e^z is periodic with period $2\pi i$.

This means that $e^{z+2\pi i} = e^z$ for all $z \in \mathbb{C}$ — it's a function where adding $2\pi i$ doesn't affect its value.

Remark 7.10. Later in the class we'll try to find functions which are *doubly* periodic, meaning that there are two distinct periods. We're working with functions over a complex variable; so the fact that e^z is $2\pi i$-periodic means that if we go upwards by $2\pi i$, then the function repeats itself — there's one period direction. A function is *doubly periodic* if there are two distinct directions in which the function repeats itself — for example, it might be $2\pi i$-periodic as well as ω-periodic where ω is not purely imaginary.

It's a really interesting question to try to find such a function; we'll get to it much later. (It turns out that the function will need to have *singularities* (and therefore not exactly be analytic), which we'll see later.)

The function e^z is entire, so it has a power series expansion:

Theorem 7.11

For each $z \in \mathbb{C}$, we have

$$e^z = \sum_{n=0}^{\infty} \frac{z^n}{n!} = 1 + z + \frac{z^2}{2} + \frac{z^3}{3!} + \cdots.$$

Proof. We proved the Cauchy–Taylor theorem and the fact that $(e^z)' = e^z$; this means $(e^z)^{(n)} = e^z$, and $e^{z_0} = e^0 = 1$. So all the coefficients in the Taylor expansion are 1. □

Remark 7.12. Note that e^z is not a polynomial. For example, polynomials only have a finite Taylor expansion, and e^z doesn't. For example, we saw that f is a polynomial if and only if there exists $B > 0$ such that $|f(z)| \leq B(1 + |z|)^k$ for all z; and e^z doesn't satisfy this, if we take $y = 0$ and look on the real line. So e^z is not a polynomial because it grows too fast.

Alternatively, every polynomial has a root (by the fundamental theorem of algebra); e^z has no root, so it can't be a polynomial.

From the exponential function, we can build up a whole bunch of generalizations of functions from calculus.

Definition 7.13. We define the functions

$$\cosh(z) = \frac{1}{2}(e^z + e^{-z}) = \sum_{m=0}^{\infty} \frac{z^{2m}}{(2m)!},$$
$$\sinh(z) = \frac{1}{2}(e^z - e^{-z}) = \sum_{m=0}^{\infty} \frac{z^{2m+1}}{(2m+1)!},$$
$$\cos(z) = \frac{1}{2}(e^{iz} + e^{-iz}) = \sum_{m=0}^{\infty} (-1)^m \frac{z^{2m}}{(2m)!},$$
$$\sin(z) = \frac{1}{2i}(e^{iz} - e^{-iz}) = \sum_{m=0}^{\infty} (-1)^m \frac{z^{2m+1}}{(2m+1)!}.$$

If we restrict to z real, we recover the definitions from calculus of cos and sin.

We can use the properties of exp to write down some nice properties of these functions:

- $e^{iy} = \cos y + i \sin y$ (Euler's formula).
- $\cosh(iz) = \cos z$ and $\sinh(iz) = i \sin z$.
- The functions cosh, sinh, cos, and sin are all entire. (This follows from the fact that e^z is entire, and potentially the chain rule.)
- In contrast to when viewed as functions of a real variable, cos and sin are necessarily unbounded (because of Liouville's theorem — cos and sin are clearly not constant, since they're not constant on the real line). The same is true of cosh and sinh (though those are already unbounded in the real case as well). In particular, both are unbounded along the imaginary axis — they become cosh and sinh, which are unbounded as real functions.
- Because of the differential equation satisfied by e^z, we get nice equations for these functions as well: $\cosh' = \sinh$, $\sinh' = \cosh$, $\cos' = -\sin$, and $\sin' = \cos$.

Theorem 7.14 (Addition formulas)

We have the following addition formulas:

- $\cosh(z_1 + z_2) = \cosh z_1 \cosh z_2 + \sinh z_1 \sinh z_2$.
- $\sinh(z_1 + z_2) = \sinh z_1 \cosh z_2 + \cosh z_1 \sinh z_2$.
- $\cos(z_1 + z_2) = \cos z_1 \cos z_2 - \sin z_1 \sin z_2$.
- $\sin(z_1 + z_2) = \cos z_1 \sin z_2 + \sin z_1 \cos z_2$.
- $\cosh^2 z - \sinh^2 z = 1$.
- $\cos^2 z + \sin^2 z = 1$.

You can deduce these all from using the definitions and the addition properties of e^z. There are more curiosities about these functions in the notes.

The reason we spend so much time on this is that these are examples of entire functions that are not polynomials, and that have a lot of nice properties.

Now we'll get back to general theory.

§7.2 Sequences of functions

We'll now review some material from real analysis and apply it here.

We'll start with sequences of *numbers*:

Definition 7.15. Let $s_1, s_2, \ldots$ be a sequence in $\mathbb{C}$, and let $s \in \mathbb{C}$. We say that $s_n \to s$ (read as 's_n converges to s') as $n \to \infty$ if for all $\varepsilon > 0$, there exists $N \in \mathbb{N}$ such that for all $n \geq N$, we have $|s_n - s| < \varepsilon$.

So given ε, as long as we look far enough in the sequence, we're within ε of the limit.

From this definition of convergence of complex numbers, we can talk about what it means for a sequence of complex *functions* to converge.

Definition 7.16. Let $S \subseteq \mathbb{C}$ and suppose that $f_n \colon S \to \mathbb{C}$ is a sequence of functions on S, and $f \colon S \to \mathbb{C}$ is also a function on S. Then we say $f_n \to f$ *pointwise* as $n \to \infty$ if for every $z \in S$, we have $f_n(z) \to f(z)$ as $n \to \infty$.

So if we look at any $z \in S$ and look at the sequence of complex numbers obtained by evaluating our functions at z, if this sequence converges to $f(z)$ as a sequence of complex numbers, then we say $f_n \to f$ pointwise.

In symbols (combining these two definitions), the definition of pointwise convergence states that for every $z \in S$ and every $\varepsilon > 0$, there exists $N \in \mathbb{N}$ such that for all $n \geq N$ we have $|f_n(z) - f(z)| < \varepsilon$.

Example 7.17

Let $S = \mathbb{D} = \{z \in \mathbb{C} \mid |z| < 1\}$. Then $f_n(z) \to 0$ as $n \to \infty$ pointwise in $\mathbb{D}$.

We'll contrast pointwise convergence with uniform convergence (this is one of the major topics in 18.100):

Definition 7.18. Let $S \subseteq \mathbb{C}$ and suppose that $f_n \colon S \to \mathbb{C}$ is a sequence of functions on S, and $f \colon S \to \mathbb{C}$ is also a function on S. Then we say that $f_n \to f$ *uniformly* on S if for every $\varepsilon > 0$, there exists $N \in \mathbb{N}$ such that $|f_n(z) - f(z)| < \varepsilon$ for *every* $z \in S$.

The distinction between the two definitions is that in the pointwise case, we first fix z, and then we choose N (so N depends on both ε and z — we might have different N depending on which z we look at). But for uniform convergence, N depends only on ε — we need one value of N that works for *all* z (so N is independent of which $z \in S$ you take).

These two notions are not the same. If you converge uniformly then of course you converge pointwise, but not vice versa. We can see this from the above example:

Example 7.19

The sequence z^n does *not* converge to 0 uniformly on $\mathbb{D}$.

So we have an example of a function that doesn't converge uniformly, but does converge pointwise. The point here is that as z gets closer and closer to the boundary of the unit disk, it takes longer and longer for $|z^n|$ to get small.

§7.3 Properties of uniform convergence

The reason we talk about uniform convergence is that a lot of things in analysis are obtained by some sort of limiting process (e.g. power series are defined by taking a limit of finite sums). You want to know what properties the limit of a sequence of functions inherits from the sequence of functions — is differentiation inherited? If we take integrals, do they converge to the right value?

For all these things, pointwise convergence isn't enough to ensure that these nice properties are inherited; but uniform convergence is. For example, a sequence of continuous functions that converges uniformly must converge to a continuous function as well.

Theorem 7.20

The uniform limit of a sequence of continuous functions is itself continuous. In other words, let $S \subseteq \mathbb{C}$ and let $f_n \colon S \to \mathbb{C}$ be a sequence of *continuous* functions, and suppose that $f_n \to f$ uniformly on S. Then f is also continuous on S.

We've probably seen the proof in real analysis.

Proof. Take $z_0 \in S$, so we want to show f is continuous at z_0. For every $z_1 \in S$, we can write

$$f(z_1) - f(z_0) = (f(z_1) - f_n(z_1)) + (f_n(z_1) - f_n(z_0)) + (f_n(z_0) - f(z_0))$$

(this is the usual trick of adding and subtracting — we know something about the sequence of functions, so we want to make it appear). To show continuity at z_0, we want to take $z_1 \to 0$ and prove that the left-hand side tends to 0. To do so, given ε, we choose sufficiently large n such that the first and last terms are both small; then we have fixed n and we know f_n is continuous, so we can choose z_1 sufficiently close to z_0 such that the second term is small as well. The key point here is that n is independent of z (since z_1 varies — we need to be able to make the first two terms small independently of z_1). □

Remark 7.21. The notion of Cauchy sequences carries over to $\mathbb{C}$ as well; the complex numbers are complete, so all convergent sequences are Cauchy. Even if we didn't prove this explicitly in 18.100, we did it implicitly — we probably proved that $\mathbb{R}^n$ is complete (as a metric space), and this is the same because we can view the complex numbers as ordered pairs of real numbers.

The next property we'll see is that integration behaves well with uniform limits — we can interchange limits and integration if our functions uniformly converge.

Theorem 7.22

Let $f_n \colon [a, b] \to \mathbb{C}$ be a sequence of functions, and $f \colon [a, b] \to \mathbb{C}$ a function, such that all are Riemann integrable. Suppose that $f_n \to f$ uniformly on $[a, b]$. Then $\int_a^b f_n \to \int_a^b f$ as $n \to \infty$.

So taking limits and integration are two processes, and you can swap them when you have uniform convergence.

Remark 7.23. This is *not* true if we only assume $f_n \to f$ pointwise on $[a, b]$. The canonical example is taking f_n to be increasingly pointy triangles — isosceles triangles of base width $\frac{1}{n}$ (starting at 0) and height $2n$, so that $\int f_n = 1$. But $f_n \to 0$ pointwise on $[0, 1]$ (since if we fix z, we can take large enough n that the triangle happens before we get to z, so the functions eventually are identically 0). We have $\int 0 = 0$, but the integral of each of these functions is 1.

Proof. We can write $\int_a^b f_n - \int_a^b f = \int_a^b (f_n - f)$, so then

$$\left| \int_a^b f_n - \int_a^b f \right| \leq \int_a^b |f_n - f| \, .$$

We want to show that given ε, for large N the quantity on the left is smaller than ε. To do so, fix ε; by uniform convergence, we can choose N such that for all $n \geq N$, we have $|f_n(x) - f(x)| < \frac{\varepsilon}{b-a}$ for all $x \in [a, b]$. The the integral is at most $(b - a) \cdot \frac{\varepsilon}{b-a}$, and we're done. □

Here we need uniform convergence because the integral sees *all* x in the interval, so we need to choose N independent of x.

Differentiation is slightly more complicated (continuity and integration have nice clean statements).

Theorem 7.24

Let $f_n \colon [a, b] \to \mathbb{C}$ be a sequence of complex-valued functions with $f_n \in \mathcal{C}^1([a, b])$ (i.e., f_n is differentiable and f_n' is continuous). Suppose that the following two conditions hold:

(a) $f_n \to f$ pointwise on $[a, b]$.

(b) $f_n' \to g$ uniformly on $[a, b]$.

Then g is continuous, and $f' = g$.

(The fact that g is continuous follows from the fact that it's a uniform limit of continuous functions.)

So to show that the limit of the derivative is the derivative of the limit, we need the *derivatives* to converge uniformly; but we only need the functions themselves to converge pointwise.

Remark 7.25. This theorem doesn't work if we don't assume that the derivatives are continuous.

Remark 7.26. With real functions, we have to be careful about the conditions on our functions. For a long time, the prevailing wisdom was that functions shouldn't be so bad (e.g. having countably many bad points). But then around the early 1900s, this view changed — there were constructions of continuous functions that were nowhere differentiable, and so on.

Remark 7.27. In analysis, you're often trying to find functions solving a differential equation. If you make the class of functions you look at really big, it's easier to find a solution using a limiting process (by e.g. taking a minimizing sequence); you want the limit to be in the set you're looking at, so you need the set to be big. But sometimes the set is so big that you don't have differentiability or continuity of derivatives, and you need separate arguments to show that the function you found is smooth for other reasons. If you take Fourier analysis, you'll see functions that have 'half of a derivative' or a 'quarter of a derivative' or 'π derivatives' — and we use all these notions.

Part of the proof of this theorem uses the fundamental theorem of calculus:

Theorem 7.28 (FTC)

- If $g \in \mathcal{C}([a, b])$, then we have
$$\frac{d}{dx}\int_a^x g(y)\,dy = g(x)$$
in $[a, b]$.
- If $g = f' \in \mathcal{C}([a, b])$, then $\int_a^b g = f(b) - f(a)$.

You can ask about weakening 'continuous,' but that's another question for Fourier analysis and we won't worry about it here.

Proof of Theorem. Take $x \in [a, b]$. Then $f_n \to f$ pointwise, which means $f_n(x) - f_n(a) \to f(x) - f(a)$. But we have
$$f_n(x) - f_n(a) = \int_a^x f_n' \to \int_a^x g$$
for each $x \in [a, b]$ uniformly, using the previous theorem and the fact that $f_n' \to g$ uniformly.

So we now know $f_n(x) - f_n(a)$ tends to both $\int_a^x g$ and $f(x) - f(a)$, which means $f(x) - f(a) = \int_a^x g$.

Now we're trying to show $f' = g$; to do so, we can use the other part of FTC and differentiate with respect to x; this gives that $f'(x) = g(x)$. In particular, f is differentiable and equal to g. □

Next time we'll start applying this in the context of holomorphicity — we'll show that a uniform limit of holomorphic functions is holomorphic. (To do so, we're going to use the Cauchy theory about integrals.)

Remark 7.29. Note that in our last theorem, we were looking at functions in real variables.

§8 October 5, 2023

Last time we reviewed uniform convergence; today we'll apply it to functions on $\mathbb{C}$.

Theorem 8.1

Let $\{f_n\}$ be a sequence of functions $f_n\colon D \to \mathbb{C}$, where D is convex and each f_n is holomorphic on D, and suppose that $f_n \to f$ uniformly on D. Then f is holomorphic.

Proof. The key point will be our characterization of holomorphic functions — to check holomorphicity, by Morera's theorem we just need to check that the integral around every closed loop (or even triangle) is 0.

And we know this for each of the f_n's — for each closed loop γ, because f_n is holomorphic, we know $\int_\gamma f_n = 0$ for each γ. Meanwhile, we saw yesterday that the fact that $f_n \to f$ uniformly implies that f is continuous (the uniform limit of continuous functions is continuous), so we can certainly integrate it.

Claim 8.2 — We have $\int_\gamma f_n \to \int f$ for any path γ in D.

It suffices to check this for smooth paths (since for a piecewise smooth path, we can just check this for each piece, and sum the pieces).

Proof. Let $\gamma\colon [a, b] \to D$ be a smooth path. Then

$$\int_\gamma f_n = \int_a^b f_n(\gamma(t))\gamma'(t)\,dt.$$

Since $\gamma \in \mathcal{C}^1([a, b])$, then γ' is continuous. Since γ is a fixed path and $f_n \to f$ uniformly, then $f_n(\gamma(t))\gamma'(t) \to f(\gamma(t))\gamma'(t)$ uniformly as well. So this converges to $\int_a^b f(\gamma(t))\gamma'(t) = \int_\gamma f$. (We saw last class that if we have a sequence of functions converging uniformly, then their integrals converge as well.) □

So since $\int_\gamma f_n = 0$ for every n and every closed loop, this implies $\int_\gamma f = 0$ for every closed loop, and therefore f is holomorphic by Morera's theorem. □

Remark 8.3. A consequence of Cauchy's integral theorem is that if we take two points p and q, and we take two different paths from p to q, then the integral of a holomorphic function will be equal along both paths — since the integral along the loop formed by the two paths (one in reverse) is 0. (This is true as long as we're inside a convex region.)

Remark 8.4. This is true even without 'convex' — given an arbitrary region, to check holomorphicity at a point z we can pick a disk around z and apply this proof on that disk (which is convex).

In the same vein, we have the following local version:

Theorem 8.5

Let $f_n\colon D \to \mathbb{C}$ where D is any region, f_n are holomorphic in D. Suppose that for each closed disk $E \subseteq D$, $f_n \to f$ uniformly on E. Then $f \in \mathcal{H}(D)$, and for each closed disk $E \subseteq D$, $f_n' \to f'$ uniformly in E.

(The first statement is what we just talked about in words.)

Corollary 8.6

In the same setup, for each closed disk $E \subseteq D$ and each k, $f_n^{(k)} \to f^{(k)}$ uniformly on E.

This is a strengthening of the previous result. We just discussed the first part — if E is a closed disk, you can perform the same proof on the interior of E.

Proof. The first statement (that $f \in \mathcal{H}(D)$) follows from the previous theorem on the interior of E.

For the second statement, we need the following lemma:

Lemma 8.7

If the closed disk $E = \{z \in \mathbb{C} \mid |z - z_0| \leq r\}$ is contained in D, then there exists $s > r$ such that $F = \{z \in \mathbb{C} \mid |z - z_0| \leq s\}$ is also contained in D.

Here we have some large region D, and a disk E inside it. The lemma says that if we have any such closed disk E inside D, we can always find a slightly larger closed disk F strictly contained in D.

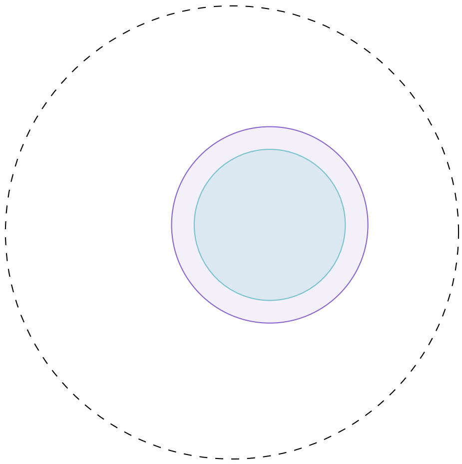

Proof. Assume not. Then we can find a sequence z_m with $|z_m - z_0| \leq r + \frac{1}{m}$ and $z_m \notin D$ for every m. (This is because for any slightly bigger radius we pick, we can find a point on that disk that's not inside the set.) But then this is a bounded sequence (because they all lie inside the disk of radius $r+1$) in the complex plane and therefore $\mathbb{R}^2$, so there exists a convergent subsequence $z_{m_k} \to z^*$. (This is because of compactness — if you have a sequence of points in a closed and bounded subset of $\mathbb{R}^n$, then it has a convergent subsequence. This is called sequential compactness.)

But then we must have $|z^* - z| \leq r$ by taking the limit. This means z^* is on the boundary of E, and is therefore inside E (because E is inside D). But since D is open, there is a ball $D' = \{z \mid |z - z^*| < \varepsilon\}$ around z^* that's also in D; and since $z_{m_k} \to z^*$ then D' must contain z_{m_k} for all sufficiently large k, which is a contradiction because z_{m_k} was chosen to be outside D. □

Now we're in business — we're trying to show that $f_n' \to f'$ uniformly, given that $f_n \to f$ uniformly. So now we can use the integral formula.

Fix this bigger disk F, and let γ be the path tracing out the boundary of F (counterclockwise). Then for every $z \in E$, we have

$$f_n'(z) = \frac{1}{2\pi i} \int_\gamma \frac{f_n(\zeta)}{(\zeta - z)^2}\, d\zeta.$$

Similarly, we showed that f is holomorphic in D (and therefore in F), so

$$f'(z) = \frac{1}{2\pi i}\int_\gamma \frac{f(\zeta)}{(\zeta - z)^2}\,d\zeta.$$

This means

$$f_n'(z) - f'(z) = \frac{1}{2\pi i}\int_\gamma \frac{f_n(\zeta) - f(\zeta)}{(\zeta - z)^2}\,d\zeta.$$

But by construction ζ is on γ, the circle of radius s around z_0, and z is inside the radius-r ball around z_0. So then $|\zeta - z| \geq |r - s|$, and therefore

$$|f_n'(z) - f'(z)| \leq \frac{1}{2\pi}\cdot 2\pi s \cdot \frac{\max_\gamma |f_n - f|}{(s-r)^2},$$

which goes to 0 as $n \to \infty$. This implies the desired uniform convergence (the right-hand side doesn't depend on z, so in particular it can be made small independently of z). □

The point of having the slightly larger disk was to have this lower bound for the denominator $|\zeta - z|$.

§8.1 Power series

Now we're going to apply this to power series.

Definition 8.8. A series $\sum_{n=1}^\infty a_n$ converges to s if the sequence of partial sums $s_N = \sum_{n=1}^N a_n$ converges to s as $N \to \infty$.

If the a_n are functions rather than numbers, then the same notion of convergence applies, and we can distinguish between pointwise and uniform convergence in the same way.

Definition 8.9. A series $\sum_{n=1}^\infty a_n(x)$ converges to $s(x)$ uniformly or pointwise if the sequence of partial sums $s_N(x) = \sum_{n=1}^N a_n(x)$ converges to $s(x)$ uniformly or pointwise.

The main tool to test for uniform convergence is called Weirstrass's test; this is probably familiar from analysis.

Theorem 8.10 (Weirstrass test)

Let $a_n\colon X \to \mathbb{C}$ be functions on a set X. Suppose that there exist positive real numbers b_n such that $|a_n(x)| \leq b_n$ for every $x \in X$ and $n \in \mathbb{N}$, and suppose that $\sum_{n=1}^\infty b - n$ converges. Then $\sum_{n=1}^\infty a_n(x)$ converges uniformly on X.

So if we can bound these functions above for each n by b_n which form convergent series (and this bound is independent of x), then this series converges uniformly.

Proof. Let $s_n(x) = \sum_{j=1}^n a_j(x)$ and $t_n = \sum_{j=1}^n b_j$. Then for all $m > n$, we have

$$s_m(x) - s_n(x) = a_{n+1}(x) + \cdots + a_m(x),$$

so by the triangle inequality

$$|s_m(x) - s_n(x)| \leq \sum_{j=n+1}^m |a_j(x)| \leq \sum_{j=n+1}^m b_j = t_m - t_n.$$

But since the b_n's converge, their partial sums form a Cauchy sequence; so for every $\varepsilon > 0$, there exists $N = N(\varepsilon)$ such that $|t_m - t_n| < \varepsilon$ for all $m, n \geq N$. And by this bound, then $|s_m(x) - s_n(x)| < \varepsilon$ for every $n, m \geq N$. (Note that the b_n's have nothing to do with x, so N is independent of x.)

So now for each x, $s_n(x)$ is a Cauchy sequence, and therefore converges to some number $s(x)$. Now we can take $m \to \infty$ and we get that

$$|s(x) - s_n(x)| \leq \varepsilon \text{ for all } n \geq N$$

(where N is just a function of ε, and this is true for all x), which gives that $s_n(x) \to s(x)$ uniformly on X. □

Remark 8.11. The point here is that we can choose N independent of x; it is possible to take the b_n's to be functions of x, as long as they converge uniformly. This is called the dominated convergence theorem.

Definition 8.12 (Power series)**.** A series of the form $\sum_{n=0}^{\infty} c_n(z - z_0)^n$ is called a power series about z_0.

Example 8.13

We have seen the Cauchy–Taylor series — if we have an analytic function f, we can associate to it its Cauchy–Taylor series, and that is a power series.

We'll now reprove another theorem from analysis class about complex functions, about radius of convergence of a power series. For this, we first need a lemma.

Lemma 8.14

If a power series $\sum_{n=0}^{\infty} c_n(z-z_0)^n$ converges at some $z_1 \neq z_0$, then it also converges for each z in the open disk $|z - z_0| < |z_1 - z_0|$; and if $r < |z_1 - z_0|$, then the convergence is uniform in the disk $|z - z_0| \leq r$.

A power series obviously converges at z_0, since all the sums are 0.

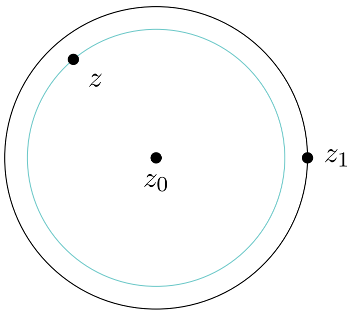

Proof. It suffices to prove the second statement. Let $r < |z_1 - z_0|$. If z is such that $|z_1 - z_0| < r$, then we can write

$$c_n(z - z_0)^n = c_n(z_1 - z_0)^n \frac{(z - z_0)^n}{(z_1 - z_0)^n}.$$

The series with just the first terms converges; so in particular, $|c_n(z_1 - z_0)^n|$ must be bounded by B for some $B > 0$. But then we can apply the Weirstrass test — let

$$\lambda = \left|\frac{z - z_0}{z_1 - z_0}\right| = \frac{r}{|z_1 - z_0|} < 1.$$

Then looking back at our expression, $|c_n(z - z_0)^n| \leq B\lambda^n$. And $\lambda < 1$, so this is a geometric series and therefore converges; then by the Weirstrass test our power series converges uniformly. □

Theorem 8.15

For each power series, exactly one of the three following scenarios holds:

(0) The poewr series converges only at $z = z_0$.

(∞) The power series converges for every $z \in \mathbb{C}$, and converges uniformly on the closed disk of radius r for every $r > 0$.

(R) There exists $R > 0$ such that the series diverges for every z with $|z - z_0| > R$ and converges for every z with $|z - z_0| < R$; furthermore, for each $r < R$ it converges uniformly on the disk $|z - z_0| \leq r$.

The number R is called the radius of convergence (it's either 0, ∞, or a finite number); in the last case you diverge outside the ball, converge inside the ball, and converge uniformly on every smaller ball.

This is very nice — it completely characterizes the behavior of power series (except when $|z - z_0| = R$).

Remark 8.16. It in fact converges absolutely on the smaller balls as well, though we won't make much use of this.

Remark 8.17. We haven't mentioned what happens on the disk of radius R. The theorem is agnostic on this, and either result is possible — if $R > 0$, then the power series may converge at some or all points z such that $|z - z_0| = R$. For example, when $R = 1$:

- $\sum z^n$ diverges for all $|z| = 1$;
- $\sum \frac{z^n}{n^2}$ converges for all $|z| = 1$;
- $\sum \frac{z^n}{n}$ diverges when $z = 1$ and converges when $z = -1$, and in fact for all $|z| = 1$ with $z \neq 1$. (For $z = -1$ it's the alternating harmonic series from calculus class; but for any $z \neq 1$ it converges by the same mechanism — if $z = e^{i\theta}$, then we have cos and sin which exhibit oscillatory behavior and cause cancellation.)

Proof. The proof is a quick consequence of the above lemma, which says that once you find one point where the series converges, you know you converge inside the disk defined by that point.

So define S to be the set of all positive numbers $|z_1 - z_0|$ for which the power series converges at $z = z_1$ (with $z_1 \neq z_0$).

If S is empty, then we are in scenario (0) — there is no z_1 for which the power series converges.

If S is unbounded, then we can find a sequence of z_n, such that $|z_n - z_0| \to \infty$ and the power series converges at those z_n; then we can apply the lemma with $z_1 = z_n$, and we are in scenario (∞). (The lemma tells us once we converge at z_1, we converge inside the disk it defines, and uniformly in any smaller disk.)

Lastly, if S is bounded, then let R be its least upper bound. (It's a set of real numbers and is bounded, so it has a least upper bound.) Now we are in scenario (R) — if z has $|z - z_0| < R$ then we have to diverge because R is an upper bound. If we're at a point z with $|z - z_0| < R$, then we can apply the lemma. And we're done. □

Remark 8.18. This is a lot like what we saw in calculus, but now 'radius' actually defines the radius of a circle, which is more satisfying.

Now let's combine this with what we saw in the beginning of the class.

Corollary 8.19

Each power series with nonzero (possibly ∞) radius of convergence must converge to a holomorphic function f in the disk $\{z \mid |z - z_0| < R\}$. Moreover, the series is necessarily the Cauchy–Taylor series of f, meaning that

$$c_n = \frac{f^{(n)}(z_0)}{n!}.$$

So once we have a nonzero radius of convergence, we have to converge to a holomorphic function, and the series is necessarily the Cauchy–Taylor series.

Proof. We have a series $f = \sum_{n=0}^{\infty} c_n(z - z_0)^n$; and we know the partial sums $f_n(z) = \sum_{j=1}^{n} c_j(z - z_0)^j$ converge uniformly on $|z - z_0| \leq r$ for any $r < R$. Each of these are holomorphic functions, so then f is holomorphic on $|z - z_0| \leq r$ for all $r \leq R$, and therefore on the open disk of radius R.

But now f is holomorphic, so all its derivatives $f^{(n)}$ exist and are holomorphic. And we can evalulate that $f^{(n)}(z_0) = c_n \cdot n!$. □

Next time we'll see some consequences, about zeroes about holomorphic functions.

§9 October 12, 2023

Next we'll look at zeros of holomorphic functions; then we'll discuss *singularities* (which occur when functions fail to be holomorphic at a point).

§9.1 Zeros

Definition 9.1. Let $f \in \mathcal{H}(D)$ for a region D, and assume that f is not identically 0. Then each $z_0 \in D$ at which $f(z_0) = 0$ is called a *zero* of f.

Remark 9.2. Constant functions have no zeros. (We exclude $f \equiv 0$ from our definition.)

Example 9.3

The zeros of a polynomial are called its roots.

Theorem 9.4

Let f be holomorphic in a region D. Suppose that there is a point $z_0 \in D$ at which f and all its derivatives are 0 — i.e.,

$$f^{(n)}(z_0) = 0$$

for all $n = 0, 1, 2, \ldots$. Then f is identically 0.

Proof. This follows from the given condition and the Cauchy–Taylor formula — we know that

$$f(z) = \sum_{n=0}^{\infty} \frac{f^{(n)}(z_0)}{n!}(z - z_0)^n.$$

But every term in this series is 0, so $f = 0$.

This formula holds at every open disk centered at z_0 inside D; so we have $f(z) = 0$ for all $z \in D_0$, where D_0 is an open disk centered at z_0 inside D.

We want to conclude that f is zero in the whole region. In order to show this, let's take any other point z in the region. We can't get to z just by finding one big disk on which this holds. But since we know the region is path-connected, we can connect z and z_0 with some path γ; then we'll tru to find overlapping disks from z_0 to z (bootstrapping the fact that f is 0 along these disks) — where the center of the next disk is inside the previous disk.

Note that the Cauchy–Taylor formula implies that for each $z \in D_0$, we actually have $f^{(n)}(z) = 0$ for all n as well (since we can differentiate the Cauchy–Taylor formula term by term).

We need a lemma that says we can find such a sequence of disks.

Lemma 9.5

Let $D \subseteq \mathbb{C}$ be any region, and let $p, q \in D$. Then there exists a finite sequence of points $z_0 = p$, z_1, $\ldots$, $z_k = q$ and open disks D_0, D_1, $\ldots$, $D_{k-1} \subseteq D$ centered at z_0, $\ldots$, z_{k-1} respectively, such that $z_{j+1} \in D_j$ for every $j = 0, \ldots, k-1$.

Assuming this, we can finish the proof of the theorem — given any $z \in D$, take $p = z_0$ and $q = z$ in the lemma, and choose disks D_0, $\ldots$, D_{k-1} as in the lemma. Then we can use the same argument repeatedly to see that for each j we have $f^{(n)}(z_j) = 0$ for all n, given this statement for z_{j-1}. This means in particular that $f(z) = 0$. $\square$

Now it remains to prove this lemma, which is just a lemma about regions in the complex plane. We'll need two smaller lemmas to prove it:

Lemma 9.6 (Worm lemma)

Let $D \subseteq \mathbb{C}$ be a region, and γ a path in D. Then there exists $r > 0$ such that for every $z \in \gamma$, we have $D(z, r) \subseteq D$ (where $D(z, r) = \{w \in \mathbb{C} \mid |w - z| < r\}$).

This amounts to saying we have a little 2-neighborhood around γ that is still inside the region D.

Lemma 9.7 (Bowstring lemma)

For every path γ from p to q in $\mathbb{C}$ of length L, we have $|p - q| \le L$.

This says that the straight line between p and q is shorter than any other path.

We'll first prove the lemma assuming these two sublemmas.

Proof of Lemma. Since D is a region, there exists a path γ from p to q. Let $r > 0$ be given by the worm lemma. Now choose points z_1, $\ldots$, z_{k-1} on γ such that the length of the pieces of γ between z_j and z_{j+1} are less than r. By the bowstring lemma, then $|z_0 - z_1|$, $\ldots$, $|z_k - z_{k-1}| < r$. Since the curved part of the path is less than r, so is the straight line segment between them; so then the disk of radius r centered at z_0 contains z_1, and so on. This gives the desired sequence of disks. $\square$

Proof of worm lemma. For this, we can use a compactness argument. Suppose not — suppose there is no such radius r. This means we can find a sequence of points and a sequence of collapsing disks that don't quite fit inside D — there exist points $z_n \in \gamma$ and $z_n' \notin D$ with $|z_n - z_n'| \to 0$ as $n \to \infty$.

Now consider a parametrization $\gamma\colon [a, b] \to \mathbb{C}$. Then we can find $t_n \in [a, b]$ such that $\gamma(t_n) = z_n$ for every n. Since $[a, b]$ is compact and we have a sequence of points, we can find a convergent subsequence; so there exists a subsequence $t_{n_k} \to t^*$ for some $t^* \in [a, b]$. Then we can consider the point $z^* = \gamma(t^*)$; this is a point on γ and therefore in D.

But γ is a continuous function, so $z_{n_k} = \gamma(t_{n_k}) \to \gamma(t^*) = z^*$. But now we have a contradiction — we have $z'_{n_k} \to z^*$ as well, and each z'_{n_k} is supposed to not be in D. But since $z^* \in D$ and D is open, we can find a small disk around z^* fully contained in D; eventually the z'_{n_k} will have to be inside this disk and therefore inside D, which is a contradiction. □

Remark 9.8. The existence of a convergent subsequence is a property of compactness — any sequence of points inside a closed and bounded interval has a convergent subsequence.

Proof of bowstring lemma. First let's assume that $\gamma\colon [a, b] \to \mathbb{C}$ is a smooth path. Then

$$q - p = \gamma(b) - \gamma(a) = \int_a^b \gamma'(t)\, dt,$$

which means

$$|q - p| = \left| \int_a^b \gamma'(t)\, dt \right| \le \int_a^b |\gamma'(t)|\, dt = L$$

by the definition of length. If we have a piecewise smooth path, then this holds for each of the pieces, and the length of a piecewise smooth path is the sum of the lengths. So if our break-points are p, p_1, p_2, $\ldots$, q, then

$$|q - p| = |q - p_2 + p_2 - p_1 + p_1 - p| \le |q - p_2| + |p_2 - p_1| + |p_1 - p|,$$

and each of these is at most the length of the corresponding piece of γ by the previous argument. □

§9.2 Some consequences

This theorem has some striking consequences:

Definition 9.9. Suppose that $f \in \mathcal{H}(D)$ is not identically 0. Let $z_0 \in D$ be a zero of f. Then the *multiplicity* of z_0 is the integer $m \in \mathbb{N}$ so that $f(z_0) = f'(z_0) = \cdots = f^{(m-1)}(z_0) = 0$ and $f^{(m)}(z_0) \neq 0$.

Note that such m always exists by the above theorem — if there were not such an m, then we would need to have $f(z_0) = f'(z_0) = \cdots$, and then f would have to be identically 0 by the theorem.

We also had something similar for multiplicities of roots of a polynomial; as an exercise you can check that the two notions of multiplicity for polynomials are the same.

Here's another consequence:

Corollary 9.10

Let $f \in \mathcal{H}(D)$ for a region D, and let $z_0 \in D$ be a zero of f of multiplicity m. Then there exists a function $g \in \mathcal{H}(D)$ such that:

- $g(z_0) \neq 0$.
- $f(z) = (z - z_0)^m g(z)$ for all $z \in D$.

This looks a lot like factoring out zeros of a polynomial; this states that you can do the same for holomorphic functions. (For polynomials g was another polynomial not vanishing at z_0; here g is a holomorphic function vanishing at z_0.)

Proof. This will follow from the fact that we have a power series representation of f. Take z_0, and let D_0 be a disk centered at z_0 and contained in D. By the Cauchy–Taylor formula, we can write

$$f(z) = \sum_{n=0}^{\infty} \frac{f^{(n)}(z_0)}{n!}(z - z_0)^n.$$

But since z_0 is a zero of f with multiplicity m, the first m of these terms have coefficient 0; so this is equal to the sum

$$\sum_{n=m}^{\infty} \frac{f^{(n)}(z_0)}{n!}(z - z_0)^n.$$

Each of the terms in the sum is divisible by $(z - z_0)^m$, so we can pull out $(z - z_0)^m$ from each to write

$$f(z) = (z - z_0)^m \varphi(z),$$

where

$$\varphi(z) = \sum_{n=m}^{\infty} \frac{f^{(n)}(z_0)}{n!}(z - z_0)^{n-m}.$$

Note that the power series for φ converges in D_0, because it converged before and all we did was factor out a constant from each term. So this means φ is holomorphic in D_0; and $\varphi(z_0) \neq 0$.

We're not quite done yet, because we need this formula to hold in all of D, and our candidate φ is something we can only get our hands on inside the smaller disk D_0. So how should we define g outside this disk? Well, we can just divide f by $(z - z_0)^m$ — set

$$g(z) = \begin{cases} \frac{f(z)}{(z-z_0)^m} & \text{if } z \in D \setminus \{z_0\} \\ \varphi(z) & \text{if } z \in D_0. \end{cases}$$

In $D_0 \setminus \{z_0\}$ we have competing definitions for g, but the two definitions agree because we have $f(z) = (z - z_0)^m \varphi(z)$. So g is well-defined; and we have $g(z_0) = \varphi(z_0) \neq 0$.

To check that g is holomorphic, first g is holomorphic in $D \setminus \{z_0\}$ (because $(z - z_0)^{-m}$ and f both are), so it suffices to check that it is holomorphic in D_0. But φ is holomorphic in D_0 and $g = \varphi$ on D_0, so g is holomorphic on D_0 as well, and therefore at z_0. $\square$

Remark 9.11. Why doesn't this work in $\mathbb{R}$? We don't have a well-defined notion of multiplicity — you can have $\mathcal{C}^\infty$-smooth functions (all derivatives exist) where all the derivatives vanish at 0, for example

$$f(x) = \begin{cases} e^{-1/x^2} & x \neq 0 \\ 0 & x = 0. \end{cases}$$

But if we assume that there *is* a multiplicity (we have a zero with only finitely many derivatives), can we do this? If we have a power series representation, then yes.

Remark 9.12. Why is $\varphi(z_0) \neq 0$? We defined

$$\varphi(z) = \sum_{n=m}^{\infty} \frac{f^{(n)}(z_0)}{n!}(z - z_0)^{n-m} = \frac{f^{(m)}(z_0)}{m!} + \frac{f^{(m+1)}(z_0)}{(m+1)!}(z - z_0) + \cdots.$$

But each of the later terms all have $z - z_0$, so when we plug in $z = z_0$, they all go away; and we're just left with the first term. And since we assumed f is a zero with multiplicity m, the first $m - 1$ derivatives are all 0 and $f^{(m)}(z_0)$ is *not* 0; so then

$$\varphi(z_0) = \frac{f^{(n)}(z_0)}{m!} \neq 0.$$

Corollary 9.13 (Isolation of zeros)

Let $f \in \mathcal{H}(D)$ for a region D, and assume f is not identically 0. Then the zeros of f are *isolated* — i.e., for every zero z_0, there exists $\delta > 0$ such that f has no zeros in the punctured disk $0 < |z - z_0| < \delta$.

Proof. The function g above is holomorphic, and therefore continuous; so since $g(z_0) \neq 0$, we can find a small disk around z_0 on which $g \neq 0$. Then $f = (z - z_0)^m g$, and the two pieces are both nonzero on that disk; so f is nonzero on that disk as well. □

Remark 9.14. Is there an example of such a f where there is no universal δ for all zeros? You could potentially have a function that has isolated 0's getting closer and closer with no accumulation point — for example $\sin(z^2)$. But the next corollary is along these lines.

Corollary 9.15 (Identity theorem)

Let $f, g \in \mathcal{H}(D)$, and suppose that $f = g$ on a line segment $[p, q]$ with $p \neq q$. Then $f = g$ on D.

So if two holomorphic functions agree on a line segment, then they have to agree everywhere.

Proof. Let $h = f - g$. If $f \not\equiv g$, then $h \not\equiv 0$. But we have $h = 0$ on $[p, q]$, so its zeros are not isolated; this is a contradiction, so we must have $h \equiv 0$. □

Example 9.16

There is only one entire function f that agrees with $\cos(z)$ on the real line.

So there is only one way to holomorphically extend cos, and it's the one we defined.

Proof. The function $f(z) = \frac{1}{2}(e^{iz} + e^{-iz})$ (which we defined as $\cos z$) is entire and agrees with cos on $\mathbb{R}$. Suppose that g is any other entire function agreeing with cos on $\mathbb{R}$. Then $f - g = 0$ on $\mathbb{R}$, which means it must be identically 0 and we must have $g \equiv f$. □

So if you find a way to extend a function on $\mathbb{R}$ to one on $\mathbb{C}$, then that's the *only* way — this is one example of the rigidity of complex functions.

Remark 9.17. Note that a function like $f(x + iy) = x + y$ is *not* holomorphic. (You can check it fails the Cauchy–Riemann equations.)

Remark 9.18. Given some isolated set, can we find a function whose zeros are that set? The answer is yes — in fact, we'll be able to do this given zeros and poles, and knowing both, there'll be a lot of rigidity.

Given *finitely* many zeros, you can take the function $f(z) = \prod_{j=1}^{n}(z - z_j)$; and any other function with those zeros will have some relation to f. With infinite zeros it's harder.

Remark 9.19. Is there a description of when a real function can be extended to an entire one?

Not all (even very nice-looking) functions can — for example

$$f(x) = \frac{1}{1+x^2}$$

is a beautiful function with nothing wrong with it (it's smooth and goes to 0 at ∞), but then you see that its power series expansion is

$$1 - x^2 + x^4 - x^6 + \cdots.$$

There's nothing wrong with the function at $x = 1$ or $x = -1$, but the power series has radius of convergence 1 — this doesn't make any sense in calculus or real analysis.

You don't even think about this until you take complex analysis. But this function considered over $\mathbb{C}$ has a big problem at i and $-i$ — we get division by 0. And i is on a disk of radius 1 — that's why the real power series has a radius of convergence 1.

§10 October 17, 2023

§10.1 Reflection principle

Last time we forgot to give one more application of the theory of zeros for holomorphic functions.

Corollary 10.1 (Reflection principle)
Suppose that $f \in \mathcal{H}(\mathbb{C})$ be entire and $f(x) \in \mathbb{R}$ for all $x \in \mathbb{R}$. Then $f(\overline{z}) = \overline{f(z)}$ for all $z \in \mathbb{C}$.

There are two very clean arguments that prove this.

Proof 1. This proof doesn't require the theory about zeros, but it does use the fact we have Cauchy–Taylor series. Since f is entire, we can expand f as a Taylor series about 0 — so we have

$$f(z) = \sum_{n=0}^{\infty} \frac{f^{(n)}(0)}{n!} z^n$$

for all $z \in \mathbb{C}$. But then we have

$$\overline{f(z)} = \overline{\sum_{n=0}^{\infty} \frac{f^{(n)}}{(}0)n!z^n} = \sum_{n=0}^{\infty} \overline{\frac{f^{(n)}(0)}{n!}} \overline{z}^n.$$

But f is 0 on $\mathbb{R}$, so $f^{(n)}$ is also 0 on $\mathbb{R}$; and this is equal to

$$\sum_{n=0}^{\infty} \frac{f^{(n)}(0)}{n!} \overline{z}^n = f(\overline{z}).$$ □

Remark 10.2. We did prove that the Taylor series for an entire function converges everywhere.

Proof 2. For *any* entire function, we claim that the function

$$f^*(z) = \overline{f(\overline{z})}$$

is also entire. To see this, we just need to show that its derivative exists for all z; we have

$$\frac{f^*(z_1) - f^*(z)}{z_1 - z} = \frac{\overline{f(\overline{z_1})} - \overline{f(\overline{z})}}{z_1 - z} = \overline{\frac{f(\overline{z_1}) - f(\overline{z})}{\overline{z_1} - \overline{z}}} \to \overline{f'(\overline{z})}$$

as $z_1 \to z$. So the limit of this expression exists as $z_1 \to z$ for all z, and so f^* is holomorphic.

Now if f is real-valued on $\mathbb{R}$, then $f^*(x) = f(x)$ for all $x \in \mathbb{R}$ (since $\overline{x} = x$ and $\overline{f(x)} = f(x)$). So we have two entire functions that agree on $\mathbb{R}$; and by the identity principle, they must be the same function. So $f^* = f$, and taking the conjugate of both sides gives that $f(\overline{z}) = \overline{f(z)}$. $\square$

This wraps up the section on zeros. Next, we'll start treating isolated singularities.

§10.2 Isolated singularities

Definition 10.3. Let D be a region, and let $z_0 \in D$. If $f \in \mathcal{H}(D \setminus \{z_0\})$, then we say that z_0 is an *isolated singularity* of f.

(If a function is defined at a point, this is not a singularity.)

Remark 10.4. If we take the function

$$f(z) = \begin{cases} z & z \neq 0 \\ 7 & z = 0 \end{cases}$$

then we do *not* say 0 is a singularity; it's simply not holomorphic at 0. We only refer to singularities when we haven't yet defined the function at that point.

So if a function is holomorphic except at a point, that point is called an isolated singularity.

Example 10.5

The functions $\frac{\sin z}{z}$, $\frac{1}{z}$, and $e^{1/z}$ are all holomorphic in $\mathbb{C} \setminus \{0\}$, and therefore have an isolated singularity at $z_0 = 0$.

It turns out that these are three good examples to look at. We'll classify the types of isolated singularities into three types:

Definition 10.6. Let z_0 be an isolated singularity of f (i.e., f is holomorphic in $D \setminus \{z_0\}$ for some region D).

- We say z_0 is a *removable singularity* if it is possible to define $f(z_0)$ in such a way that the new function f defined on D is in $\mathcal{H}(D)$.
- We say z_0 is a *pole* of f if $|f(z)| \to \infty$ as $z \to z_0$ (for $z \neq z_0$).
- We say z_0 is an *essential singularity* if it is neither removable nor a pole.

So z_0 is a removable singularity if the singularity can be removed — the issue is just that f isn't defined at z_0, but we can simply define it. It's a pole if the function becomes unbounded as we approach the isolated singularity in all directions.

We'll learn that removable singularities and poles are perfectly fine, but essential singularities are real problems.

§10.3 Removable singularities

The following theorem says that we can characterize removable singularities.

Theorem 10.7 (Riemann)
Let z_0 be an isolated singularity of f. Then z_0 is a removable singularity if and only if f is bounded in some punctured disk $\{z \mid 0 < |z - z_0| < r\}$.

Proof. The forwards direction is on the homework. (You can define f at that point to be holomorphic; and then by continuity this function f must be bounded.)

The backwards direction is trickier (and is the one involving complex analysis). Suppose that f is bounded in $\{0 < |z - z_0| < r\}$ (it's also holomorphic on this disk by assumption); we want to try to prove that we can make it holomorphic.

Claim 10.8 — For $z \neq z_0$, we have

$$f(z) = \frac{1}{2\pi i} \int_\gamma \frac{f(\zeta)\, d\zeta}{\zeta - z}$$

for all $0 < |z - z_0| < r$, where γ is the circle $\{z \mid |z - z_0| = r\}$.

Proof. Consider an open disk of radius $r_1 > r$ containing γ; we restrict to this disk (so that we have suitable convex regions containing everything relevant).

Note that this looks like the Cauchy integral formula, but we don't yet know this — we only proved the Cauchy integral formula for holomorphic functions in a region, and f isn't yet defined at z_0 (and we're not claiming this for $z = z_0$, because f is of course not defined at z_0).

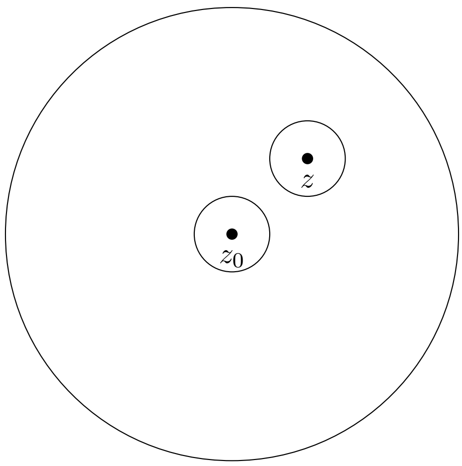

First we claim that

$$\int_\gamma \frac{f(\zeta)}{\zeta - z}\, d\zeta = \int_{C_0} \frac{f(\zeta)}{\zeta - z}\, d\zeta + \int_C \frac{f(\zeta)}{\zeta - z}\, d\zeta,$$

where γ is the large circle, C_0 is a small circle around z_0, and C is a small circle around z.

The way we prove this is by connecting everything with some overlapping paths.

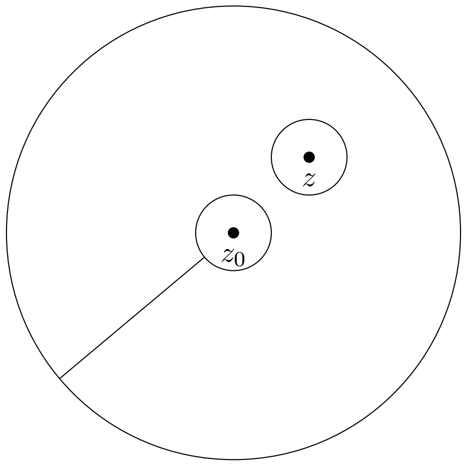

On all of the straight-line paths, the integral cancels out; so

$$\int_\gamma - \int_{C_0} - \int_C = \sum_{j=1}^{6} \int_{\gamma_j}.$$

But for each of the γ_j, we can find a convex region containing it; and since f is holomorphic in this region and it doesn't contain z, then $\frac{f(\zeta)}{\zeta - z}$ is a holomorphic function in each of these regions. So these integrals are all 0 by the Cauchy integral theorem.

Now for the second term, we have

$$\int_C \frac{f(\zeta)}{\zeta - z}\, d\zeta = 2\pi i f(z)$$

by the Cauchy integral formula. So the second term is equal to the thing we want; this means we just need to show the first term is 0.

And we know that f is bounded near z_0! And the denominator is also bounded (below). So we can simply shrink C_0 — let the radius of C_0 tend to 0. Then we have

$$\int_{C_0} \frac{f(\zeta)}{\zeta - z}\, d\zeta \to 0,$$

since we can bound it above by its maximum times its length.

So we can disregard this term, and we have

$$\int_\gamma \frac{f(\zeta)}{\zeta - z}\, d\zeta = 2\pi i f(z). \qquad \square$$

So now we can simply define

$$\widetilde{f}(z) = \frac{1}{2\pi i} \int_\gamma \frac{f(\zeta)\, d\zeta}{\zeta - z}$$

for all z with $|z - z_0| < r$, including z. This agrees with f on the disk, and defines it at z_0.

So now we just need to prove that f is analytic. But we can do this by the same method that we used to prove Cauchy's derivative formula (consider a difference quotient), and we're done. $\square$

§10.4 Poles

Now we'll talk about poles.

Theorem 10.9

If $f \in \mathcal{H}(D - \{z_0\})$ and z_0 is a pole, then for some $m \in \mathbb{N}$ the following holds:

(1) We have
$$f(z) = \frac{g(z)}{(z - z_0)^m}$$
for all $z \in D \setminus \{z_0\}$, for some $g \in \mathcal{H}(D)$ with $g(z_0) \neq 0$.

(2) We can write
$$f(z) = p\left(\frac{1}{z - z_0}\right) + h(z)$$
where p is a polynomial of degree m with no constant term, and $h \in \mathcal{H}(D)$.

(The first statement is analogous to factoring out zeros.)

Definition 10.10. We call the integer m as above the *multiplicity* of the pole.

Definition 10.11. The polynomial
$$p\left(\frac{1}{z - z_0}\right) = \frac{c_{-m}}{(z - z_0)^m} + \frac{c_{-m+1}}{(z - z_0)^{m-1}} + \cdots + \frac{c_{-1}}{(z - z_0)}$$
is called the *principal part* of the pole at z_0, and we call c_{-1} the *residue* of the pole.

We'll see that c_{-1} appears a lot (though it might seem random right now to give it a name).

Now let's prove the theorem.

Proof. Since z_0 is a pole, we have $|f(z)| \to \infty$ as $z \to \infty$; this means close enough to z_0, we can bound f away from 0. So there exists $r > 0$ such that $\{|z - z_0| < r\} \subseteq D$ and $f(z) \neq 0$ on $\{0 < |z - z_0| < r\}$. So the function
$$\varphi = \frac{1}{f}$$
is holomorphic in this punctured disk $0 < |z - z_0| < r$, and it has an isolated singularity at z_0. And this singularity is removable, since the fact that f tends to ∞ means φ is bounded. So we can simply remove it, and when we remove it we must have $\varphi(z_0) = 0$ (since $\varphi(z) \to 0$ as $z \to z_0$, and φ has to be continuous).

Now z_0 is a zero of φ, so it has a well-defined multiplicity m (as a zero). And moreover, from our theory of zeros we know that we can write
$$\varphi(z) = (z - z_0)^m \psi(z)$$
for some function ψ with $\psi(z_0) \neq 0$ which is holomorphic in $|z - z_0| < r$.

Now we can simply move things around. First, ψ has no zeros in this disk. So we can simply define
$$g(z) = \begin{cases} (z - z_0)^m f(z) & \text{for } z \in D \setminus \{z_0\} \\ \frac{(z-z_0)^m}{\varphi(z)} = \frac{1}{\psi(z)} & \text{for } |z - z_0| < r. \end{cases}$$

Since ψ is holomorphic and nonzero, its inverse is also holomorphic (and the two definitions agree).

To prove (2), we can simply apply (1). We know that $g \in \mathcal{H}(D)$, so in particular it's holomorphic in a small disk $|z - z_0| < r$. Then we can apply the Cauchy–Taylor formula for $g(z)$; this gives
$$f(z) = \frac{g(z)}{(z - z_0)^m} = \frac{1}{(z - z_0)^m} \sum_{n=0}^{\infty} \frac{g^{(n)}(z)}{n!}(z - z_0)^n = \sum_{-m}^{\infty} c_n (z - z_0)^n$$

(for some coefficients c_n). We can then simply separate out the part with powers between $-m$ and 1, which is a polynomial in $\frac{1}{z-z_0}$ with no constant term — so we can write this as

$$p\left(\frac{1}{z-z_0}\right)+h(z)$$

where h is holomorphic on this region.

So near z_0, we have defined h using this power series expansion. We need to define h on all of D; to do this, we simply can do so using f — we take

$$h(z)=\begin{cases}\sum_{n=0}^{\infty}c_n(z-z_0)^n & \text{for } |z-z_0|<r\\ f(z)-p\left(\frac{1}{(z-z_0)}\right) & \text{for } z\in D\setminus\{z_0\}.\end{cases}$$

These definitions agree on the overlap, so h is holomorphic everywhere on D. □

Next time we'll see partial fractions, another theorem about poles (which says how to express rational functions in terms of their poles and zeros). But first we'll talk about what an essential singularity does.

§10.5 Essential singularities

We've discussed poles and removable singularities; the last category of singularities is essential singularities. To get a small taste of what these are:

Theorem 10.12 (Casuratti–Weirstrass)
If $f\in\mathcal{H}(D\setminus\{z_0\})$ has an essential singularity at z_0, then all $\varepsilon>0$, all $\delta>0$, and all $w\in\mathbb{C}$, there exists z such that $0<|z-z_0|<\delta$ such that $|f(z)-w|<\varepsilon$.

What this means is that if f has an essential singularity at z_0, then for *any* complex number w you like and any ε-ball around w and any δ-ball around z_0, we can find some z inside the δ-ball sent to the ε-ball. So no matter how small a disk you look around z_0, the function goes *everywhere* — it's completely uncontrolled. There's even wilder theorems; but this shows this function is completely wild, and there's no taming it — it goes all over the complex plane, no matter how small of a neighborhood you look at. So you can do basically nothing with these singularities.

Proof. Assume that this is not true; then we can find some ε, δ, and w for which this is not true — so there are some $\varepsilon>0$, $\delta>0$, and $w\in\mathbb{C}$ for which there is no $z\in\{0<|z-z_0|<\delta\}$ for which $|f(z)-w|<\varepsilon$.

Now the function $g=\frac{1}{f-w}$ is holomorphic and bounded in $0<|z-z_0|<\delta$, so it has a removable singularity at z_0 — this means g can be defined at z_0 in a way that makes it holomorphic.

There are two possibilities. If $g(z_0)\neq 0$, then $f-w$ is bounded, so f is bounded; and therefore z_0 was a removable singularity of f by Riemann's theorem.

On the other hand, if $g(z_0)=0$, then $|f(z)-w|\to\infty$ as $z\to z_0$, which (since w is fixed) means that $|f(z)|\to\infty$ as well, i.e., z_0 is a pole of f.

This is a contradiction in either case, since we assumed z_0 was neither a removable singularity nor a pole. □

This tells us a function with an essential singularity does everything all at once around that singularity.

Example 10.13
The function $e^{1/z}$ has an essential singularity at 0.

§11 October 19, 2023

§11.1 Poles

Today we'll discuss partial fractions; first we'll summarize what we've seen about poles so far.

Definition 11.1. If $f \in \mathcal{H}(D \setminus \{z_0\})$, we say that f has a *pole* at z_0 if $|f(z)| \to \infty$ as $z \to z_0$.

Last time, we showed that we can write f as

$$f(z) = p\left(\frac{1}{z - z_0}\right) + h(z)$$

for $z \in D \setminus \{z_0\}$, where h is holomorphic in D and p is a polynomial

$$p\left(\frac{1}{z - z_0}\right) = \frac{c_{-m}}{(z - z_0)^m} + \cdots + \frac{c_{-1}}{(z - z_0)}.$$

We call the above the *principal part* of the pole; we call m the *multiplicity* of the pole; and we define c_{-1} as the *residue* $\operatorname{Res}(f; z_0) = c_{-1}$.

§11.2 Partial fractions

Theorem 11.2

Let $f = \frac{g}{h}$ where g and h are polynomials with no common roots, and with $h \neq 0$; let $z_1, \ldots, z_r$ be the zeros of h (equivalently, the poles of f), and let $p_1(\frac{1}{z - z_1}), \ldots, p_r(\frac{1}{z - z_r})$ be their principal parts (as poles of f). Then for all $z \in \mathbb{C} \setminus \{z_1, \ldots, z_r\}$, we have

$$f(z) = \sum_{j=1}^{r} p_j\left(\frac{1}{z - z_j}\right) + p(z)$$

where p is a polynomial.

The point is that we can pull out *all* the singular parts of f. (You use partial fractions in a similar way to compute integrals in calculus.)

Proof. Define the function

$$p(z) := f(z) - \sum_{j=1}^{r} p_j\left(\frac{1}{z - z_j}\right)$$

in $\mathbb{C} \setminus \{z_1, \ldots, z_r\}$. Note that p is holomorphic in $\mathbb{C} \setminus \{z_1, \ldots, z_r\}$, since f and the relevant polynomials only have problems at the z_j's. But moreover, all the singularities are removable, since we've subtracted out the principal parts of the poles. To see this more explicitly, near z_1 we have

$$p(z) = f(z) - p_1\left(\frac{1}{z - z_1}\right) - \sum_{j=2}^{r} p_j\left(\frac{1}{z - z_j}\right).$$

The second sum is holomorphic (these polynomials only have problems at z_j for $j \geq 2$, and we're only considering the area around z_1); and $f(z) - p_1(\frac{1}{z - z_1})$ has a removable singularity at z_1 (by the definition of a pole — which states that this difference is given by a holomorphic function on the neighborhood).

This means p is a rational function (since it's obtained by adding and subtracting polynomials) and has no poles; this means the denominator has no roots and is therefore constant, so p is a polynomial. □

Remark 11.3. For a general function with poles at $z_1, \dots, z_r$, the same argument gives that we can write $f(z) = \sum_{j=1}^r p_j(\frac{1}{z-z_j}) + h(z)$ where h is holomorphic in D (by again using the previous theory on small disks around each of the poles).

Remark 11.4. Can singularities cluster? We've learned everything there is to say about removable singularities; we're now going to talk about poles for a while. And then we'll maybe see the Picard theorem about essential singularities; but it's nice to keep these things separate.

§11.3 Calculating residues

First, let's see how to find the residue of a pole. This turns out to be easy.

Theorem 11.5

(a) If f has a simple pole at z_0 (meaning that $m = 1$), then

$$\operatorname{Res}(f; z_0) = \lim_{z \to z_0} (z - z_0) f(z).$$

(b) If $f = \frac{g}{h}$ where $g, h \in \mathcal{H}(D)$ and h has a *simple* zero in D at z_0, then

$$\operatorname{Res}(f, z_0) = \frac{g(z_0)}{h'(z_0)}.$$

(c) If f has a pole of multiplicity m at z_0, then

$$\operatorname{Res}(f; z_0) = \frac{1}{(m-1)!} \left(\frac{d}{dz}\right)^{m-1} \{(z - z_0)^m f(z)\}\bigg|_{z_0}.$$

The last case means that you consider the function $(z - z_0)^m f(z)$, evaluate its $(m-1)$th derivative, and plug in z_0. Note that (a) is a special case of both (b) and (c).

These formulas are easy to both compute and prove.

Proof. We'll first prove (a) — we have

$$f(z) = \frac{c_{-1}}{z - z_0} + h(z)$$

where h is holomorphic, so

$$(z - z_0) f(z) = c_{-1} + (z - z_0) h(z) \to c_{-1}$$

as $z \to z_0$ (since h is holomorphic at z_0, and therefore has a limit).

THe proof of (b) is similar — in this case $f = \frac{g}{h}$, where $h(z_0) = 0$ and $h'(z_0) \neq 0$ (since z_0 is a simple zero), then f has a simple pole, so we can apply (a): we have

$$(z - z_0) f(z) = \frac{g(z)}{\frac{h(z) - h(z_0)}{z - z_0}} \to \frac{g(z_0)}{h'(z_0)}$$

as $z \to z_0$.

For (c), we need to do a bit more work, but we can still use the formula — we have

$$f(z) = \frac{c_{-m}}{(z - z_0)^m} + \cdots + \frac{c_{-1}}{(z - z_0)} + h(z).$$

Now we multiply by $(z-z_0)^m$; this gives

$$(z-z_0)^m f(z) = c_{-m} + \cdots + (z-z_0)^{m-2}c_{-2} + (z-z_0)^{m-1}c_{-1} + (z-z_0)^m h(z),$$

and since h is holomorphic, we can express it using its Cauchy–Taylor formula as $c_0 + (z-z_0)c_1 + \cdots$; this means

$$(z-z_0)^m f(z) = c_{-m} + \cdots + (z-z_0)^{m-1}c_{-1} + c_0(z-z_0)^m + \cdots.$$

This is a power series on some small disk about z_0 (where the power series of h converges). So this function has a removable singularity, which we can remove by giving it the value c_{-m}.

Then after removing the singularity, this series has to be the Cauchy–Taylor series of $(z-z_0)^m f(z)$; and since it's a Cauchy–Taylor series, we know what its coefficients are — they come from differentiating this function and evaluating the derivatives at z_0. In particular, c_{-1} is the $(m-1)$st coefficient, so

$$c_{-1} = \frac{1}{(m-1)!} \cdot \left(\frac{d}{dz}\right)^{m-1} \{(z-z_0)^m f(z)\}\Big|_{z=z_0}. \qquad \square$$

Remark 11.6. We call the series

$$\frac{c_{-m}}{(z-z_0)^m} + \frac{c_{-1}}{(z-z_0)^{m-1}} + \cdots$$

the *Laurent series* for f.

Example 11.7

The function

$$f(z) = \frac{1}{z^2+1}$$

at $z = i$. To compute its pole, we can write it as $\frac{g}{h}$ where $g = 1$ and $h(z) = z^2+1$; then by (b)

$$\operatorname{Res}(f; i) = \frac{g(i)}{h'(i)} = \frac{1}{2i}$$

(since $h'(z) = 2z$).

Example 11.8

Let $0 < a < b$ be real numbers and consider the function

$$f(z) = \frac{1}{(z^2+a^2)(z^2+b^2)},$$

which has simple poles at ia and ib. To compute $\operatorname{Res}(f, ia)$, we can take $g(z) = \frac{1}{z^2+b^2}$ and $h(z) = z^2+a^2$; then

$$\operatorname{Res}(f; ia) = \frac{1}{(ia)^2+b^2} \cdot \frac{1}{2ia}.$$

Example 11.9

Take the function

$$f(z) = \frac{1}{(z^2+1)^m},$$

which has poles at $z = \pm i$. We can then compute the residue at i using (c) — we have

$$(z-i)^m f(z) = \frac{1}{(z+i)^m},$$

and after differentiating $m-1$ times, we have

$$\mathrm{Res}(f; i) = \frac{1}{(m-1)!} \cdot \frac{(-m)(-m-1)\cdots(-m-(m-2))}{(2i)^{2m-1}} = \frac{1}{2i} \cdot \frac{1}{2^{m-2}} \binom{2m-2}{m-1}.$$

§11.4 Cauchy residue theorem

Why do we care about computing residues? The reason is the Cauchy residue theorem.

Theorem 11.10 (Cauchy residue theorem)

Let D be convex, and let $f \in \mathcal{H}(D \setminus \{z_1, \dots, z_r\})$ (where $z_1, \dots, z_r$ are distinct). Let γ be a counter-clockwise closed 'reasonable' loop not passing through any of the poles. Then

$$\int_\gamma f(z)\, dz = 2\pi i \sum \mathrm{Res}(f, z_j),$$

where the sum is over all z_j's inside γ.

By 'reasonable' we mean any closed loop for which the proof we'll give applies (e.g. a circle, rectangle, semicircle, and so on).

So we take our closed loop γ, and there'll be some poles inside and some poles outside; the integral just sees the poles that are inside and ignores the poles that are outside. This is a generalization of the Cauchy integral formula (where there are no poles, so the right-hand side is zero).

Remark 11.11. This can be generalized to D being simply connected, but we won't do so in this class. We're avoiding this word because 18.901 is not a prerequisite for this class. Some of these theorems generalize, but their interestingness doesn't really increase with the complicatedness of the region.

Proof. We'll consider the case where γ is a rectangle (for simplicity). Let $p_1(\frac{1}{z-z_1})$, ..., $p_r(\frac{1}{z-z_r})$ be the principal parts of the poles, and let

$$\varphi(z) = f(z) - \sum_{j=1}^r p_j\left(\frac{1}{z-z_j}\right)$$

for $z \in D \setminus \{z_1, \dots, z_r\}$. By the same argument as we used to prove the partial fractions theorem, φ has removable singularities at all the poles, so we can let $\widetilde{\varphi}$ be the function φ with its singularities removed. Then

$$\int_\gamma \varphi = \int_\gamma \widetilde{\varphi} = 0$$

by the Cauchy integral theorem.

But we know φ is the difference of f and its principal parts, so $\int_\gamma f$ is equal to the sum of integrals of its principal parts; and the point is that we can compute these integrals, since the p's are polynomials. We have

$$\int_\gamma f = \sum_{j=1}^{r} \int_\gamma p_j\left(\frac{1}{z-z_j}\right)\, dz.$$

Now it suffices to look at each of these individually; we will show that for any pole z_0 with principal part

$$p\left(\frac{1}{z-z_0}\right) = \frac{c_{-m}}{(z-z_0)^m} + \cdots + \frac{c_{-1}}{(z-z_0)},$$

we have that

$$\int_\gamma p\left(\frac{1}{z-z_0}\right)\, dz = \begin{cases} 2\pi i c_{-1} & \text{if } z_0 \text{ is inside } \gamma \\ 0 & \text{if } z_0 \text{ is outside } \gamma. \end{cases}$$

First suppose that z_0 is outside γ. Then we can find a slightly larger convex region containing γ but not z_0, and our integral is 0 by the Cauchy integral theorem (since this function is holomorphic in any region that doesn't contain z_0).

Now suppose that z_0 is inside γ. Then we do the standard trick — let c be a small circle centered around z_0. Then we can reduce the integral to one around c by adding in cancelling paths; so

$$\int p\left(\frac{1}{z-z_0}\right)\, dz = \int_c p\left(\frac{1}{z-z_0}\right)\, dz.$$

Finally, we will show that

$$\int_c \frac{dz}{(z-z_0)^\ell} = \begin{cases} 2\pi i & \text{if } \ell = 1 \\ 0 & \text{if } \ell > 1 \end{cases}.$$

To see the latter case, note that $(z-z_0)^{-\ell}$ has an antiderivative if $\ell > 1$ — we have

$$(z-z_0)^{-\ell} = \left(\frac{1}{-\ell+1}(z-z_0)^{-\ell+1}\right)'.$$

And any function with an antiderivative has integral 0 around any closed path, so the integral is 0.

Meanwhile if $\ell = 1$, we computed the integral last class — so we're done. □

Remark 11.12. What's an example of a path for which this doesn't work? The part that potentially breaks is reducing the integral around γ to an integral around the circle. Any type of path we can draw on the board will be suitable; you might have to worry about paths that weirdly self-intersect so that there isn't a circle like this. But such an example has to be pretty pathological.

In particular, an issue in general is the inside vs. outside of the path — there's some hypothesis about paths in which you can give a well-defined inside and outside, and then the proof should work. (You'll probably need to chop up the path in more complicated ways, but this should be possible using compactness.)

§11.5 Some applications

Here's one example: let γ be the semicircle of radius R centered at 0 (in the upper half-plane). Then we have

$$\int_\gamma \frac{dz}{z^2+1} = 2\pi i \operatorname{Res}\left(\frac{1}{z^2+1}, i\right) = \frac{2\pi i}{2i} = \pi.$$

But we can split up the integral into two parts — this gives us

$$\int_{-R}^{R} \frac{dx}{x^2+1} = \pi - \int \frac{dz}{z^2+1},$$

where the second integral is along just the arc of the semicircle. But the integral along the arc goes to 0 — the length of the arc is bounded by πR, and the maximum the function takes on the arc is $R^2 - 1$. So the integral is at most

$$\frac{\pi R}{R^2 - 1} \to 0$$

as $R \to \infty$. This means

$$\int_{-\infty}^{\infty} \frac{dx}{x^2+1} = \pi.$$

This is a lot easier than doing the integral in calculus.

And we can do more complicated examples using the same logic — we can compute any integral of a rational function on the real axis, just by using the residue theorem (with no more calculus).

Remark 11.13. Earlier, we mentioned how the pole at i also ruins the convergence of the power series of this function.

We can generalize this example.

Example 11.14

Let f be any rational function with no poles on $\mathbb{R}$, and suppose that $f = g/h$ with $\deg(h) \geq \deg(g) + 2$. Then we have

$$\int_{-\infty}^{\infty} f(x)\, dx = 2\pi i \sum \operatorname{Res}(f, z_j)$$

where z_j are the poles of f above $\mathbb{R}$.

(We can do the same with the poles below R, but we'll get a minus sign.)

Proof. The proof is the same as in the example — take the same contour (the semicircle with radius R). Then we have

$$\int_{-R}^{R} f(x)\, dx = 2\pi i \sum \operatorname{Res}(f, z_i) - \int \frac{g(z)}{h(z)}\, dz$$

where the integral is over the half-circle arc (we just need to take R large enough to capture all the poles above $\mathbb{R}$; but there's finitely many poles, so this is certainly possible). But the integral overthe arc goes to 0, by the same argument — for R large g and h are both asymptotically governed by their leading term, so

$$f(z) \approx \frac{c_k z^k}{d_\ell z^\ell} \sim z^{k-\ell}$$

(for $|z|$ large). If $\ell \geq k+2$, then this is at most $|z|^{-2}$ (which is R^{-2}). And the length of the arc contributes R; so

$$\left| \int_\gamma \frac{g(z)}{h(z)}\, dz \right| \leq \frac{CR}{R^2} \to 0$$

as $R \to \infty$. ☐

§12 October 24, 2023

Last time we saw the Cauchy residue theorem.

Theorem 12.1

If $D \subseteq \mathbb{C}$ is convex and $f \in \mathcal{H}(D \setminus \{z_1, \ldots, z_r\})$ where the z_j are all poles of f, then for any 'sufficiently nice' counterclockwise closed loop γ not passing through any of the poles, we have

$$\int_\gamma f(z)\,dz = 2\pi i \sum \operatorname{Res}(f; z_j)$$

where the sum is over all poles z_j *inside* γ.

So we don't count the poles on the outside, but we do count the poles on the inside.

We started applying this by computing some integrals using the residue theorem; we'll now continue this.

Last time we considered cases $\int_{-\infty}^{\infty} \frac{h(x)}{g(x)}\,dx$ where h and g are polynomials (with no common roots) and $f(z) = \frac{h(z)}{g(z)}$ has no real poles (i.e., g has no real roots); we were able to evaluate such integrals as long as $\deg(g) \geq \deg(h) + 2$. We saw that such an integral is $2\pi i$ times the sum of the residues at the poles above $\mathbb{R}$.

§12.1 Cauchy principal values

Now we'll consider the case where f does have real poles. Then $\int_{-\infty}^{\infty} f(x)\,dx$ is improper in two ways — it's improper because we're taking $x \to \infty$, but it's also improper because f has poles on the real axis. We can still try to assign this integral a number, which we call the *principal value*.

Definition 12.2. Let $x_1 < x_2 < \cdots < x_s$ be the poles of f on $\mathbb{R}$. We define the *Cauchy principal value* of $\int_{-\infty}^{\infty} f(x)\,dx$ as

$$⨍_{-\infty}^{\infty} f(x)\,dx = \lim_{\substack{R\to\infty \\ \delta\to 0}} \left(\int_{-R}^{x_1-\delta} f(x)\,dx + \int_{x_1+\delta}^{x_2-\delta} f(x)\,dx + \cdots + \int_{x_s+\delta}^{R} f(x)\,dx \right),$$

if the limit exists.

Here f isn't defined on these points $x_1, \ldots, x_s$; so we first compute the integral from R to $x_1 - \delta$, then $x_1 + \delta$ to $x_2 - \delta$, and so on.

This integral might not be finite, or might be difficult to compute; the integral doesn't have to exist.

Remark 12.3. Can we define this for a function with infinitely many poles, e.g. $\frac{1}{\sin x}$? Yes, if the poles are discrete.

Now we'll consider the example of rational functions (similarly to before).

Theorem 12.4

Let $f = g/h$ be a rational function with $\deg(h) \geq \deg(g) + 2$, and such that all its real poles are simple. Then the Cauchy principal value $⨍_{-\infty}^{\infty} f(x)\,dx$ exists and is equal to

$$2\pi i \left(\sum \text{Res}(f, z_j) + \frac{1}{2}\sum \text{Res}(f, x_j)\right)$$

where the first term sums over poles *above* $\mathbb{R}$, and the second term sums over poles *on* $\mathbb{R}$.

So as before, we pick up a term for all the poles above $\mathbb{R}$; we also pick up a term for all the poles on $\mathbb{R}$, but we get a coefficient of $\frac{1}{2}$.

Proof. We'll try to use the same contour. But instead of taking our contour *through* the poles, we'll hop over them with little half-circles, of radius δ.

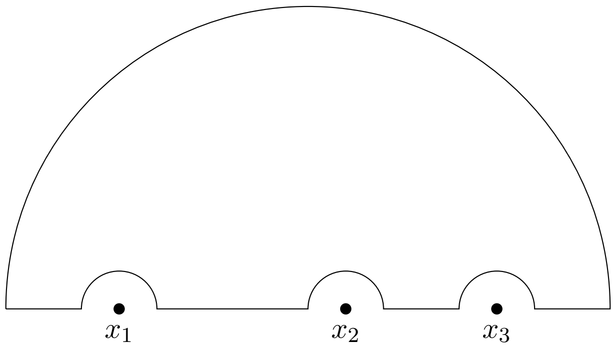

By the residue theorem, we have

$$2\pi i \sum_{z_j \text{ poles over } \mathbb{R}} \text{Res}(f, z_j) = \int_\gamma f(z)\,dz.$$

But we can break γ into several pieces — as

$$\int_\gamma f(z)\,dz = \int_{-R}^{x_1 - \delta} + \cdots + \int_{x_s + \delta}^{R} + \left(\int_{\gamma_1} f + \cdots + \int_{\gamma_s} f\right) + \int_{\gamma_R} f,$$

where γ_i is the little semicircle at x_i and the final integral is the integral over the large half-circle.

By the same argument as last class, the integral over the large half-circle goes to 0 as $R \to \infty$. The first term is the principal value (as $R \to \infty$ and $\delta \to 0$). So now we have to figure out the integrals over the little half-circles; the idea is that since we're integrating over half a circle, we get half the residue.

Lemma 12.5

Let $f \in \mathcal{H}(D \setminus \{z_0\})$ where z_0 is a simple pole, and let C_δ be a counterclockwise half-circle centered at z_0 of radius δ. Then $\int_{C_\delta} f(z)\,dz \to \pi i \text{Res}(f, z_0)$ as $\delta \to 0$.

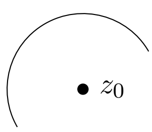

Note that C_δ is not closed; it's just half a circle.

Proof. Look at a small disk around z_0 of radius r; then in $0 < |z - z_0| \leq r$ we can write

$$f(z) = \frac{c_{-1}}{z - z_0} + h(z)$$

where $h(z)$ is holomorphic in D. Since h is holomorphic, it'll be bounded on this disk; so $\int_{C_\delta} h(z)\,dz \leq |C_\delta| \max_{|z-z_0|\leq r} |h(z)| \to 0$ as $\delta \to 0$. So h doesn't contribute anything, and we're left to compute the integral of $\frac{1}{z-z_0}$. But this is now an explicit computation — we can parametrize C_δ as $z(\theta) = z_0 + \delta e^{i\theta}$ for $\theta \in [\theta_0, \theta_0 + \pi]$. Then we can simply compute the integral. We have $z'(\theta) = i\delta e^{i\theta} = i(z(\theta) - z_0)$, so this integral is just equal to

$$c_{-1} \int_{\theta_0}^{\theta_0+\pi} \frac{i(z(\theta) - z_0)}{z(\theta) - z_0}\, d\theta = c_{-1}\pi i = \pi i \text{Res}(f; z_0). \qquad \square$$

Once we know this, we've now computed each of the limits of the $\int_{\gamma_j} f$ — as we let $\delta \to 0$, these tend to $-\pi i \sum \text{Res}(f; x_j)$ (since the half-circles are clockwise). And now we're done — taking the limit as $R \to \infty$ and $\delta \to 0$, we have

$$2\pi i \sum_{\text{poles above } \mathbb{R}} \text{Res}(f, z_j) = ⨍_{-\infty}^{\infty} f(z)\, dz - \pi i \sum \text{Res}(f, x_j) + 0,$$

which rearranges to the desired formula. $\square$

It's easy to remember because we pick up all the residues above $\mathbb{R}$, and since we have half-circles we get half of the residues on $\mathbb{R}$.

Remark 12.6. We could also have captured the poles below $\mathbb{R}$ instead, by taking the downwards semicircle; this would involve some more negative signs. In practice, if there's 10 poles above $\mathbb{R}$ and one below $\mathbb{R}$, you may want to do the contour below $\mathbb{R}$ so you only have to compute the residue of one pole.

Is there intuition for why these integrals spit out the same value? Not sure. If f is real on $\mathbb{R}$, then by the reflection principle we can see this symmetry; but if it's not then this is less obvious (what we've done doesn't require our rational functions to have real coefficients).

Remark 12.7. Can the condition that f is rational be replaced with an asymptotic condition? Yes, using the same proof, as long as f decays enough on the large semicircle (e.g. as $|z|^{-2}$, or even $|z|^{-(1+\varepsilon)}$ or even $(|z| \log |z|)^{-1}$).

§12.2 Some examples

Example 12.8

Compute $⨍_{-\infty}^{\infty} \frac{dx}{x^3-1}$.

Solution. The function $f(z) = \frac{1}{z^3-1}$ has poles at the cube roots of unity — i.e., 1, ζ, and ζ^2 where $\zeta = e^{2\pi i/3} = -\frac{1}{2} + \frac{\sqrt{3}}{2}$. The degree of the denominator is 3 and the degree of the numerator is 0, so we can just compute this using the theorem — it's equal to

$$2\pi i \text{Res}(f, \zeta) + \pi i \text{Res}(f, 1).$$

And now we can compute these residues. We can write

$$f(z) = \frac{1}{(z-1)(z-\zeta)(z-\zeta^2)} = \frac{\frac{1}{(z-1)(z-\zeta)^2}}{z-\zeta}.$$

Then we have

$$\operatorname{Res}(f;\zeta) = \lim_{z\to\zeta} \frac{1}{(z-1)(z-\zeta^2)} = \frac{1}{(\zeta-1)(\zeta-\zeta^2)} = \frac{1}{-\zeta(1-\zeta)^2}$$

We have $\zeta^2+\zeta+1=0$, so $(1-\zeta)^2 = \zeta^2 - 2\zeta + 1 = -3\zeta$, and this is equal to

$$\operatorname{Res}(f;\zeta) = \frac{1}{3\zeta^2}.$$

Similarly we can compute

$$\operatorname{Res}(f;1) = \frac{1}{3}.$$

So the answer ends up being

$$2\pi i\left(\frac{1}{3\zeta^2} + \frac{1}{2}\cdot\frac{1}{3}\right) = -\frac{\pi}{\sqrt{3}}.$$ □

Here's another cool example (not about principal values).

Example 12.9

Compute $\int_0^\infty \frac{dx}{x^3+1}$.

Solution. Consider $f(z) = \frac{1}{z^3+1}$. We have $z^6 - 1 = (z^3+1)(z^3-1)$, so we can write

$$f(z) = \frac{1}{z^3+1} = \frac{z^3-1}{z^6-1},$$

and therefore the poles of f are the 6th roots of unity that are not also 3rd roots of unity.

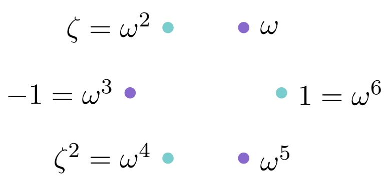

(The poles are the points in purple.)

We're interested in the integral from 0 to ∞; we take the following contour.

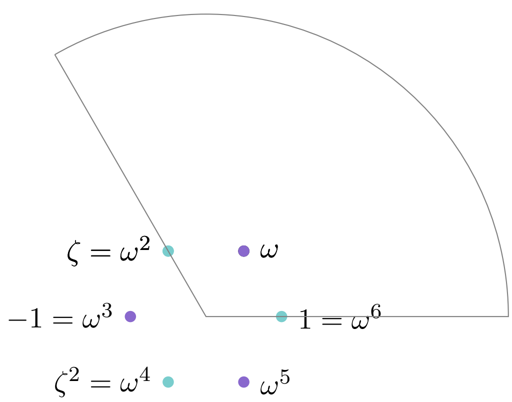

This contour traps only one pole, so we have

$$2\pi i \text{Res}(f, \omega) = \int_\gamma f(z)\,dz = \int_0^\infty f(x)\,dx + \int f(z)\,dz + \int f(z)\,dz$$

where the first integral is over the arc, and the second over the diagonal line.

The integral over the arc goes to 0 as $R \to \infty$, because the denominator has degree 3 and the numerator has degree 0 (so the function decays much faster than the length of the arc).

On the other hand, consider the integral

$$\int_{[0,\omega^2 R]} f(z)\,dz.$$

This path is a straight line segment, and we can parametrize it by $z(x) = x\omega^2$ for $x \in [0, R]$. Then $z'(x) = \omega^2$, and we get

$$\int_0^R \frac{\omega^2 x}{(x\omega^2)^3 + 1} = \omega^2 \int_0^R \frac{dx}{x^3+1}.$$

So we actually get another piece of the thing we were trying to compute! And we're left with

$$2\pi i \text{Res}(f;\omega) = (1-\omega^2)\int_0^\infty \frac{dx}{x^3+1}.$$

(the minus sign comes from the fact that the segment has the opposite orientation). Finally, it remains to compute this residue. We can do this simply by taking the derivative of the denominator, so this is $\frac{2\pi i}{3\omega^2}$. Now we have

$$\int_0^\infty \frac{dx}{x^3+1} = \frac{2\pi i}{3\omega^2(1-\omega^2)} = \frac{2\pi i}{3(\omega^2-\omega^4)} = \frac{2\pi i}{3(\zeta - \zeta^2)} = \frac{2\pi i}{3\sqrt{3}i} = \frac{2\pi}{3\sqrt{3}}.$$ □

§13 October 31, 2023

We'll discuss another family of applications to the residue theorem — computing Fourier transforms.

§13.1 Fourier transforms

Definition 13.1. Given a continuous function $f\colon \mathbb{R} \to \mathbb{C}$ with $\int_{-\infty}^\infty |f| < \infty$, its *Fourier transform* $\hat{f}\colon \mathbb{R} \to \mathbb{C}$ is defined as

$$\hat{f}(t) = \int_{-\infty}^\infty f(x)e^{2\pi i t x}\,dx.$$

There are various different conventions for the Fourier transform; we'll use the same convention as Gallagher (who was an analytic number theorist).

There's a lot of places Fourier transforms comes up; it's essential in PDEs, and also comes up in number theory and algebraic geometry. It takes a function and spits out another function; it's usefulness in PDEs is that it turns derivatives into multiplication — we have

$$\hat{f'}(t) = \pm 2\pi i t \hat{f}(t).$$

The hard thing about differential equations is that there's derivatives; and if you take the Fourier transforms then the derivative disappears and you get multiplication by t, which is much easier to deal with. (This holds for sufficiently nice functions.)

For example, you can use this to solve any constant-coefficient linear differential equation. (You will see this in 18.103.) Fourier introduced this procedure to solve the heat equation. This is a continuous version of the Fourier series — if $f\colon S^1 \to \mathbb{C}$, then you can reconstruct f via a series $f(x) = \sum_{n=-\infty}^{\infty} \hat{f}(n)e^{-inx}$, where $\hat{f}(n) = \int_{-\pi}^{\pi} f(n)e^{in\theta}\, d\theta$. The point is that these $\hat{f}(n)$'s are just numbers, and this gives a series; and you can reconstruct the whole function f using these series (if f is sufficiently nice). This is nice for many applications — it says essentially that you can write a periodic function as a superposition of sines and cosines. Here t is an analog of n, and there's an inversion formula — you can recover f according to its Fourier transform, as

$$f(x) = \int_{-\infty}^{\infty} \hat{f}(t)e^{-2\pi itx}\, dt.$$

So the Fourier series gives an isometry between L^2 and ℓ^2, and the Fourier transform an isometry between L^2 and L^2 (where L^2 are the square-integrable functions, meaning $\int |f|^2\, dx < \infty$). This is a very condensed version, but this is an important object — it has applications to a bunch of different fields in math, as well as the sciences and engineering.

So the Fourier transform is an important object in math; then there's the question of, given a function, how do we compute its Fourier transform? This isn't super clear, because we're taking f, multiplying it by some oscillatory function $e^{2\pi ixt}$ (which winds around the circle), and trying to integrate.

But now we have contour integration and the residue theorem, and we can use this to explicitly compute some Fourier transforms.

§13.2 Computing Fourier transforms

Example 13.2 (1) For $f(x) = \frac{1}{1+x^2}$ we have $\hat{f}(t) = \pi e^{-2\pi|t|}$.

(2) For $f(x) = e^{-\pi x^2}$ we have $\hat{f}(t) = e^{-\pi t^2}$.

(3) For $f(x) = \frac{1}{\cosh(\pi x)}$ we have $\hat{f}(t) = \frac{1}{\cosh(\pi t)}$.

We'll do these three computations. Embedded here is a general rule of Fourier transforms — the decay of the function leads to smoothness of the Fourier transform, and smoothness of the function leads to decay of the Fourier transform. For example $\frac{1}{1+x^2}$ is nicely smooth, so its Fourier transform is quickly decaying.

§13.3 First example

Proof of (1). We want to compute the Fourier transform of $\frac{1}{1+x^2}$ — we want to show that

$$\int_{-\infty}^{\infty} \frac{e^{2\pi itx}}{1+x^2}\, dx = \pi e^{-2\pi|t|}.$$

Consider the function

$$f(z) = \frac{e^{2\pi izt}}{1+z^2}.$$

We want to know its integral on $\mathbb{R}$. First suppose t is positive. This function has poles at $\pm i$; so we take the usual contour consisting of a semicircle with radius R, where $R \to \infty$ (consisting of $[-R, R]$ and the arc above the x-axis). By the residue theorem, we know

$$\int_{-R}^{R} \frac{e^{2\pi ixt}}{1+x^2}\, dx + \int_{\text{arc}} \frac{e^{2\pi izt}}{1+z^2}\, dt = 2\pi i \text{Res}\left(\frac{e^{2\pi izt}}{1+z^2}; i\right) = 2\pi i \cdot \frac{e^{-2\pi t}}{2i} = \pi e^{-2\pi t}.$$

Now we want to show that the integral on the arc goes to 0. For $t > 0$, we have $\left|e^{2\pi itz}\right| = e^{-2\pi yt}$, where $z = x + iy$; and since $y \geq 0$, this is at most 1. So then we can remove it; we can bound the numerator by 1, and the denominator is R^2, and the length of the path is R, so this integral goes to 0 as $R \to \infty$.

We'd be in trouble if $t < 0$, because we can't upper bound this function anymore. (Our upper bound would be growing in y, which can be R.) But we've proven the formula for $t \geq 0$; for $t \leq 0$ we can instead take a circle below the real axis, and do the same argument (because we want y to be negative). Now we have a pole at $-i$. We also need to go counterclockwise, so we get minus the integral; but that takes care of it. □

Before we do (2), here's a few generalizations of (1).

Theorem 13.3

For each rational function f with no real poles and such that $f(\infty) = 0$ (i.e., the degree of the denominator is bigger than the degree of the numerator), for $t > 0$ we have

$$\lim_{R\to\infty} \int_{-R}^{R} f(x)e^{2\pi itx}\, dx = 2\pi i \sum_{\text{poles above } \mathbb{R}} \operatorname{Res}(f(z)e^{2\pi i z_j t}, z_j).$$

We're going to need to do some work here, since $f(\infty) = 0$ is a weaker assumption than what we used above. We're going to try to exploit extra decay coming from the $e^{2\pi ixt}$ factor. (We've covered the case where the degree of the denominator is at least 2 more than that of the numerator, but this allows even 1 more.)

Remark 13.4. This holds even if f itself doesn't have finite integral on $\mathbb{R}$.

Proof. The residue theorem gives us that for $R > \max |z_j|$ (where z_j are the poles over $\mathbb{R}$), we have

$$\int_{-R}^{R} f(x)e^{2\pi ixt}\, dx + \int_{\text{arc}} f(z)e^{2\pi izt}\, dz = 2\pi i \sum \operatorname{Res}(f; z_j).$$

So it's enough to prove that (for $t > 0$) this integral goes to 0 as $R \to \infty$. And we've seen that just bounding the integral by the sup of the integrand times the length of the path is insufficient; so we have to be more careful than that.

Claim 13.5 — For every $t > 0$ w ehave $\int_{\text{arc}} f(z)e^{2\pi izt}\, dz \to 0$ as $R \to \infty$.

Proof. The fact that $f(\infty) = 0$ means that the degree of the denominator is at least that of the numerator plus one, so $|zf(z)|$ is bounded — there exists $B > 0$ such that $|zf(z)| \leq b$ for all $z \in \mathbb{C}$. We don't want to put absolute values inside the integral, so we'll need to look at what the integral is by parametrizing the integral — we can parametrize the path as $z(\theta) = Re^{i\theta}$ for $\theta \in [0, \pi]$. Then $z'(\theta) = iRe^{i\theta}$. So we have

$$\int_{\text{arc}} f(z)e^{2\pi izt} = \int_0^{\pi} f(Re^{i\theta})e^{2\pi iRe^{i\theta}t} \cdot iRe^{i\theta}\, d\theta.$$

Now we have $Re^{i\theta \cdot F(Re^{i\theta})} = zf(z)$. Now at *this* point we put absolute values in; we have

$$B \int_0^{\pi} e^{-2\pi t \sin\theta}\, d\theta.$$

(We put absolute values inside the integral —

$$\left| \int_0^{\pi} F(-)e^{2\pi iRe^{i\theta}t} iRe^{i\theta} \right| \leq \int_0^{\pi} \left| iRe^{i\theta}F(Re^{i\theta}) \right| \left| e^{2\pi i(\cos\theta + i\sin\theta)Rt} \right| \leq B \int_0^{\pi} e^{-2\pi Rt\sin\theta}$$

(separating the exponent into its real and imaginary parts — only the real part affects the magnitude). Now t is fixed; and we need to prove that

$$\int_0^{\pi} e^{-2\pi Rt\sin\theta}\, d\theta \to 0.$$

Lemma 13.6 (Jordan's lemma)

For $t > 0$, we have $\int_0^\pi e^{-2\pi Rt\sin\theta}\, d\theta \to 0$ as $R \to \infty$.

Proof. We have $\sin\theta \geq 2\theta/\pi$ for $0 \leq \theta \leq \pi/2$. (We can see this by graphing.) So we can lower-bound sin in this way; and then the integral from 0 to π is the same as twice the integral from 0 to $\pi/2$. So our integral is

$$2\int_0^{\pi/2} e^{-2\pi Rt\sin\theta}\, d\theta \leq 2\int_0^{\pi/2} e^{-4Rt\theta}\, d\theta \leq 2\int_0^\infty e^{-4Rt\theta}\, d\theta$$

(since we're making the numerator less negative). But now we can just integrate it, and it's equal to $1/2Rt$, which goes to 0 as $R \to \infty$ (for each fixed $t > 0$). □

□

This handles the case of f not decaying too fast. □

We can further generalize, and even allow for poles on the real axis as long as they're simple.

Theorem 13.7

For each rational f with no multiple real poles with $f(\infty) = 0$, we have

$$⨍_{-\infty}^{\infty} f(x)^{2\pi itx}\, dx = 2\pi i\left(\sum_{\text{poles above } \mathbb{R}} \operatorname{Res}(f(z)e^{2\pi izt}; z_j) + \frac{1}{2}\sum_{\text{poles on } \mathbb{R}} \operatorname{Res}(f(z)e^{2\pi izt}, z_j)\right).$$

(Since we have real poles, we take the Cauchy principal value.)

The right-hand side comes from the residue theorem for the case hwere we have simple poles on $\mathbb{R}$ (which is what we did last class); the fact that the integral over the arc goes to 0 is the same as what we just did.

§13.4 Second example

Example 13.8

We have

$$⨍_{-\infty}^{\infty} \frac{e^{2\pi ixt}}{x}\, dx = 2\pi i\left(\frac{1}{2}\operatorname{Res}(e^{2\pi izt}/z, 0)\right) = \pi i.$$

But now we can split this into two pieces based on the real and imaginary part — we have

$$⨍_{-\infty}^{\infty} \frac{\cos(2\pi xt)}{x}\, dx = 0 \text{ and } ⨍_{-\infty}^{\infty} \frac{\sin(2\pi xt)}{x}\, dx = \pi$$

for $t > 0$. The first is unsurprising because cos is even and $1/x$ is odd; and by the definition of the Cauchy principal value (which is symmetric). For the second, we can change variables nad this says that

$$\int_{-R}^{R} \frac{\sin\theta}{\theta}\, d\theta \to \pi,$$

or in other words $\int_0^\infty \frac{\sin\theta}{\theta} = \frac{\pi}{2}$. That's neat — the oscillations of sin are enough for this thing to converge (kind of like the alternating series test).

This gives the Fourier transform when $t > 0$, and for $t < 0$ we have $\hat{f}(-t) = -\overline{\hat{f}}(t)$. If f is real then you can just remove the bars.

Then you have $t = 0$, which is different. There we want to know the Fourier transform for $t = 0$; it's possibethat the integral fo f converges absolutely or not.

§13.5 Second example

Now we'll look at (2), which involves Prof. Lawrie's favorite trick he learned as an undergrad. At some point in the proof, we'll need the following computation.

Claim 13.9 — We have $\int_{-\infty}^{\infty} e^{-\pi x^2}\,dx = 1$.

This is in every book on probability. It's hard because you don't have an antiderivative, but you can do it by an amazing trick.

Proof. Instead of trying to compute this integral, we'll try to compute its square — and using the freedom to rename our integration variables, we can rewrite this square as

$$\left(\int_{-\infty}^{\infty} e^{-\pi x^2}\,dx\right)^2 = \left(\int_{-\infty}^{\infty} e^{-\pi x^2}\,dx\right)\left(\int_{-\infty}^{\infty} e^{-\pi y^2}\,dy\right).$$

But now this looks like a double integral that happens to separate into two integrals, so we can rewrite it as

$$\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{-\pi(x^2+y^2)}\,dx\,dy.$$

And now this looks like polar coordinates — for $(x, y) \in \mathbb{R}$ we can pass to (r, θ) coordinates, where $r^2 = x^2 + y^2$. Then the volume form $dx\,dy$ becomes $r\,dr\,d\theta$. So in polar coordinates, this is equal to

$$\int_0^{2\pi}\int_0^{\infty} e^{-\pi r^2} r\,dr\,d\theta = 2\pi \int_0^{\infty} e^{-\pi r^2} r\,dr$$

(since the integrand doesn't depend on θ). Now we can use the chain rule — we can do u-substitution with $u = \pi r^2$ and $du = 2\pi r\,dr$, so this is equal to

$$\int_0^{\infty} e^{-u}\,du = 1.$$

(This is amazing — in calculus if you can't find an antiderivative, or the function has no closed-form antiderivative, what on earth are you supposed to do with it? But somehow this works.) □

Remark 13.10. We only really showed the square of the integral is 1, but it is clearly positive.

Now we can compute the Fourier transform. (We'll see later where this formula comes in.)

Claim 13.11 — We have $\int_{-\infty}^{\infty} e^{2\pi i x t} e^{-\pi x^2}\,dx = e^{-\pi t^2}$.

Proof. This is kind of asking for us to complete the square — we can write

$$-x^2 + 2ixt = -(x - it)^2,$$

so then this becomes

$$-e^{-\pi t^2}\int_{-\infty}^{\infty} e^{-\pi(x-it)^2}\,dx.$$

If we believe the answer, we'd like to show this integarl is equal to 1. And this now resembles the earlier computation, but shifted by it. That tells us which contour we should pick (because we'd like to turn it into this integral that we know) — we want to show that $\int_{-\infty}^{\infty} e^{-\pi(x-it)^2}\,dx = \int_{-\infty}^{\infty} e^{-\pi x^2}\,dx$. To do this, we integrate the function over a rectangle with corners R, $-R$, $R-it$, and $-R-it$. The function $e^{-\pi z^2}$ is holomorphic in this rectangle, so by the Cauchy integral theorem its integral around the rectangle is zero. This tells us

$$\int_{[-R-it,R-it]} + \int_{[R-it,R]} - \int_{[-R,R]} + \int_{[-R,-R-it]} e^{-\pi z^2}\,dz = 0.$$

The first integral isi $\int_{-R}^{R} e^{-\pi(x-it)^2}\,dx$, since the path is parametrized by $z(x) = x - it$ for $x \in [-R, R]$, where $z'(x) = 1$. So we get

$$\int_{-R}^{R} e^{-\pi(x-it)^2} - \int_{-R}^{R} e^{-\pi x^2}\,dx = i\int_{0}^{-t} e^{-\pi(-R+iy)^2}\,dy - i\int_{-t}^{0} e^{-\pi(R-iy)^2}\,dy$$

(by parametrizing each of the integrals).

For $y \in [-t, 0]$, we have $\left|e^{-\pi(\pm R+iy)^2}\right| = e^{-\pi(R^2-y^2)}$, because we only look at the real part. This is at most $e^{-\pi(R^2-t^2)}$. And so then both of these vertical integrals go to 0 as $R \to \infty$ (and t is fixed) — the length of the integral is $|t|$. ∎

So the Fourier transform of the Gaussian function is the Gaussian functionb back again. (3) is done in the notes by a similar argument (its Fourier transform is also itself).

§14 November 2, 2023

§14.1 Counting zeros

Another cool application of the theory we've developed is counting zeros of functions.

Theorem 14.1

Let D be a convex region, let f be holomorphic in D and not identically zero, and let C be a circle in D not passing through any zeros of f. Let $n(f, C)$ be the number of zeros of f inside C, counted with multiplicity. Then $n(f, C)$ is given by

$$n(f, C) = \frac{1}{2\pi i}\int_C \frac{f'(z)}{f(z)}\,dz.$$

(If f has a zero with multiplicity 2, then it's counted twice.)

This is a remarkable formula. In particular, this expression is an integer. (This looks like the log derivative of f.)

Proof. We can replace D with a sub-disk containing C so that f only has finitely many zeros in D. (We'll call the new region D as well.)

Claim 14.2 — The zeros of f with multiplicity m are precisely the simple poles of $\frac{f'}{f}$, and have residue m.

Proof. To see this, near a zero z_0 we can write $f(z) = (z - z_0)^m g(z)$ where g is holomorphic and $g(z_0) \neq 0$. Now we can compute $\frac{f'}{f}$ using this representation of f — we have $f'(z) = m(z-z_0)^{m-1}g(z)+(z-z_0)^m g'(z)$, so then

$$\frac{f'(z)}{f(z)} = \frac{m}{z - z_0} + \frac{g'(z)}{g(z)}.$$

But now $\frac{m}{z-z_0}$ is the principal part of the pole — this is because $\frac{g'(z)}{g(z)}$ is holomorphic (at least, in a small neighborhood around around z_0 — we know $g(z_0) \neq 0$, so g is nonzero in a small disk around z_0, and since g is holomorphic so is g'), and this is exactly how we defined the principal part of a pole.

This means that z_0 is a simple pole of $\frac{f'}{f}$ with residue m. □

Then we have

$$\int_C \frac{f'(z)}{f(z)}\,dz = 2\pi i \sum_{z_j \text{ poles inside } C} \operatorname{Res}\left(\frac{f'}{f}, z_j\right) = 2\pi i \sum m_j$$

by the Cauchy residue theorem, where the second sum is over the distinct zeros of f inside C (and m_j are their multiplicities). This completes the proof. □

Let's now have some fun with this.

Theorem 14.3 (Glicksberg 1976)

Let f and g be holomorphic in a convex region D, and let C be a circle in D. Suppose that for all $z \in C$ we have

$$|f(z) - g(z)| < |f(z)| + |g(z)|\,. \tag{$\star$}$$

Then f and g have the same number of zeros inside C.

This allows us to count the zeros of f by comparing it to another function that we can easily count the zeros of — if we can check this inequality and we know how many zeros one has, then we know how many zeros the other has.

Proof. The assumption says that we don't have equality in the triangle inequality; what does this mean? The weaker inequality $|f(z) - g(z)| \leq |f(z)|+|g(z)|$ is *always* true — this is the triangle inequality. Equality happens when $f(z)$, 0, and $g(z)$ lie on a line (in that order) — i.e., equality holds if and only if $0 \in [f(z), g(z)]$. So the statement ($\star$) is equivalent to the statement that $0 \notin [f(z), g(z)]$ for every $z \in C$.

Now for $z \in D$ and $t \in [0, 1]$, set

$$h_t(z) = (1 - t)f(z) + tg(z).$$

For each fixed t this is a holomorphic function in D. And this parametrizes the line segment between f and g; our hypothesis says that h_t has no zeros on C for any t. (We're trying to apply our original theorem to h_t.) In particular h_t is not identically 0.

So the number of zeros of h_t in C is

$$n(h_t, C) = \frac{1}{2\pi i}\int \frac{h_t'(z)}{h_t(z)}\,dz$$

by the theorem. (This number exists and is well-defined by our theorem.)

We'll now use a *continuity argument* — the theorem implies that this integral is always an integer. And when $t = 0$, this spits out the number of zeros of f — i.e., $n(h_0, C) = n(f, C)$ — and similarly $n(h_1, C) = n(g, C)$.

But this is a continuous function (of t) that's always integer-valued, so it can't jump from one integer to another — and this means it has to be constant. So these two numbers are the same.

(The expression $\frac{h_t'(z)}{h_t(z)}$ is manifestly continuous in t, so the integral is as well.) $\square$

As a corollary, we can prove Rouche's theorem.

Theorem 14.4 (Rouche 1862)

If f and φ are holomorphic in a convex region D and C is a circle in D on which $|\varphi(z)| < |f(z)|$ for all $z \in C$, then f and $f + \varphi$ have the same number of zeros inside C.

So we have a circle on which φ is dominated by f.

Proof. Take $g = f + \varphi$ in Glicksberg's theorem. Then f and φ are holomorphic, so g is too, and the given condition becomes that $|f(z) - g(z)| = |\varphi(z)| < |f(z)| \leq |f(z)| + |g(z)|$ for $z \in C$. So this verifies the conditions of Glicksberg's theorem, and therefore f and $g = f + \varphi$ have the same number of zeros. $\square$

§14.2 Roots of polynomials

As a corollary, we can give a third (slicker) proof of the fundamental theorem of algebra.

Corollary 14.5 (Fundamental theorem of algebra)

Every polynomial of degree $n \geq 1$ has exactly n roots in $\mathbb{C}$ (counted with multiplicity).

Proof. Let $p(z) = c_n z^n + \cdots + c_1 z + c_0$. Take $f(z) = c_n z^n$ and $\varphi(z)$ to be the remainder — i.e., $\varphi(z) = c_{n-1} z^{n-1} + \cdots + c_0$.

To make φ dominated by f, we should take a really big circle — take C to be a circle of radius R about zero. If R is big enough, then $|\varphi(z)| < |f(z)|$ on C. So $f(z)$ and $p(z) = f(z) + \varphi(z)$ have the same number of roots, counted with multiplicity. But f is a scaled version of the function z^n, which has one root of order n at $z = 0$; so this entire thing has n roots (with multiplicity).

This counts the number of roots inside $\{|z| < R\}$, and we can then take R really big to ensure that this captures all the roots. $\square$

We can have a slightly different generalization of this argument.

Corollary 14.6

A polynomial $p(z) = c_n z^n + c_{n-1} z^{n-1} + \cdots + c_0$ of degree $n \geq 1$ has exactly k roots, counted with multiplicity, in the unit disk $\{|z| < 1\}$, if $|c_k| > \sum_{j \neq k} |c_j|$.

So if we take a polynomial and we want to know how many roots it has in the unit disk, if we can find one coefficient bigger than the sum of the rest then we know how many there are.

Proof. We just need to define φ and f for Rouche's theorem. Take the circle $C = \{|z| = 1\}$; then this condition implies that on C we have $\left|\sum_{j \neq k} c_j z^j\right| \leq \sum_{j \neq k} |c_j| < |c_k| = \left|c_k z^k\right|$. So this gives us the condition for Rouche's theorem — and this means $p(z)$ has the same number of roots inside $\{|z| < 1\}$ as $c_k z^k$ (taking $f(z) = c_k z^k$ and φ to be the remainder). $\square$

§14.3 Inverse function theorem

Theorem 14.7

Suppose that $f(z) - w_0$ has a zero of multiplicity m at a point z_0. Then:

(1) For every sufficiently small $\delta > 0$, there exists $\varepsilon > 0$ (depending on δ) such that for every w with $|w - w_0| < \varepsilon$, the function $f(z) - w$ has exactly m zeros in $|z - z_0| < \delta$. Moreover, for $w \neq w_0$, these zeros are all simple.

(2) We can ensure that $\varepsilon \to 0$ as $\delta \to 0$.

(3) (Inverse function theorem) Suppose that $m = 1$. Then for $|w - w_0| < \varepsilon$, the unique zero $z(w)$ of $f(z) - w$ in $|z - z_w| < \delta$ is a holomorphic function on $\{|w - w_0| < \varepsilon\}$.

So $f(z) = w_0$, and it's hit with multiplicity m. (1) tells us that every other point w sufficiently close to w_0 is *also* hit m times by points around z_0 (which are all distinct).

So w_0 is hit m times by the same point z_0. And then if we pick any other w, there are m different points sent to w.

The inverse function theorem is the special case where $m = 1$. We'd like to say something about the inverse function — and in particular, we can conclude that it's a holomorphic function.

Proof of (1). We'll choose our small δ in a specific way, and then pick ε and check that the conclusion is satisfied.

Take δ small enough so that $f(z) - w_0$ has no zeros in the punctured disk $0 < |z - z_0| \leq \delta$ and f' also has no zeros there. We can do this by the isolation of zeros (for both f and f').

Now set $\varepsilon = \min_{|z-z_0|=\delta} |f(z) - w_0|$. (This exists and is positive, since $f(z) - w_0$ is nonzero on this disk.)

Now suppose that $|w - w_0| < \varepsilon$. Then

$$|w - w_0| < \varepsilon \leq |f(z) - w_0|$$

for z with $|z - z_0| = \delta$ (since we defined ε as the minimum of the right-hand side on the circle). Now we can apply Rouche's theorem (with f taken to be $f(z) - w_0$ and $\varphi = w - w_0$, which is (for each w) a constant function in z); this implies that $f(z) - w$ has the same number of zeros in $|z - z_0| < \delta$ as $f(z) - w_0$.

Finally, these zeros are all simple because f' has no zeros. □

This is about counting the number of pre-images — so the number of pre-images is fixed.

Proof of (2). We know $f \in \mathcal{H}(D)$ is continuous, so $f(z) - w_0 \to 0$ as $z \to z_0$. And we defined ε as the minimum; so ε as we defined it here tends to 0. □

Proof of (3). The existence of $z(w)$ is a special case of (1). Now consider

$$\frac{1}{2\pi i}\int_{C_\delta}\frac{zf'(z)}{f(z)-w}\,dz.$$

(Here C_δ is the circle of radius δ.) We'll compute this using the residue theorem — the denominator has a unique zero, at $z(w)$ (the $z(w)$ given by (1)), so this is

$$\operatorname{Res}\left(\frac{zf'(z)}{f(z)-w}, z(w)\right).$$

But we have a simple zero in the denominator, and the numerator doesn't share it; so we can just evaluate the numerator and take the derivative of the denominator, and we just get

$$\frac{z(w)\cdot f'(z(w))}{z(w)} = z(w).$$

So this function $z(w)$, which assigns to each w its pre-image, is expressible as this integral by the residue theorem.

Now to check that this is a holomorphic function of w, it's enough to check that $\int \frac{zf'(z)}{f(z)-w}$ is holomorphic. And we can do this by hand, using the same argument as for the Cauchy derivative formula — taking a difference quotient — to show that the LHS is holomorphic in $|w-w_0|<\varepsilon$. □

Remark 14.8. For this to work, we ned $z(w)\neq 0$ (otherwise we don't have an actual pole). But you can rig things so that it's not — we can just shift the function. (Alternatively, there isn't any pole, but that's fine because then the integral is 0, and that's what we need.)

We'll finish this argument in the exercises.

§14.4 More corollaries

Another corollary is an upgrade to our theorem on essential singularities.

Theorem 14.9

If f has an essential singularity at z_0, then for each $w_0\in\mathbb{C}$ and every $\varepsilon>0$ and $\delta>0$, the number of zeros of $f(z)-w$ in the punctured disk $0<|z-z_0|<\delta$ is an unbounded function of w in $|w-w_0|<\varepsilon$.

So if we take any ε-ball around w_0 and any δ-ball around our essential singularity, and we take w inside the ε-ball and look at how many times it gets hit, we can always find w hit a million times, or ten million times, or so on. So essential singularities are super wild.

Remark 14.10. This means that if you give me any ε, δ, w_0, and a number, e.g. one million, I can find w with $|w-w_0|<\varepsilon$ that gets hit a million times from $|z-z_0|<\delta$.

This is an exercise (with a hint) on the problem set. You can argue by contradiction and use some of this theory about the number of pre-images.

§15 November 16, 2023

Today we'll talk about the elliptic integral

$$x(u) = \int_0^u \frac{dt}{\sqrt{p(t)}}\, dt$$

where p is a degree 3 or 4 polynomial.

To study this, we developed the theory of elliptic functions.

Definition 15.1 (Weirstrass $\wp$-function)**.** Fix $\omega_1, \omega_2 \in \mathbb{C}^\times$ with non-real ratio, and let $L = \omega_1\mathbb{Z} + \omega_2\mathbb{Z}$ be the lattice they generate. Then

$$\wp(z) = \frac{1}{z^2} + \sum_{\omega \in L \setminus \{0\}} \left(\frac{1}{(z-\omega)^2} - \frac{1}{\omega^2} \right).$$

Last time we proved the following theorem.

Theorem 15.2

(1) $\wp$ is meromorphic (i.e., entire away from poles).

(2) $\wp$ has poles of order 2 at every $\omega \in L$.

(3) We have

$$\wp'(z) = -z \sum_{w \in L} \frac{1}{(z-\omega)^w} \text{ for } z \in \mathbb{C} \setminus L.$$

(4) $\wp$ and $\wp'$ are elliptic with periods ω_1 and ω_2.

We saw that if we have an elliptic function, it has to have degree at least 2 (it has to have poles). So this one has poles; its poles are at the lattice points. And both $\wp$ and $\wp'$ are doubly periodic (in contrast to sin or cos, which only have one period).

In effect, as we saw last class, this means we just need to look at the *period parallelogram* P (spanned by one ω_1 and ω_2) — then by periodicity, if we understand the function on P we understand it everywhere.

We need to develop the theory more to circle back to elliptic integrals. (At the beginning of the class we said that if $p(t)$ is a quadratic polynomial, then we get the arcsin function — the inverse is periodic. We'll see that the Weirstrass $\wp$-function is the inverse of an elliptic integral.)

Theorem 15.3

Let $d = \min_{\omega \in L\setminus\{0\}} |\omega|$. Then $\wp(z)$ has the series expansion

$$\wp(z) = \frac{1}{z^2} + \sum_{k=1}^{\infty} c_{2k} z^{2k} \text{ for } 0 < |z| < d,$$

where we have

$$c_{2k} = (2k+1) \sum_{\omega \in L \setminus \{0\}} \frac{1}{\omega^{2k+2}}.$$

Note that $\wp$ is even, so we only get even powers.

Where does this come from? We can unpack the expression inside the sum (inside the disk of radius d).

Proof. For $|\zeta| < 1$, we have the geometric series formula

$$\frac{1}{1-\zeta} = \sum_{k=0}^{\infty} \zeta^k$$

(which converges uniformly for any closed disk inside the unit circle). We can differentiate term-by-term to get that

$$-\frac{1}{(1-\zeta)^2} = \sum_{k=1}^{\infty} k\zeta^{k-1} = \sum_{k=0}^{\infty} (k+1)\zeta^k.$$

Now we can rewrite

$$\frac{1}{(z-\omega)^2} - \frac{1}{\omega^2} = \frac{1}{\omega^2}\left(\frac{1}{(1-(z/\omega))^2} - 1\right).$$

We can expand using our formula, and subtracting off the 1 from the first term, we get

$$\frac{1}{\omega^2}\sum_{k=1}^{\infty}(k+1)\frac{z^k}{\omega^k} = \sum_{k=1}^{\infty}\frac{z^k}{\omega^{k+2}}.$$

Now let's try to rewrite the Weirstrass function — we have

$$\wp(z) = \frac{1}{z^2} + \sum_{\omega \in L\setminus\{0\}}\left(\frac{1}{(z-\omega)^2} - \frac{1}{\omega^2}\right) = \frac{1}{z^2} + \sum_{\omega\in L\setminus\{0\}}\sum_{k=1}^{\infty}(k+1)\frac{z^k}{\omega^{k+2}}.$$

We have a double summation, so we can switch the order — we have to justify this, but the justification comes from the absolute summability of the Weirstrass function (away from finitely many bad points), which we proved in the first class — as long as we take ω sufficiently large relative to z, we get absolute summability. That allows us to switch all the orders of summation; so then we get

$$\wp(z) = \frac{1}{z^2} + \sum_{k=1}^{\infty}(k+1)\sum_{\omega\in L\setminus\{0\}}\frac{1}{\omega^{k+2}}z^k.$$

These should be the c_k's, but we need to show that the odd powers disappear. To see this, note that the operation $\omega \mapsto -\omega$ simply permutes the lattice points (and doesn't affect the sum). But if we do this change in the sum, we get a factor of $(-1)^{k+2} = -1$. So when k is odd, we get $\pm$ of the same terms in the sum, and all of them cancel with each other; this means only the even-k terms survive. □

Remark 15.4. The point of d is that we look at the biggest open disk that contains 0 and none of the other lattice points.

Now we have a series expansion of $\wp$ around one of its poles (namely, $z = 0$). We're now going to try to derive various functional equations for $\wp$. First, here's a first-order ODE.

Theorem 15.5

We have $\wp'(z)^2 = 4\wp(z)^3 - g_2\wp(z) - g_3$ in $\mathbb{C}\setminus L$, where $g_2 = 60\sum_{w\in L\setminus\{0\}}\frac{1}{\omega^4}$ and $g_3 = 140\sum_{\omega\in L\setminus\{0\}}\frac{1}{\omega^6}$.

This states $\wp$ is a solution to a certain nonlinear first-order differential equation (this is where we might start to see the integrand for our elliptic integral).

Proof. Where does this come from? We're going to obtain this from our series expansion for $\wp$. We can start by obtaining an expansion for $\wp'$ by differentiating term-by-term — in $0 < |z| < d$ we have

$$\wp'(z) = -\frac{2}{z^3} + \sum_{k=1}^{\infty} 2kc_{2k}z^{2k-1}.$$

Squaring this gives

$$\wp'(z)^2 = \left(-\frac{2}{z^3} + 2c_2 z + 4c_4 z^3 + 6c_6 z^5 + \cdots\right).$$

We'll just keep track of the terms that don't have a z (i.e., that don't go to 0 as $z \to 0$) — we'll just throw the terms that go to 0 with z into an error. This gives us

$$\wp'(z)^2 = \frac{4}{z^6} - \frac{8c_2}{z^2} - \frac{16c_4}{z^4} + O(|z|^2).$$

(We say $f = O(g)$ if $|f| \le C\,|g|$.)

We can similarly compute

$$\wp(z)^3 = \frac{1}{z^6} + \frac{3c_2}{z^2} + 3c_4 + O(z^2).$$

Now we subtract (we're going to try to kill off the singular terms one by one) — so we get

$$\wp'(z)^2 - 4\wp(z)^3 = -\frac{20c_2}{z^2} - 28c_4 + O(|z|^2).$$

Now we have a $\frac{1}{z^2}$ term; and we can kill that by adding a multiple of $\wp$ — we have

$$\wp'(z)^2 - 4\wp(z)^3 + 20c_2\wp(z) + 28c_4 = O(|z|^2).$$

Now let's look at the function on the left-hand side, which is $O(|z|^2)$ (this is relevant near 0). But the left-hand side is defined away from all lattice points; and it's elliptic (because $\wp$ is elliptic and $\wp'$ is, and squaring an elliptic function gives you another).

So we have an elliptic function with periods ω_1 and ω_2 on the left-hand side. But it's bounded as $z \to 0$. So it doesn't have a pole there! That means it's constant. And the constant has to be 0, since we can evalulate at $z = 0$.

So the left-hand side is elliptic, bounded (it has no poles in P), and is 0 at $z = 0$; that means it must equal 0 identically. Now that finishes the theorem (we can plug in our expressions for c_2 and c_4 to get g_2 and g_3). □

This is an ODE for $\wp$. It's a bit unsatisfactory because we have a square in our top-order derivative; but we can obtain a nicer-looking system by just differentiating (giving a second-order equation).

Corollary 15.6

We have $\wp''(z) = 6\wp(z)^2 - 10c_2$ (or alternatively $6\wp(z)^2 - \frac{1}{2}g_2$).

This can be seen just by differentiating the original (on the left-hand side we get $2\wp'(z)\wp''(z)$, and on the right-hand side we get a bunch of $\wp'(z)$s that we can cancel out).

Now we'll wrap this up; we'll go back to our ODE for $\wp'$. (The notation g_2 and g_3 is for historical reasons.) We have

$$\wp'(z)^2 = 4\wp(z)^3 - g_2\wp(z) - g_3,$$

and we can now integrate this — we have

$$\frac{\wp'(z)}{\sqrt{4\wp(z)^2 - g_2\wp(z) - g_3}} = 1,$$

and so integrating both sides, we have

$$\int_{[0,z]} \frac{\wp'(z)}{\sqrt{4\wp(z)^2 - g_2\wp(z) - g_3}}\,dz = \int_{[0,z]} 1\,dz = z.$$

Now change variables, replacing $z \mapsto \wp(z)$ (using u-substitution). As $z \to 0$ we have a pole, so $\wp(z) \to \infty$; this means we get

$$\int_{[\infty, \wp(z)]} \frac{d\wp}{\sqrt{4\wp^3 - g_2\wp - g_3}} = z,$$

which can be written as

$$-\int_{[u,\infty]} \frac{du}{\sqrt{4u^3 - g_2 u - g_3}} = z.$$

The left-hand side is the elliptic integral $I(u)$ we're interested in; and we have $I(\wp(z)) = z$. So this says that the inverse function of our integral is the Weirstrass $\wp$ function. (This is just like how if we take $\sqrt{1-x^2}$ we get arcsin, which is the inverse function of sin — here the elliptic integral is the inverse function of the doubly periodic $\wp$.)

So this wraps the circle around.

There's a lot of function theory, from Abel's point of view and Jacobi's point of view (for example, you can look at associated Riemann surfaces).

This might look odd at first, but one way to think about it is to put yourself in the 19th century, where you're trying to develop a theory of functions. What types of functions are out there? You have the exponential and trigonometric functions appearing naturally; but these functions are just as natural (they're doubly periodic). So there's this whole classical theory of finding these functions — there's $\wp$, θ functions, the Γ function, the ζ function, and so on; and people studied their property, just as in high school we study sin and cos.

§15.1 More on meromorphic functions

This is some specific function theory; now we'll come back to the more abstract setting of developing the theory of complex analysis. One area we haven't touched too much on is the general theory of meromorphic functions which are not rational functions. In particular, we'll try to study the behavior as $z \to \infty$.

Theorem 15.7

Let f be a meromorphic function on $\mathbb{C}$. Then f is a rational function if and only if f has finitely many (possibly 0) poles and for $|z| \to \infty$, either $|f(z)|$ is bounded or $|f(z)| \to \infty$.

The extra condition essentially can be thought of as saying that at the point at ∞ we have a removable singularity or a pole (you can't be bounded along some direction and unbounded along another direction).

Remark 15.8. The last condition rules out things like e.g. $e^z/p(z)$ — if we take $z = it$ then $|e^{it}| = 1$, but if we take $z = t$ then $|e^t| \to \infty$. So we have two different behaviors as $z \to \infty$, violating the second condition.

Remark 15.9. This implies that if we have an entire function that tends to ∞ as $z \to \infty$, then it has to be a polynomial.

This is a really rigid setting — if you have finitely many poles and have a limit at ∞, then you have to be rational. Next we're going to leave this world and consider meromorphic functions which are not rational, and we'll see they have pretty extreme behavior (like Casorati–Weirstrass on steroids).

Proof. If f is rational, then we can write $f = g/h$ for polynomials g and h; the poles are the zeros of h which aren't zeros of g, so there can be only finitely many of them by the fundamental theorem of algebra.

Furthermore, we have a well-defined limit as $z \to \infty$ — it's 0 if $\deg(g) < \deg(h)$, some nonzero complex number if $\deg(g) = \deg(h)$, and ∞ if $\deg(g) > \deg(h)$.

For the backwards direction, suppose f has finitely many poles. Then we can remove the principal parts of these poles — let

$$\varphi(z) = f(z) - \sum_{k=1}^{n} p_k \left(\frac{1}{z - z_k} \right),$$

where $z_1, \ldots, z_n$ are the poles of f and $p_1, \ldots, p_n$ their principal parts.

Then this function is entire — the poles are isolated, so if we look in a neighborhood of one of the poles (the principal parts only have problems around those poles as well), we can write f as its principal part plus something holomorphic, and subtracting the principal part kills off the singularity. So $\varphi(z)$ is entire.

Now we can look at the behavior of $\varphi(z)$ at ∞; this is determined by that of $f(z)$, since all the principal parts tend to 0 as $|z| \to \infty$ (they're polynomials with z in the denominator).

So $\varphi(z)$ is bounded as $|z| \to \infty$ if and only if f is. If f is bounded, then $\varphi(z)$ is bounded as $z \to \infty$, so it has to be a constant function by Liouville's theorem (it's entire and bounded, so therefore constant).

In the second case, if $|f| \to \infty$ as $|z| \to \infty$, then the same is true of φ. But then $\varphi(1/z)$ is holomorphic in $\mathbb{C} \setminus \{0\}$ with a pole at $z = 0$ (since $\varphi(1/z) \to \infty$ as $z \to 0$). Let m be the multiplicity of this pole. We showed that then we have $|\varphi(1/z)| \leq c(1/|z| + 1)^m$ (the multiplicity of the pole tells us how we tend to ∞). (This bound is true for $|z| \leq 1$, but it's also true everywhere because the behavior of $\varphi(1/z)$ outside the unit disk is the same as the behavior of φ inside.)

But now we can apply the generalized Liouville's theorem — we have a polynomial bound on the growth rate of φ, which implies that φ is a polynomial.

So φ is a polynomial; this means f minus some rational function is a polynomial, so f is a rational function. $\square$

This is maybe a bit disappointing — the functions with these nice properties are just the rational functions. But rational functions are actually quite interesting; there's lots of questions about them we don't yet understand.

§15.2 The Riemann sphere

Imagine a sphere with its south pole sitting at the origin of the complex plane (imagine the complex plane is the floor, and we put a basketball on the floor).

Then we can fix the point at the north pole; for each point on the sphere, we can draw a straight line down from the north pole, and it'll hit the complex plane somewhere. So we can associate each point on the sphere with a unique point in $\mathbb{C}$.

How do we compute this function? Let's write $\mathbb{S}^2 = \{(u_1, u_2, u_3) \in \mathbb{R}^3 \mid u_1^2 + u_2^2 + u_3^2 = 1\}$.

Definition 15.10. The *stereographic projection* is the map $\zeta \colon \mathbb{S}^2 \setminus \{\mathrm{NP}\} \to \mathbb{C}$ given by

$$\zeta(u_1, u_2, u_3) = \frac{1}{1 - u_3}(u_1 + iu_2).$$

(The formula is not that important; all you need to know is that it's a one-to-one correspondence.)

The north pole doesn't get anywhere, because it's where we're sending out this ray from.

We can write down the inverse of this map as

$$z \mapsto \frac{1}{1 + |z|^2}(2\,\mathrm{Re}(z), 2\,\mathrm{Im}(z), 1 - |z|^2).$$

Now what happens as $z \to \infty$? (This is actually easier to see in the picture.) If we take a point really far away on the plane, we draw a really flat line coming out, which hits the sphere somewhere close to the north pole — and this will be true no matter which direction we go in. So as z gets big, we're in a small neighborhood in the north pole — as $z \to \infty$ in $\mathbb{C}$, we're converging to the north pole on the sphere.

So you can think of adding one more point to the complex plane, a point at ∞, and you get the whole sphere. (The real line is ordered, so it's sometimes helpful to think about $\pm\infty$, but you can make them the same point by identifying with a circle. In $\mathbb{C}$ there's no $\pm\infty$; but you can still think of ∞ as being the 'missing point' to which the north pole is defined.)

Now we can identify $\mathbb{C}$ with the Riemann sphere minus a point; this makes sense geometrically, because a sphere is not the same as a flat plane (so you cannot get the whole thing). But if we consider this new object $\mathbb{C}_\infty = \mathbb{C} \cup \{\infty\}$, this is the same thing as $\mathbb{S}^2$.

The point is that now in the earlier theorem, we can think of rational functions as functions $\mathbb{C}_\infty \to \mathbb{C}_\infty$ — if a function has a pole, then the function goes to ∞ in the range, so it tends to the north pole in the sphere. And similarly, when we compactify the domain and think of the *domain* as a Rieman sphere, the second line in the theorem is a statement about the behavior of the function at the north pole. Thinking of functions on the sphere, the north pole isn't special — it's just another point.

So you can also think about holomorphic functions $\mathbb{C}_\infty \to \mathbb{C}_\infty$ (where 'poles' are points where the function tends to ∞, i.e., the north pole; and the behavior as $z \to \infty$ has to be tamed to do this, since we still have to land in the sphere, so we need to have either a pole or some value).

> **Remark 15.11.** The reason $e^{1/z}$ has an essential singularity at 0 is that e^z has one at ∞ — earlier we thought e^z was great and $e^{1/z}$ was terrible. But there's no difference between the south pole and north pole here — e^z is fine at all the complex numbers but has an essential singularity at the north pole.

So now we can think about, what are the holomorphic functions $\mathbb{C}_\infty \to \mathbb{C}_\infty$? Earlier, if we had an entire bounded function, it had to be constant. And the earlier theorem says that a function f is a holomorphic function $\mathbb{C}_\infty \to \mathbb{C}_\infty$ if and only if f is rational.

What if we want functions $f\colon \mathbb{C}_\infty \to \mathbb{C}$? This is just constants — you can think of $\mathbb{C}_\infty$ as a compact set, and a continuous function on a compact set has to be bounded. So this is a very uninteresting class of functions.

In geometry and topology, people like studying functions between objects. You get nowhere thinking about function theory from the Riemann sphere. So what people started doing (and what we'll get close to) is that in some cases studying functions isn't enough. And this is where the theory of differential forms comes up. There aren't good holomorphic functions on $\mathbb{C}_\infty$, other than constants; so then you start looking for holomorphic 1-forms.

> **Remark 15.12.** What happens to the 'finitely many poles' condition? If we had infinitely many poles on a sphere, they'd have to cluster, and that would contradict isolation of zeros.

We'll get back to this more when we talk about conformal mappings and the Riemann mapping theorem. The point about this map that makes it nice is that it preserves angles.

> **Remark 15.13.** The formulas for stereographic projection are both if the plane goes through the middle of the sphere.

§16 November 28, 2023

In the last few classes we'll prove the Riemann mapping theorem. Before that, we'll see an easier but fundamental problem.

> **Remark 16.1.** The next homework will be due next Friday; there is no homework due this Friday.

§16.1 The Riemann sphere

Last time, we introduced the Riemann sphere, which we'll think of as $\mathbb{C} \cup \{\infty\}$. We can think of this via the stereographic projection map — imagine that we have the unit sphere $\mathbb{S}^2 = \{u \in \mathbb{R}^3 \mid |u| = 1\}$ sitting inside $\mathbb{R}^3$. Then the stereographic projection map $\Phi_{\text{NP}}: \mathbb{S}^2 \setminus \{(0,0,1)\} \to \mathbb{C}$ (where $(0,0,1)$ is the *north pole* of $\mathbb{S}^2$) is defined by taking the north pole and drawing straight lines emanating out of the north pole. These lines intersect the xy-plane; and we identify the point of intersection of each line with the xy-plane with the point on $\mathbb{S}^2$ at which it punctures the sphere.

To figure out an explicit formula for Φ, we can look at a planar slice and use similar triangles; we end up with

$$\Phi(u_1, u_2, u_3) = \frac{1}{1 - u_3}(u_1 + iu_2) \in \mathbb{C}.$$

This is a well-defined map as long as $u_3 \neq 1$ (which is true at all points in $\mathbb{S}^2$ except the north pole). This has an inverse map $\zeta_{\text{NP}}: \mathbb{C} \to \mathbb{S}^2$, where

$$\zeta(z) = \frac{1}{1 + |z|^2}(2\operatorname{Re}(z), 2\operatorname{Im}(z), |z|^2 - 1).$$

We let $\mathbb{C}_\infty = \mathbb{C} \cup \{\infty\}$ be the Riemann sphere. Under this map, we identify the north pole of the sphere with the point at ∞ (this is called the *one-point compactification* of $\mathbb{C}$).

§16.2 Riemann surfaces

The Riemann sphere is a basic example of something called a *Riemann surface*. We won't go into the details of what a Riemann surface is, but the point is that the map ζ_{NP} doesn't quite cover the sphere if we just restrict its domain to $\mathbb{C}$ — it misses a point — so we need two of these maps to cover the whole sphere, and we could define another one by doing a stereographic projection from a different point (e.g., the south pole). For example, we could define $\Phi_{\text{SP}}: \mathbb{S}^2 \setminus \{(0,0,-1)\} \to \mathbb{C}$ given by

$$\Phi_{\text{SP}}(u_1, u_2, u_3) = \frac{1}{1 + u_3}(u_1 - iu_2).$$

The inverse of this map is $\zeta_{\text{SP}}: \mathbb{C} \to \mathbb{S}^2$ given by

$$\zeta_{\text{SP}}(z) = \frac{1}{1 + |z|^2}(2\operatorname{Re}(z), 2\operatorname{Im}(z), 1 - |z|^2)$$

(now this map will miss the south pole).

These functions mapping $\mathbb{C}$ to the sphere are called *coordinate charts* — both ζ_{NP} and ζ_{SP} are called coordinate charts. The sphere is a 2-dimensional object (it sits in $\mathbb{R}^3$ but is a 2-dimensional surface), so we should be able to describe portions of it using only two coordinates (i.e., by a single complex number); and this is a way to describe the sphere using complex numbers.

If we take these two charts together, we cover the whole sphere. Then we have to worry, are the charts compatible with each other? This means we want a *transition map* — we have one map ζ_{NP} from one copy of $\mathbb{C}$ to the sphere that misses the north pole, and another map ζ_{SP} from one copy of $\mathbb{C}$ to the sphere that misses the south pole. Every other point on the sphere is hit by both maps, so looking at the overlap, we get a map between these two copies of the complex plane if we exclude 0 — we get a map $\mathbb{C}^\times \to \mathbb{C}^\times$ where we go up to the sphere with ζ_{NP} and then down to the other copy of $\mathbb{C}$ with Φ_{SP}. This means we consider the map $\Phi_{\text{SP}} \circ \zeta_{\text{NP}}$. We can figure out what this map is — if we do the algebra, we find it's the map $\frac{1}{z}$. So this gives us the transition map $\mathbb{C}^\times \to \mathbb{C}^\times$ (where $\mathbb{C}^\times = \mathbb{C} \setminus \{0\}$ — the map isn't defined at the origin because we should only think about it as a map on the overlap region).

And the same computation shows you that

$$\Phi_{\text{NP}} \circ \zeta_{\text{SP}}(z) = \frac{1}{z}$$

as well. These two compositions are called *transition maps*.

Definition 16.2 (Riemann surfaces, roughly)**.** A *Riemann surface* is a surface that can be covered by coordinate charts in such a way, such that the transition maps between the coordinate charts are holomorphic.

§16.3 Holomorphic functions on the Riemann sphere

We'll now extend the notion of holomorphic functions to three new possibilities — we can think about holomorphic functions $f\colon \mathbb{C}_\infty \to \mathbb{C}$, $\mathbb{C} \to \mathbb{C}_\infty$, and $\mathbb{C}_\infty \to \mathbb{C}_\infty$. (So we're extending either the domain or the range of a holomorphic function.)

What should we require in these three cases? First, of course we want holomorphic functions to be continuous, in all these cases — this means we can think about these functions within our coordinate charts. And we want them to be holomorphic relative to the charts. What does this mean? For the first group — functions $f\colon \mathbb{C}_\infty \to \mathbb{C}$ — we know what it means to be holomorphic at points $z \in \mathbb{C}$, and we also need to require the function to be holomorphic in the neighborhood of ∞.

Definition 16.3. We say $f\colon \mathbb{C}_\infty \to \mathbb{C}$ is *holomorphic near* $z = \infty$ if $f(\frac{1}{z})$ is holomorphic near $z = 0$.

This means we're playing well with the transition maps — so what we're trying to say here is that for maps $\mathbb{C}_\infty \to \mathbb{C}$, we need to understand the notion of being holomorphic at or near ∞, and here we're saying that's the same as saying $f(\frac{1}{z})$ is holomorphic near 0 (which is something we know how to deal with).

One way to think about $\mathbb{C}_\infty$ is as $\mathbb{C} \cup \{\infty\}$; we know what it means to be holomorphic at $z \in \mathbb{C}$, so all we need to understand is what it means to be holomorphic at ∞, and we can think of that as the holomorphicity of $f(\frac{1}{z})$. This is equivalent to putting ourselves in the other chart, where now the north pole corresponds to 0 in the south-pole chart.

Question 16.4. What are the holomorphic functions $\mathbb{C}_\infty \to \mathbb{C}$?

These are just the constant maps, by Liouville's theorem — they're entire functions, and they have to be bounded (since $\mathbb{C}_\infty$ is compact). So by Liouville's theorem $\mathcal{H}(\mathbb{C}_\infty, \mathbb{C})$ is just the set of constants, which is pretty boring.

Now for the second class of functions — $f\colon \mathbb{C} \to \mathbb{C}_\infty$ — we have to think about what it means for some z_0 to be sent to the north pole, i.e., to ∞. So we need to define what it means to be holomorphic in that case.

Definition 16.5. If $f\colon \mathbb{C} \to \mathbb{C}_\infty$ and $f(z_0) = \infty$, then we say f is holomorphic at $z = z_0$ if $f(z)^{-1}$ is holomorphic at z_0.

The holomorphic functions $\mathbb{C} \to \mathbb{C}_\infty$ are the *meromorphic* functions. So this is another way of thinking about meromorphic functions — holomorphic functions, but we've extended the range to allow ∞ (this is exactly what it means to have a pole).

And for the third class — functions $\mathbb{C}_\infty \to \mathbb{C}_\infty$ — all we need to deal with is the case if $f(\infty) = \infty$.

Definition 16.6. For functions $f\colon \mathbb{C}_\infty \to \mathbb{C}_\infty$, if $f(\infty) = \infty$, then we say f is holomorphic near $z = \infty$ if and only if $f(\frac{1}{z})^{-1}$ is holomorphic near 0.

So we've extended the definition of holomorphicity by using the transition function, both on the domain and the range.

These functions are what we discussed last class — last class we proved the theorem that $\mathcal{H}(\mathbb{C}_\infty, \mathbb{C}_\infty)$ is exactly the set of rational functions.

§16.4 Conformality

Definition 16.7. For a region $D \subseteq \mathbb{C}$, we say a function $f \in \mathcal{C}^1(D)$ is *conformal on* D if $df \neq 0$ in D and df preserves angles and orientation.

What does this mean? For a function f differentiable at z_0, df is the matrix-valued function that gives the first-order linear expansion of f —

$$f(z) = f(z_0) + df(z_0)h + r(z_0)(h),$$

where $h = z - z_0$ and $r(z_0)$ vanishes appropriately. We saw that $f \in \mathcal{H}(D)$ if and only if $df = \rho A$ where $\rho > 0$ and A is a rotation matrix — i.e., df is given by rotations and dilations (which preserve angles and orientations). This followed from the Cauchy–Riemann equations — if $f = u + iv$, then

$$df = \begin{bmatrix} \partial_x u & \partial_y u \\ \partial_x v & \partial_y v \end{bmatrix},$$

and the Cauchy–Riemann equations state that

$$\partial_x u = \partial_y v \text{ and } \partial_x v = -\partial_y u,$$

which means this matrix is given by

$$df = \begin{bmatrix} \partial_x u & -\partial_x v \\ \partial_x v & \partial_x u \end{bmatrix}.$$

Conformal is more restrictive than being holomorphic, since we also require $df \neq 0$ — in particular, if f is holomorphic in Ω, then it's conformal at every point except where $df = 0$.

Example 16.8

The map $z \mapsto z^2$ is holomorphic (it's an entire function); it preserves angles except at the origin, where it doubles angles (if you compute this matrix and evaluate at $z = 0$). So it's conformal in any region of $\mathbb{C}$ excluding 0.

Remark 16.9. What does preserving orientation mean — is it the same as saying that the determinant is positive? It should be — for example, the map $z \mapsto \overline{z}$ reverses orientation.

Remark 16.10. The values of ρ and A will change with z — neither ρ nor A are constant functions (they change at each point). But each point should have that $\rho(z) > 0$ and that $A(z) \in \mathrm{SO}(2, \mathbb{R})$.

Remark 16.11. What does $df = 0$? We can think of df as a 2×2 matrix; this means every value of that matrix is 0. Equivalently, $df \neq 0$ means that $f'(z_0) \neq 0$. (Here f' has to take the form of multiplication by a complex number; we can view that either as a matrix or as a complex number.)

Remark 16.12. How does the expansion $f(z) = f(z_0) + df(z_0)h + r(z_0)(h)$ work? We can identify a function on D with a function $\mathbb{R}^2 \to \mathbb{R}^2$ — we can view z as the vector $(x, y) \in \mathbb{R}^2$, and view f as the function $f(z) = (u(x, y), v(x, y))$. Then df is the Jacobian matrix of this function. If we do Taylor expansion, then

$$f(z) = f(z_0) + \begin{bmatrix} \partial_x u & \partial_y u \\ \partial_x v & \partial_y v \end{bmatrix} (z_0) \begin{bmatrix} h_1 \\ h_2 \end{bmatrix} + \cdots .$$

Remark 16.13. Why does $z \mapsto z^2$ preserve angles at $z \neq 0$ — what does preserving angles mean? Its derivative is $2z$; that means multiplication by the derivative just multiplies by $2z$, which preserves angles. The point of preserving angles is that the matrix df preserves angles, i.e., is given by multiplication by a complex number.

§16.5 Conformal automorphisms of the Riemann sphere

We want to solve the following problem.

Question 16.14. Find all conformal automorphisms of $\mathbb{C}_\infty$.

What is a conformal automorphism? These are bijections (i.e., one-to-one and onto maps) such that both the map and its inverse are conformal.

We already know that these functions have to be holomorphic, so they have to be rational functions. But they also need to be bijections. So which rational functions are allowed? Both the numerator and denominator have to have degree at most 1. And that's all.

To repeat, let $\mathrm{Aut}(\mathbb{C}_\infty)$ be the set of conformal automorphisms. Since $\mathcal{H}(\mathbb{C}_\infty, \mathbb{C}_\infty)$ is the set of rational functions, we know all $f \in \mathrm{Aut}(\mathbb{C}_\infty)$ must be rational. And then by the fundamental theorem of algebra, if $f = p/q$ for polynomials p and q, we must have $\deg(p) \leq 1$ and $\deg(q) \leq 1$ — otherwise f wouldn't be a bijection.

These functions have a name — they're called *Möbius transformations* or *fractional linear transformations*.

Definition 16.15. We define $\mathrm{GL}_2(\mathbb{C})$ as the set of invertible 2×2 matrices with entries in $\mathbb{C}$ (i.e., matrices with nonzero determinant).

Theorem 16.16

Every $A \in \mathrm{GL}_2(\mathbb{C})$ defines a *Möbius transformation*

$$T_A(z) = \frac{az + b}{cz + d}, \text{ where } A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}.$$

(i) $T_A \in \mathcal{H}(\mathbb{C}_\infty, \mathbb{C}_\infty)$.

(ii) The map $A \mapsto T_A$ depends only on the equivalence class of A in $\mathrm{GL}_2(\mathbb{C})$ under the relation $A \sim B$ if and only if $A = \lambda B$ for $\lambda \in \mathbb{C}$. In other words, the set of Möbius transformations can be identified with

$$\mathrm{PSL}_2(\mathbb{C}) = \mathrm{SL}_2(\mathbb{C})/\{\pm \mathrm{Id}\}.$$

(iii) $T_A \circ T_B = T_{AB}$ and $T_A^{-1} = T_{A^{-1}}$. In particular, each T_A is a conformal automorphism.

We already know (i) — it's a rational function.

For (ii), we associate to each A such a rational function, so we have a map between $\mathrm{GL}_2(\mathbb{C})$ and the Möbius transformation. We can see that if we multiply A by some complex number ρ, then ρ factors out of the numerator and denominator, so we get back the same function again. So two matrices different just by multiplication by a number give the same linear transformation. This means the map $A \mapsto T_A$ only depends on an *equivalence class* of matrices in GL_2. Here $\mathrm{SL}_2(\mathbb{C})$ is the subgroup of $\mathrm{GL}_2(\mathbb{C})$ consisting of 2×2 matrices A with $\det(A) = 1$; and then we can also flip signs (Id and $-\mathrm{Id}$ give rise to the same transformation).

The point (ii) says that these maps play well with matrix multiplication. Then since T_A is conformal and has an inverse, which is also given by the same type of object (so is also conformal), this means it's a conformal automorphism.

Proof. We already proved (i). Then (iii) is just a computation, which will be left as an exercise.

For (ii), on one hand, if $A = \lambda B$ then $T_A = T_B$ (this can be seen by computation). For the converse, suppose that $T_A = T_{\tilde{A}}$; by rescaling (using the first statement), we can assume that $\det(A) = \det(\tilde{A}) = 1$. We want to now find a relationship between A and $\tilde{A}$.

Since $T_A = T_{\tilde{A}}$ and both are holomorphic functions, their *derivatives* also have to be equal to each other. We have

$$T_A(z) = \frac{az + b}{cz + d},$$

so then we can compute

$$T'_A(z) = \frac{a(cz + d) - c(az + b)}{(cz + d)^2} = \frac{ad - bc}{(cz + d)^2}.$$

If we take another matrix $\tilde{A}$, then we have

$$T'_{\tilde{A}}(z) = \frac{\tilde{a}\tilde{d} - \tilde{b}\tilde{c}}{(\tilde{c}z + \tilde{d})^2}.$$

By the first step we can restrict to matrices with determinant 1, so the numerator in each of these is equal to 1; this means we need to have

$$\frac{1}{(cz + d)^2} = \frac{1}{(\tilde{c}z + \tilde{d})^2},$$

and therefore $cz + d = \pm(\tilde{c}z + \tilde{d})$. If they're equal, then we get $d = \tilde{d}$ and $c = \tilde{c}$, and we can go back and see that $a = \tilde{a}$; and if they're negatives, we get $a = -\tilde{a}$. So this means two matrices correspond to the same transformation if and only if they're in the same equivalence class in $\mathrm{GL}_2(\mathbb{C})$ as described. □

This might seem tangential to the class, but these objects are pretty fundamental, so we're going to study them for a bit before we get to the Riemann mapping theorem. This gives us the conformal automorphisms of the Riemann sphere $\mathbb{C}_\infty$; we're going to want to do something much more complicated, which is to try to understand functions like z^2 or $\sqrt{z}$ — we're going to try to understand which regions of $\mathbb{C}$ are conformally equivalent to other regions, in particular the unit disk.

§16.6 More about Möbius transformations

Lemma 16.17

Every Möbius transformation is generated by a composition of four elementary transformations:

(1) Translation — $z \mapsto z + z_0$ for $z_0 \in \mathbb{C}$.

(2) Dilation — $z \mapsto \lambda\mathbb{C}$ for $\lambda \in \mathbb{R}_{>0}$.

(3) Rotation — $z \mapsto e^{i\theta}z$ for $\theta \in \mathbb{R}$.

(4) Inversion — $z \mapsto \frac{1}{z}$.

So we can make up every Möbius transformation by composing these four.

Proof. Consider some Möbius transformation

$$T_A = \frac{az+b}{cz+d}.$$

If $c = 0$, then this is given by

$$\frac{a}{d}z + \frac{b}{d},$$

which is a composition of the first three — multiplication by a complex number is dilation plus rotation (we can see this by using polar coordinates), and then adding $\frac{b}{d}$ corresponds to translation.

If $c \neq 0$, then we also need inversion. We can write

$$T_A(z) = \frac{bc - ad}{c^2}\left(\frac{1}{z + \frac{d}{c}}\right) + \frac{a}{c},$$

which consists of all four operations. □

Remark 16.18. We can think of inversion as flipping the inside to the outside, so it's not a composition of the first three — the inside of the unit disk gets sent to the outside, and vice versa.

Lemma 16.19

The map T_A takes circles and lines to circles and lines.

We'll prove this next time. (A line might become a circle, and a circle might become a line.)

Next time, we'll also see that these maps are completely determined by their action on three points — they can have at most two fixed points, and then they're completely determined by their action on three distinct points.

§17 November 30, 2023

§17.1 More on Möbius transformations

For a matrix $A \in \mathrm{GL}_2(\mathbb{C})$, we defined

$$T_A(z) = \frac{az+b}{cz+d} \text{ where } A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}.$$

(Here T_A is a map $\mathbb{C}_\infty \to \mathbb{C}_\infty$)

Fact 17.1 — Any non-identity Möbius transformation T_A has at most 2 fixed points.

Proof. If z is a fixed point, then we must have

$$z = T_A(z) = \frac{az+b}{cz+d}.$$

If we multiply through, we get a quadratic polynomial in z (or a linear polynomial, if $c = 0$), so we have at most two fixed points. □

Example 17.2

The map $T(z) = \frac{1}{z}$ has two fixed points (± 1), while $T(z) = z + 1$ has only one (∞).

Theorem 17.3

(i) A Möbius transformation is completely determined by its action on three distinct points.

(ii) Given $z_1, z_2, z_3 \in \mathbb{C}_\infty$, there exists a unique Möbius transformation such that $T(z_1) = 0$, $T(z_2) = 1$, and $T(z_3) = \infty$.

Proof. For (i), suppose that we have two Möbius transformations T and S which agree at three points. Then we can consider the map $S^{-1} \circ T$. This map will fix each of those three points, so it must be the identity, which means $S = T$.

For (ii), uniqueness follows from (i), so it suffices to prove existence. To do so, take the map

$$T(z) = \frac{z - z_1}{z - z_3} \cdot \frac{z_2 - z_3}{z_2 - z_1}.$$

When we plug in z_1 we get 0; when we plug in z_2 everything cancels and we get 1; and when we plug in z_3 we get ∞.

We have to adjust this a bit if one of z_1, z_2, or z_3 is ∞. In that case, you just take the limit — if $z_3 = \infty$, then we take the limit of this fractional linear transformation as $z_3 \to \infty$, and we get something that works. □

§17.2 Automorphisms of other objects

We've talked about automorphisms of $\mathbb{C}_\infty$ — we found that $\operatorname{Aut}(\mathbb{C}_\infty) = \mathrm{PSL}_2(\mathbb{C}) = \mathrm{SL}_2(\mathbb{C})/\{\pm \mathrm{Id}\}$. But you can also think about automorphisms of the unit disk $\mathbb{D} = \{z \mid |z| < 1\}$, and the upper half plane $\mathbb{H} = \{z \mid \operatorname{Im}(z) > 0\}$. In an exercise on the homework, we'll show that

$$\operatorname{Aut}(\mathbb{D}) = \left\{ e^{i\theta} \frac{z - \alpha}{1 - \overline{\alpha} z} \mid \theta \in \mathbb{R}, \alpha \in \mathbb{D} \right\}$$

(these are a subset of all Möbius transformations; the key ingredient in the proof is what's called the Schwarz lemma).

Remark 17.4. By *automorphisms*, we mean conformal bijections with conformal inverse.

There's a couple of ways to describe $\mathrm{Aut}(\mathbb{H})$. One is to use $\mathrm{Aut}(\mathbb{C}_\infty)$, and note that if we want to preserve $\mathbb{H}$ then we have to preserve the real axis — so we end up with

$$\mathrm{Aut}(\mathbb{H}) = \mathrm{PSL}_2(\mathbb{R}) = \mathrm{SL}_2(\mathbb{R})/\{\pm \mathrm{Id}\}.$$

What this means is that $\mathrm{Aut}(\mathbb{H})$ is the set of maps

$$T_A(z) = \frac{az+b}{cz+d} \text{ for } a, b, c, d \in \mathbb{R} \text{ with } ad - bc = 1.$$

One way to prove this, once we know $\mathrm{Aut}(\mathbb{D})$, is to prove that $\mathbb{D}$ is conformally equivalent to $\mathbb{H}$ — denoted $\mathbb{D} \simeq \mathbb{H}$ — meaning that there exists a conformal bijection with conformal inverse between the two sets. This might seem surprising since $\mathbb{D}$ is bounded and $\mathbb{H}$ is not; but you can see this by just writing down a conformal equivalence. We can define

$$T(z) = \frac{i-z}{i+z}.$$

Then $T\colon \mathbb{H} \to \mathbb{D}$. To see this, if z is in the upper half plane, then we're closer to i than to $-i$; so $|T(z)| < 1$. So T does map into $\mathbb{D}$. Meanwhile, we can write down its inverse

$$S(w) = i \cdot \frac{1-w}{1+w}.$$

Then we can check that S maps $\mathbb{D}$ to $\mathbb{H}$. You can just check by function composition that these two are inverses of each other; and they're clearly holomorphic (T only has a pole at $-i$, which doesn't matter if we restrict to $\mathbb{H}$; and S only has an issue at $w = -1$, which isn't in $\mathbb{D}$) and you can check that their derivatives don't vanish.

§17.3 Conformal equivalences

So you can start by studying maps from an object to itself — e.g. the Riemann sphere, $\mathbb{D}$, or $\mathbb{H}$. But then you can expand this question and ask which regions of $\mathbb{C}$ are conformally equivalent to each other.

Example 17.5

We showed earlier that $\mathbb{D}$ is conformally equivalent to $\mathbb{H}$.

Example 17.6

The map $z \mapsto z^2$ takes the quadrant $A = \{x > 0, y > 0\}$ and sends it to $\mathbb{H} = \{y > 0\}$; so A is conformally equivalent to $\mathbb{H}$.

Example 17.7

Consider the map $f(z) = e^z$. If we take a horizontal line $y = y_0$, then we have $e^{x+iy_0} = e^x e^{iy_0}$; so we get a half-ray coming out of the origin. And so a horizontal box is mapped by e^z to a portion of an annular region. (e^z is not a conformal map on the whole of $\mathbb{C}$, but it gives a conformal map between regions like this.)

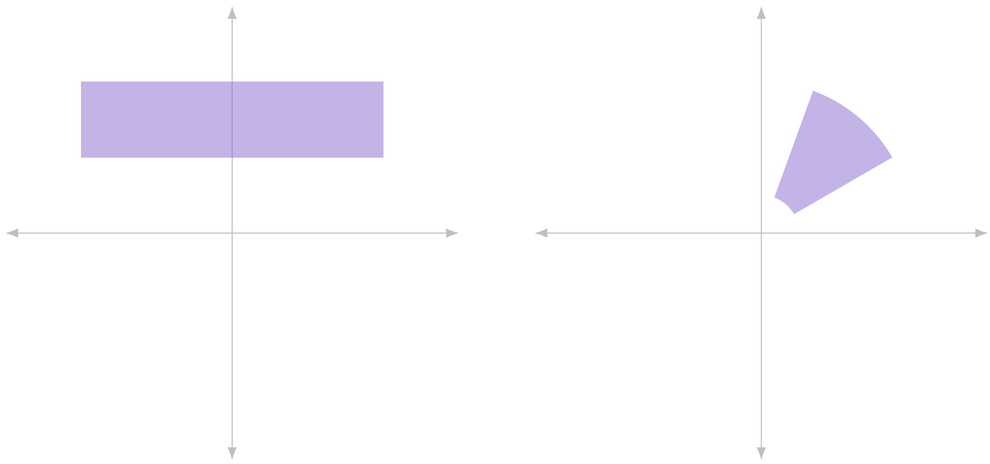

Question 17.8. Which regions D and Ω can be mapped onto each other conformally in bijective fashion? In particular, what if $D = \mathbb{D}$?

First, what are the obstacles — can we map $\mathbb{C}$ onto $\mathbb{D}$? The obstruction is Liouville's theorem — this would lead to a bounded entire function, which can't exist by Liouville's theorem (the only such things are constant functions). So we cannot have $\mathbb{C} \simeq \mathbb{D}$.

Another obstruction is that if we have a conformal map from one region onto another, it certainly needs to be continuous. So if we have a region Ω with a hole, we can't map it onto $\mathbb{D}$ — this would violate continuity.

We can't take the unit disk and turn it into a region with a hole, since that amounts to ripping something (which is bad for continuity). To prove this, think about the intermediate value theorem. (This is more suited to a topology class, but you can use a version of the intermediate value theorem to show this — these two regions are not homeomorphic to each other because of the hole. A conformal bijection has to be continuous, so it'll in particular give a homeomorphism between the two regions.)

So we've found two distinct obstructions — one is Liouville's theorem, and another is that every conformal mapping is a homeomorphism.

Remark 17.9. Another way to see that you can't have a hole is the residue theorem — once you have a conformal equivalence, Cauchy theory has to transfer, so in particular the fact that integrals around closed loops are 0.

The amazing thing is that these are the *only* two obstructions — the region has to miss at least one point, and it 'can't have holes.' So once we remove these two obstructions, our region is immediately conformally equivalent to $\mathbb{D}$.

To make this precise, we need to define simple connectivity. It essentially means that a region is path-connected and has no holes; but here's a complex analysis way of defining it.

Definition 17.10. A region Ω is *simply connected* if every point $z_0 \in \Omega^c = \mathbb{C} \setminus \Omega$ can be connected via a path to ∞ — i.e., for every $z_0 \in \Omega^c$ and every $\varepsilon > 0$, there exists a $\mathcal{C}^1$ path $\gamma\colon [0, \infty) \to \mathbb{C}$ such that:

(1) $\operatorname{dist}(\gamma(t), \Omega^c) < \varepsilon$ for all $t \geq 0$.

(2) $\gamma(0) = z_0$.

(3) $\lim_{t\to\infty} \gamma(t) = \infty$.

The first condition looks a bit strange; but in particular, it's sufficient if this path stays out of Ω.

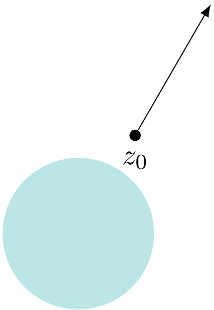

So if there's a path that lies completely within the complement, that works. But the first condition relaxes that — the path just has to be within ε of the complement.

There's an equivalent definition, which is perhaps more familiar:

Definition 17.11. A region is simply connected if any loop in Ω can be continuously deformed to a point.

The idea you should keep in your head is that simple connectivity means no holes — a loop around the hole can't be contracted to a point.

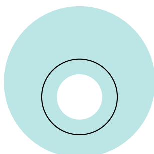

All the theory we developed for convex regions also holds for simply connected regions.

Now with this definition, we can state the Riemann mapping theorem.

Theorem 17.12 (Riemann mapping theorem)

Given any $\Omega_1, \Omega_2 \subseteq \mathbb{C}$ such that $\Omega_1, \Omega_2 \neq \mathbb{C}$ and both Ω_1 and Ω_2 are simply connected, we have $\Omega_1 \simeq \Omega_2$ (i.e., Ω_1 is conformally equivalent to Ω_2, meaning there exists $f\colon \Omega_1 \to \Omega_2$ such that f is a bijection, and both f and its inverse are conformal).

Remark 17.13. It suffices to consider the case where $\Omega_2 = \mathbb{D}$; once we know Ω_1 and Ω_2 are both conformally equivalent to $\mathbb{D}$, we can then compose the maps to get a conformal equivalence between Ω_1 and Ω_2.

So this shows that these are the *only* two obstructions.

Remark 17.14. We already showed that $\mathbb{H} \simeq \mathbb{D}$; that is a special example of this.

§17.4 Proof of Riemann mapping theorem

We'll actually prove a stronger statement that will immediately imply the theorem.

Theorem 17.15

For every simply connected region $\Omega \neq \mathbb{C}$ and every $z_0 \in \Omega$, there exists a unique conformal bijection $\varphi: \Omega \to \mathbb{D}$ such that $\varphi(z_0) = 0$ and $\varphi'(z_0) \in \mathbb{R}_{\geq 0}$.

Remark 17.16. The condition $\varphi'(z_0) \in \mathbb{R}_{\geq 0}$ is very much needed to prove uniqueness — otherwise you could just rotate.

There are some slightly slicker arguments using more sophisticated machinery; we'll just do a real analysis proof, but we need a bit of theory that might not have been covered in our real analysis class.

The strategy is that we're going to first try to find one map into the unit disk — we're going to take Ω and try to send it into the unit disk somehow. And then we'll try to maximize this somehow, to fill out the disk by expanding things — and we'll expand by maximizing φ'. We'll take the biggest derivative, and show that'll allow us to fill up the whole unit disk.

First, here's the outline: fix $z_0 \in \Omega$. Let $\mathcal{F}$ be the family of all maps $f: \Omega \to \mathbb{D}$ such that f is holomorphic in Ω, f is injective (but not necessarily surjective), and $f'(z_0) > 0$. We're going to prove three properties of this family $\mathcal{F}$.

(a) $\mathcal{F} \neq \emptyset$ (i.e., there is at least one map in $\mathcal{F}$).

(b) $\sup_{f \in \mathcal{F}} f'(z_0)$ exists and is some finite m, and it is realized — i.e., there exists $\varphi \in \mathcal{F}$ with $\varphi'(z_0) = m$.

(c) This map φ is conformal and surjective, and satisfies $\varphi(z_0) = 0$. (It is an element of $\mathcal{F}$, so it automatically satisfies the other desired properties.)

Proof of (a). Since $\Omega \neq \mathbb{C}$, there exists a point $p_0 \in \Omega^c$. We'll divide into two cases to show some intuition.

First let's suppose there is an open disk $\mathbb{D}(p_0, \delta)$ containing p_0 that's fully inside Ω^c. So here we're assuming there's a disk of radius δ fully inside the complement.

In this case, we can just set

$$f(z) = \frac{\delta}{z - p_0}.$$

Since $z \notin \mathbb{D}(p_0, \delta)$, we must have $|z - p_0| > \delta$, so this maps into the unit disk.

The slogan of the Riemann mapping theorem can be thought of as 'find a square root' — now let's think about what happens if we don't have such a disk (this is certainly possible — for example, we can take the slit plane). The trick is that if we take a square root, we can expand out this slit — if we take a square root of a slit circle, we get a semicircle.

By taking a square root, we really mean we take a *branch* of the square root — $re^{i\theta} \mapsto \sqrt{r} \cdot e^{i\theta/2}$. So angles near π get sent to $\pi/2$ — they get halved. (Specifically, the angle $\pi - \varepsilon$ is sent to $\frac{\pi-\varepsilon}{2}$.) This is the main idea — if we have this problem where our complement doesn't have a disk, we're going to open it up by taking a square root.

So now let's consider the harder case, where Ω^c has no disks.

Fact 17.17 — There exists a function g such that

$$g(z)^2 = \frac{z - p_0}{z_0 - p_0} \text{ and } g(z_0) = 1.$$

In other words, $g(z)$ is a square root of the function

$$\frac{z - p_0}{z_0 - p_0}.$$

This uses the fact that Ω is simply connected.

We proced this a month ago — we have this region Ω, and the function we're trying to square root is holomorphic in Ω; and we proved a while back that you can find a couple of different g's, by requiring $g(z_0)$ to be any root of unity. But it is nontrivial. (We proved it for convex regions, but the same argument works for simply connected regions.)

Claim 17.18 — g is bounded away from -1.

Proof. If not, there exists a sequence ζ_m such that $g(\zeta_n) \to -1$. This would imply that

$$g(\zeta_n)^2 = \frac{\zeta_n - p_0}{z_0 - p_0} \to 1$$

(since this is the square of g). But this implies that $\zeta_n \to z_0$; and since $g(z_0) = 1 \neq -1$ this violates continuity. □

So then this is nice because now we've created the picture we had before, where we have space to take a reciprocal — now we're bounded away from the point -1. So we have $\eta > 0$ such that $|g(z) + 1| > \eta$ for all $z \in \Omega$. Now set

$$f(z) = \frac{\eta}{g(z) + 1}.$$

Then $|f(z)| < 1$ for all $z \in \Omega$, so f takes Ω to $\mathbb{D}$; and it is a composition of one-to-one functions, so f is one-to-one.

Finally, we can always multiply by $e^{i\theta}$ to obtain $f'(z_0) > 0$. (We can see that f' doesn't vanish for any $z \in \Omega$, using the fact that $p_0 \notin \Omega$.)

So $f \in \mathcal{F}$, and therefore $\mathcal{F}$ is nonempty. □

Now we get to (b), which is the part involving an analysis trick.

Proof of (b). First we want to show that $f'(z_0)$ is bounded over functions in $\mathcal{F}$, so that $\sup_{f \in \mathcal{F}} f'(z_0) < \infty$.

Since Ω is open (it's a region) and $z_0 \in \Omega$, there exists some $\delta > 0$ so that the disk $\mathbb{D}(z_0, 2\delta)$ is contained in Ω. So in particular, we can use the Cauchy derivative formula for the circle of radius δ — we have

$$f'(z_0) = \frac{1}{2\pi i} \int_{|z - z_0| = \delta} \frac{f(z)}{(z - z_0)^2} \, dz.$$

And we know $|f| \le 1$ and $|z - z_0| = \delta$, so we get $|f'(z_0)| \le \frac{1}{\delta}$.

This holds for every $f \in \mathcal{F}$ (z_0 and Ω are fixed, and δ just depends on them); so this gives us an upper bound on $|f'(z_0)|$.

Now we know $|f'(z_0)|$ is bounded on $\mathcal{F}$, so we can take a maximizing sequence — we know that

$$\sup_{f \in \mathcal{F}} |f'(z_0)| = m < \infty,$$

so we can take a sequence $f_n \in \mathcal{F}$ such that $|f_n'(z_0)| \to m$ as $n \to \infty$. Our goal is to find a convergent subsequence. But we want a lot from this subsequence — we're going to obtain φ by taking a limit of a sequence, so we need uniform convergence (to ensure that φ is holomorphic). So what we want is really a convergent subsequence that converges *uniformly* on all compact (closed and bounded) subsets of Ω.

This might seem daunting; we're going to need to prove a bunch of properties about this family $\mathcal{F}$ that will allow us to get this (in particular *equicontinuity*). We'll do this next class. □

§18 December 5, 2023

Definition 18.1. We say two regions $U, V \subseteq \mathbb{C}$ are *conformally equivalent* (written $U \simeq V$) if we can find $f\colon U \to V$ and $g\colon V \to U$ such that f and g are both holomorphic, and are inverses of each other — i.e., $f(g(z)) = z$ and $g(f(z)) = z$.

We're in the middle of proving the Riemann mapping theorem.

Theorem 18.2 (Riemann mapping theorem)

Let $\Omega_1, \Omega_2 \subseteq \mathbb{C}$ be regions such that Ω_1 and Ω_2 are simply connected and both not equal to $\mathbb{C}$. Then $\Omega_1 \simeq \Omega_2$.

We're taking one of these regions to be the unit disk $\mathbb{D}$ (which satisfies these constraints); we're trying to prove that if $\Omega \neq \mathbb{C}$ is simply connected, then $\Omega \simeq \mathbb{D}$. We're proving the following stronger statement:

Theorem 18.3

If $\Omega \neq \mathbb{C}$ is simply connected, then for all $z_0 \in \Omega$, there exists a unique conformal bijective map $\varphi\colon \Omega \to \mathbb{D}$ such that $\varphi(z_0) = 0$ and $\varphi'(z_0) \in \mathbb{R}_{>0}$.

Last time, we considered the family of functions

$$\mathcal{F} = \{f\colon \Omega \to \mathbb{D} \mid f \text{ is holomorphic and 1-to-1, and } f'(z_0) > 0\}.$$

The proof consists of three steps:

(a) $\mathcal{F}$ is nonempty — we proved this last class.

(b) Then we can consider $\sup_{f \in \mathcal{F}} f'(z_0)$. We want to show that this supremum exists and is some finite number m, and there exists $\varphi \in \mathcal{F}$ such that $\varphi'(z_0) = m$ (i.e., this supremum is realized).

(c) Finally, we'll show that this function φ is our desired conformal mapping — we'll show that φ is conformal, $\varphi(z_0) = 0$, $\varphi'(z_0) > 0$ (these statements are automatic, since $\varphi \in \mathcal{F}$), and that $\varphi\colon \Omega \to \mathbb{D}$ is onto.

Last time we proved (a), and started (b). We'll do this the only way you might imagine — we'll take a maximizing sequence and try to find a convergent subsequence.

Last time we also showed that $\sup_{f \in \mathcal{F}} f'(z_0)$ is finite; this followed from the Cauchy derivative formula.

Proof of (b). Let $m = \sup_{f \in \mathcal{F}} f'(z_0)$, and let $\{f_n\}$ be a sequence in $\mathcal{F}$ such that $f_n'(z_0) \to m$ as $n \to \infty$. (The definition of a supremum means that there exists such a sequence.)

Our goal is to find a subsequence of $\{f_n\}$ converging *uniformly* on all compact subsets $K \subseteq \Omega$; then we know the limit function is holomorphic, and that'll be our candidate φ.

We're going to appeal to a general procedure for doing this — we'll do a diagonalization argument. First, let $\{\zeta_m\}$ be a countable dense subset of Ω (for example, we can take the set of points in Ω with rational coordinates — the set of all rational points in $\mathbb{R}^2$ is countable, so we can enumerate them in some order).

Now we can look at the sequence $\{f_n(\zeta_1)\}_{n=1}^{\infty}$. This is some sequence of numbers inside $\mathbb{D}$, so it's bounded; in particular it has a convergent subsequence. So we can find a subsequence $\{f_{1n}\}$ of our functions such that $f_{1n}(\zeta_1)$ converges to some value; let its limit be $\varphi(\zeta_1)$.

So that's our first step. We're now going to do this for every subsequent point — but restricting to a subsequence every point. We found a subsequence of the f_n's here; now we throw away the rest, just look at the subsequence we extracted, and evaluate that subsequence at ζ_2. And we keep on doing this.

So we can look at the sequence $\{f_{1n}(\zeta_2)\}$, and we can find a convergent subsequence consisting of the functions $\{f_{2n}\}$; call the limit $\varphi(\zeta_2)$. And we go on doing this — we continue doing this forever, and we eventually obtain a sequence $\{\{f_{kn}\}_{n=1}^{\infty}\}_{k=1}^{\infty}$ such that the sequence $\{f_{kn}\}_{n=1}^{\infty}$ converges at each of $\zeta_1, \ldots, \zeta_k$ (to $\varphi(\zeta_1), \ldots, \varphi(\zeta_k)$).

Now imagine that we put these functions in a table, in the obvious way.

$$\begin{matrix} f_{11} & f_{21} & f_{31} & f_{41} & f_{51} \\ f_{12} & f_{22} & f_{32} & f_{42} & f_{52} \\ f_{13} & f_{23} & f_{33} & f_{43} & f_{53} \\ f_{14} & f_{24} & f_{34} & f_{44} & f_{54} \\ f_{15} & f_{25} & f_{35} & f_{45} & f_{55} \end{matrix}$$

And the trick is that we single out the diagonal of this table. The point is that we have all these subsequences which converge, and we extract this diagonal sequence — let $\varphi_n = f_{nn}$. Then $\varphi_n(z)$ converges for every $z \in \{\zeta_m\}_{m=1}^{\infty}$ by construction — because each of the further subsequences is extracted from the previous one.

> **Remark 18.4.** Do we need to worry about the fact that the limit might not be in the unit disk — i.e., it might drift to the boundary?
>
> We'll learn some more about φ first, and maybe that'll deal with this issue. (One way to see this is from the holomorphicity of φ, which we'll prove soon; images of open sets are open.)

So we've found this convergent subsequence, but only on this countable dense subset; the first thing we need to do is extend to the rest of Ω and show that it converges uniformly on every compact subset.

> **Claim 18.5** — The sequence φ_n converges on Ω, and converges uniformly on every compact $K \subseteq \Omega$.

Proof. It's a fact that any compact $K \subseteq \Omega$ is contained in a finite union of closed disks; so we can assume that K is a closed disk.

Since K is compact and Ω is open, the distance between K and Ω^c must be positive — so we can let

$$\operatorname{dist}(K, \Omega^c) = 2d > 0.$$

Then since $|\varphi_m| \leq 1$, we know by the Cauchy derivative formula on the contour $|\zeta - z| = d$ that

$$|\varphi_m'(z)| = \frac{1}{2\pi i} \int_{|\zeta - z| = d} \frac{\varphi_n(\zeta)}{(\zeta - z)^2} \, d\zeta,$$

which means

$$|\varphi'(z)| \leq \frac{1}{2\pi} \cdot \frac{2\pi d}{d^2} = \frac{1}{d}.$$

So we have a uniform bound on the derivatives — this is true for *every* $z \in K$ (and it's uniform in n and z).

> **Remark 18.6.** The functions φ_n are a subsequence of our original $\{f_n\}$, and the f_n's are all holomorphic — all we did was pass to nested subsequences repeatedly and then choose a specific diagonal one, so the φ_n all come from the original sequence and are therefore holomorphic.

This means the φ_n are an 'equicontinuous family,' meaning the following:

> **Claim 18.7** — For every $\varepsilon > 0$ and for every n,
>
> $$|\varphi_n(z_1) - \varphi_n(z_2)| \leq \varepsilon \text{ as long as } |z_1 - z_2| \leq \varepsilon d.$$

Proof. We can write

$$|\varphi_n(z_1) - \varphi_n(z_1)| = \left| \int_{[z_1, z_2]} \varphi_n'(z) \right| \leq |z_1 - z_2| \cdot \sup |\varphi_n'(z)| \leq \frac{|z_1 - z_2|}{d}.$$ □

This statement essentially means that these functions are not just continuous, but they're all continuous with the same δs. And now once we have convergence on a countable dense subset, we get convergence everywhere from this relationship — for any $z \in K$, we have

$$|\varphi_n(z) - \varphi_m(z)| \leq |\varphi_n(z) - \varphi_n(\zeta_k)| + |\varphi_n(\zeta_k) - \varphi_m(\zeta_k)| + |\varphi_m(\zeta_k) - \varphi_m(z)|$$

(where ζ_k comes from our countable dense family). We can bound the first and last term using equicontinuity, and the fact that for each fixed k, $\varphi_n(\zeta_k)$ is a Cauchy sequence in n — so we can first choose ζ_k such that $|\zeta_k - z| < \frac{1}{3}\varepsilon d$, and then we can choose m and n large enough such that the middle term is small as well.

So then $\{\varphi_n(z)\}_{n=1}^{\infty}$ is a Cauchy sequence, and is therefore convergent; this gives us pointwise convergence. And we also get that φ is continuous — because

$$|\varphi(z_1) - \varphi(z_2)| = \lim_{n \to \infty} |\varphi_n(z_1) - \varphi_n(z_2)| < \varepsilon \text{ if } |z_1 - z_2| < \varepsilon d.$$

So as long as $|z_1 - z_2| < \varepsilon d$, then equicontinuity gives us continuity of the limit function too.

This gives convergence and continuity, but we need *uniform* convergence. Fix $\varepsilon > 0$, and for each j, let

$$S_j = \{z \in K \mid |\varphi_n(z) - \varphi(z)| < \varepsilon \text{ for all } n > j\}.$$

What are some properties of the S_j? On one hand, if we take a union of all of them, then we cover K (since $\varphi_n \to \varphi$ pointwise). So $K \subseteq \bigcup_j S_j$. Furthermore, the sets S_j are open by equicontinuity. (The point is essentially that this is an 'open condition.') So we have a covering of K by open sets, which since K is compact means we can pass to a finite subcover — this means we can find N such that $K \subseteq \bigcup_{j=1}^{N} S_j$. Then we have

$$|\varphi_n(z) - \varphi(z)| < \varepsilon \text{ as long as } n > N,$$

which gives us uniform convergence. □

And now we're in business — $\varphi_n \to \varphi$ uniformly, which implies that φ is holomorphic in Ω. And now by the maximum principle, either φ is constant, or φ maps $\Omega \to \mathbb{D}$ (i.e., it only takes values with magnitude strictly less than 1).

To see that φ is not constant, we defined φ such that $\varphi'(z_0) = m > 0$; this means φ is nonconstant. Then the maximum principle implies that $\varphi\colon \Omega \to \mathbb{D}$. (Otherwise, if φ mapped some z onto the boundary of $\mathbb{D}$, then it'd be attaining a maximum at z.)

Remark 18.8. In more detail, last problem set we proved that if a holomorphic function takes its maximum on the interior of a set, then it has to be constant. Here when we're showing that φ maps Ω to $\mathbb{D}$, the worry is that we found φ by taking a limit of functions with values in $\mathbb{D}$; but since we took limits, we might worry that $|\varphi(z)| = 1$ for some z (we might have had a sequence $\varphi_n(z)$ tending to the boundary of the unit disk, which would cause $\varphi(z)$ to be on the boundary).

And we need to rule out this possibility. The way we rule it out is by first proving φ is holomorphic. And now it's a holomorphic function on this open set Ω; if it were to take its maximum inside, then it would be a constant function. And it can't be the constant function, since $\varphi'(z_0) > 0$.

Finally, φ is one-to-one because it is a uniform limit (on all compact $K \subseteq \Omega$) of one-to-one functions. (This is nontrivial, but we will not go through the details.)

This finishes the proof of (b). □

Proof of (c). We first need to show that $\varphi(z_0) = 0$. Assume not. If $\varphi(z_0) = \alpha$ for some $0 < |\alpha| < 1$, then we can consider the function

$$f(z) = \frac{\varphi(z) - \alpha}{1 - \overline{\alpha}\varphi(z)}.$$

This map is one-to-one (we've composed φ with a conformal automorphism of $\mathbb{D}$) and holomorphic, and f also takes $\Omega \to \mathbb{D}$ (since the map $\psi\colon z \mapsto \frac{z-\alpha}{1-\overline{\alpha}z}$ takes $\mathbb{D} \to \mathbb{D}$).

But now we have

$$f'(z) = \frac{\varphi'(z)}{1 - |\alpha|^2},$$

so in particular $f'(z_0) > \varphi'(z_0)$, which is a contradiction.

That tells us $\varphi(z_0) = 0$. Now why is φ onto? We'll again argue by contradiction — this proof is more involved, but we'll again try to break maximality. Assume φ is not onto; then there exists $w \in \mathbb{D}$ such that $\varphi(z) \neq w$ for all $z \in \Omega$. Then we can write w as $w = -t^2 e^{i\theta}$ for some $0 < t < 1$ — we know $w \neq 0$, because φ sends z_0 to 0.

Now we want to build a function in the same fashion — set $g(z) = e^{-i\theta}\varphi(z)$, so that $g(z) \neq -t^2$ for all $z \in \Omega$ (this is to get rid of the phase). And now set

$$f_1(z) = \frac{g(z) + t^2}{1 + t^2 g(z)}.$$

Then f_1 takes Ω to $\mathbb{D}$, and $f_1(z_0) = t^2$; but since g is never equal to $-t^2$, this function doesn't hit 0. So we can find a square root — since $g(z) \neq -t^2$, we know $f_1(z) \neq 0$ for all $z \in \Omega$, which means there exists a branch f_2 of the square root such that $f_2(z)^2 = f_1(z)$ and $f_2(z_0) = t$.

Finally, set

$$f_3(z) = \frac{f_2(z) - t}{1 - t f_2(z)}.$$

Then f_3 is one-to-one. We can now compute a bunch of derivatives:

- $f_1'(z_0) = g'(z_0)(1 - t^4)$,

- We have
$$f_2'(z_0) = \frac{f_1'(z_0)}{2\sqrt{f_1(z_0)}} = \frac{f_1'(z_0)}{2t}.$$
- Finally, this implies that
$$f_3'(z_0) = \frac{f_2'(z_0)}{1-t^2} = \frac{g'(z_0)(1+t^2)}{2t}.$$

And now we've got a contradiction — we know $1+t^2 > 2t$ (since $t \neq 1$), and if we set $f(z) = e^{i\theta} f_3(z)$ (undoing the rotation), then we get $f \in \mathcal{F}$ and $f'(z_0) > \varphi'(z_0)$, which is a contradiction.

So φ is onto. □

Remark 18.9. Did we prove the uniqueness of φ? No.

Proof of uniqueness. If there exist two such φ_1 and φ_2, then the map $\Phi = \varphi_1 \circ \varphi_2^{-1}$ is a map $\mathbb{D} \to \mathbb{D}$ such that $\Phi(0) = 0$ and $\Phi'(0) > 0$, which implies that $\Phi(z) = e^{i\theta} z$ (by a homework problem); then $\Phi'(0) = e^{i\theta}$. But this is only a positive real number if $e^{i\theta} = 1$, and therefore Φ is the identity. □

Remark 18.10. Did we show that φ is conformal? No.

So now we're done with the proof of the Riemann mapping theorem. The downside of the proof is that we know nothing about what the Riemann mapping (the conformal bijection) actually looks like. One line of thought you might pursue is, can we construct the Riemann maps for certain domains? There's a whole business of this, but we won't get into that.

§18.1 Analytic continuation

Fact 18.11 — For a region Ω, if $f \equiv 0$ on an open set $D \subseteq \Omega$, then $f \equiv 0$ on all of Ω.

(We discussed this earlier in the class.)

Question 18.12. Suppose we have a holomorphic function f on some region Ω. What is its largest domain of holomorphicity (or meromorphicity)?

What do we mean by this? Imagine f is defined on some region; the question is, can we extend it to a larger region? Can we find $\Omega_2 \supseteq \Omega$ such that there exists g holomorphic on Ω_2 with $g = f$ on Ω?

Example 18.13

The Γ function is defined as
$$\Gamma(s) = \int_0^\infty e^{-t} t^{s-1}\, dt.$$
For which $s \in \mathbb{C}$ does this function make sense?

First, when does this integral converge? This integral converges for all s with $\mathrm{Re}(s) > 0$. To see this, the behavior near ∞ is great — we have an exponential compared to a polynomial — so we only need to worry about what happens near 0. Let $s = \alpha + i\beta$, so that
$$t^{s-1} = t^{\alpha-1} t^{i\beta}.$$

As long as $\alpha > 0$, this integral makes sense at 0 as well.

So this means $\Gamma(s)$ is convergent and holomorphic for $\mathrm{Re}(s) > 0$.

Theorem 18.14

The function $\Gamma(s)$ admits a meromorphic continuation to $\mathbb{C}$ with simple poles at $s = 0, -1, -2, -3, \ldots$.

We'll do this next time — we'll try to find a function which agrees with Γ on the right half of the plane and is meromorphic on $\mathbb{C}$, with these poles. And after that we'll look at the ζ function.

§19 December 7, 2023

§19.1 The Riemann mapping theorem

We'll now fill in the gap in our proof of the Riemann mapping theorem regarding the injectivity of φ.

Last class, we had a sequence $\varphi_n \colon \Omega \to \mathbb{D}$ such that $\varphi_n \to \varphi$ uniformly on each compact $K \subseteq \Omega$, and the functions φ_n are holomorphic and one-to-one.

Lemma 19.1

For any such sequence φ_n, either φ is one-to-one or φ is constant.

Remark 19.2. The specific φ we had in our proof of the Riemann mapping theorem satisfied $\varphi'(z_0) > 0$, so φ is not constant.

Proof. Suppose that φ is not constant and not one-to-one; we'll try to derive a contradiction. Since φ is not one-to-one, we can find $z_1, z_2 \in \Omega$ such that $\varphi(z_1) = \varphi(z_2)$. The key to this proof is that we'll use the uniformity of the limit together with the zero-counting formula — uniformity of the limit allows us to the limit inside integrals.

We want to rig this to a statement about the zeros of some function, so we define the auxiliary sequence g_n of functions defined as $g_n(z) = \varphi_n(z) - \varphi_n(z_1)$. We're just subtracting off a sequence of points from the old sequence of functions, so the convergence properties between g_n and φ_n are the same.

But now we know that $g_n(z_1) = 0$ by construction; since the φ_n are all one-to-one functions, this is the only zero of each g_n (for all n). So for every n, this function just has a single zero.

But $g_n(z) \to g(z) = \varphi(z) - \varphi(z_1)$, and this convergence is uniform on compact subsets of Ω. And g has (at least) two zeros — we have $g(z_1) = g(z_2) = 0$.

It could be that g is identically equal to 0; then φ is a constant function. But we're assuming φ is not constant (in our contradiction argument), so g is not identically 0.

So the zero at z_2 is isolated; this means we can find a small neighborhood around z_2 in which g has no other zeros. In particular, we can find a small loop γ around z_2 such that g has no zeros on γ, and a unique zero inside γ (namely z_2). Then by the zero counting formula we have

$$1 = n(g, \gamma) = \frac{1}{2\pi i} \int_\gamma \frac{g'(w)}{g(w)}\, dw.$$

But on the other hand, the g_n have just a single zero at g_1, so if we take the same loop around z_2, we have

$$0 = n(g_n, \gamma) = \frac{1}{2\pi i} \int_\gamma \frac{g_n'(w)}{g_n(w)}\, dw.$$

But now we get a contradiction — the g_n are bounded away from zero on γ, and $g_n \to g$ uniformly on γ, so $\frac{1}{g_n} \to \frac{1}{g}$ uniformly on γ; and by the Cauchy derivative formula $g_n' \to g'$ uniformly on γ as well. This means $\frac{g_n'}{g} \to \frac{g'}{g}$ uniformly on γ. And now we get a contradiction — the uniformity of the limit implies that

$$0 = \lim_{n\to\infty} \frac{1}{2\pi i} \int_\gamma \frac{g_n'(w)}{g_n(w)}\,dw = \frac{1}{2\pi i}\int_\gamma \frac{g'(w)}{g(w)}\,dw = 1.$$

This is a contradiction. □

We also didn't prove that φ is conformal; but that's actually a consequence of injectivity.

Lemma 19.3

If φ is holomorphic in Ω and one-to-one, then $\varphi'(z) \neq 0$ for all $z \in \Omega$.

So if we have a holomorphic function that's one-to-one, then it's automatically conformal.

Proof. This can be proven by contradiction — assume we have a point where the derivative vanishes. Then using the power series expansion, if $\varphi'(z_0) = 0$ then near z_0 we can write

$$\varphi(z) - \varphi(z_0) = a(z - z_0)^k + G(z)$$

where $k \geq 2$ and z_0 is a zero of G of order $k+1$. Then you can write $\varphi(z) - \varphi(z_0) - w$, for very small w, as $F(z) + G(z)$ where F is much bigger than G; then you can use Rouche's theorem to show that $F + G$ has to have k zeros (since F does). □

§19.2 The Γ function

Definition 19.4. We define $\Gamma(s) = \int_0^\infty e^{-t}t^{s-1}\,dt$ for $\mathrm{Re}(s) > 0$.

Last class we showed that this integral does converge for $\mathrm{Re}(s) > 0$. We then stated the following theorem, without proof — this function is originally defined only on the right half of the complex plane, but we can try to extend it to the *whole* complex plane. We'll lose holomorphicity, but we can find a *meromorphic* function that agrees with it on the whole complex plane.

Theorem 19.5

$\Gamma(s)$ admits a meromorphic extension to $\mathbb{C}$ with simple poles at $s = 0, -1, -2, -3, \ldots$.

Proof. We'll try to find a nice formula for Γ that makes sense even past $\mathrm{Re}(s) = 0$. We'll do this by repeatedly shifting the domain using integration by parrts — if $\mathrm{Re}(s) > 0$ then

$$\Gamma(s) = \int_0^\infty e^{-t}t^{s-1}\,dt = \int_0^\infty e^{-t}\cdot\frac{d}{dt}\left(\frac{1}{s}t^s\right)dt = \frac{1}{s}\left([e^{-t}t^s]_0^\infty + \int_0^\infty\right)e^{-t}t^s\,dt = \frac{1}{s}\int_0^\infty e^{-t}t^s\,dt = \frac{1}{s}\Gamma(s+1).$$

Note that the final integral converges as long as $\mathrm{Re}(s) > -1$ — we've introduced a pole at $s = 0$, and we're multiplying that by a holomorphic function.

So we've developed this formula — we have $\Gamma(s+1) = s\Gamma(s)$, which gives an extension of Γ to $\mathrm{Re}(s) > -1$.

And then we can keep going — for s to the left, we can iterate this process and write

$$\Gamma(s) = \frac{1}{s}\cdot\frac{1}{s+1}\cdot\Gamma(s+2).$$

The right-hand side is now defined for $\mathrm{Re}(s) > -2$, and now we have poles at $s = 0$ and $s = -1$. And we can iterate forever, to define $\Gamma(s)$ on all of $\mathbb{C}$ (with poles at nonpositive integers). □

Exercise 19.6. The residue of Γ at $s = 0$ is 1, and $\operatorname{Res}(\Gamma, n) = (-1)^n (n!)^{-1}$.

We see a beautiful relationship — if we evaluate this formula $\Gamma(s+1) = s\Gamma(s)$ at the integers, we can note that $\Gamma(1) = \int_0^\infty e^{-t}\,dt = 1$, and so then $\Gamma(n+1) = n!$. Because of this formula, this function pops up a lot in number theory.

So what we've done is found a new function; we can compute its values using these formulas everywhere on the complex plane, and this agrees with the function we wrote down in the beginning on the right half of the plane.

Here's one more fact about the Γ function.

Fact 19.7 — The function $\Gamma(s)^{-1}$ is entire.

We'll show this using the following identity.

Claim 19.8 (Functional identity for Γ) — We have

$$\frac{1}{\Gamma(s)} = \frac{\sin \pi s}{\pi} \cdot \Gamma(1-s).$$

This gives a functional identity for the Γ function. Then $\Gamma(1-s)$ is meromorphic, with simple poles at $1 - s = -n$ for nonnegative integers n; so it has simple poles whenever s is a positive integer. But then $\frac{\sin \pi s}{\pi}$ has simple zeros on the positive integers, and those cancel out the poles. That proves $\Gamma(s)^{-1}$ is entire; and it has zeros at $s = 0, -1, -2, \ldots$.

Proof. We can write

$$\Gamma(s)\Gamma(1-s) = \int_0^\infty e^{-t} t^{s-1}\,dt \cdot \int_0^\infty e^{-r} r^{1-s-1}\,dr,$$

and $1 - s - 1 = -s$. The first integral converges for $\operatorname{Re}(s) > 0$ and the second integral converges for $\operatorname{Re}(s) < 1$, and so there's an overlap where they both converge, and this formula makes sense on the overlap $0 < \operatorname{Re}(s) < 1$. So let's think about it there.

We have a product of integrals, so we can think of it as a double integral

$$\Gamma(s)\Gamma(1-s) = \int_0^\infty \int_0^\infty e^{-(t+r)} \left(\frac{t}{r}\right)^{s-1} \cdot r\,dt\,dr.$$

And now we can do a change of variables — substitute $u = t + r$ and $v = \frac{t}{r}$ (unsurprisingly). Then we can compute the Jacobian

$$du\,dv = |\det \mathbf{J}|\ dt\,dr,$$

where the Jacobian is

$$\mathbf{J} = \begin{bmatrix} \partial u/\partial t & \partial u/\partial r \\ \partial v/\partial t & \partial v/\partial r \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 1/r & -t/r^2 \end{bmatrix},$$

which has determinant

$$\det \mathbf{J} = -\frac{1}{r^2}(t + r).$$

We want to write this in terms of u and v instead; and we can compute that

$$\frac{1}{r}\,dt\,dr = \frac{1}{v+1}\,du\,dv.$$

This means our original integral is equal to

$$\int_0^\infty \int_0^\infty e^{-u} v^{s-1} \cdot \frac{1}{v+1}\,du\,dv = \int_0^\infty \frac{v^{s-1}}{v+1}\,dv.$$

This can be found using contour integration; it's a challenging exercise, and a first example of what's called a *keyhole contour*. We want to evaluate an integral from 0 to ∞; what do we need to avoid? For $s \in \mathbb{C}$ we define v^{s-1} using log, so we need to make sure we have some contour where we can avoid log. So we use the following contour:

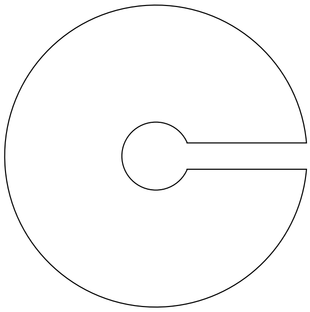

Now we can turn the function we're integrating into

$$z^{s-1} = e^{(s-1)\log z}. \qquad \square$$

So these are some amazing computations. Remarkably, they were done by Riemann.

Remark 19.9. Why is analytic continuation unique? If we have two holomorphic functions and they agree on an open set (or even on enough spots), then they have to agree everywhere. (Their difference would have to be the zero function.)

§19.3 The Riemann ζ function

Definition 19.10. We define $\zeta(s) = \sum_{n=1}^{\infty} \frac{1}{n^s}$.

This series converges for all $\mathrm{Re}(s) > 1$. Our goal is to find a meromorphic continuation to $\mathbb{C}$. We know we'll have a problem at $s = 1$ — we get the harmonic series. So the extension we'll find will have a pole at $s = 1$; in fact that will be the only pole.

There are two ways. The first, as done by Riemann, is to make use of the following fact.

Claim 19.11 — We have

$$n^{-s} = \frac{1}{\Gamma(s)} \int_0^{\infty} e^{-nt} t^{s-1}\, dt.$$

Proof. The right-hand side looks like the Γ function except for the n, so let's get rid of it — let $u = nt$ so that $du = ndt$. Then we get

$$\frac{1}{\Gamma(s)} \int_0^{\infty} e^{-u} \left(\frac{u}{n}\right)^{s-1} \cdot \frac{1}{n}\, du = n^{-s} \cdot \frac{1}{\Gamma(s)} \int_0^{\infty} e^{-t} t^{s-1}\, dt = n^{-s}. \qquad \square$$

Now we can sum over n on the right-hand side, and we get

$$\zeta(s) = \sum_{n=1}^{\infty} n^{-s} = \frac{1}{\Gamma(s)} \int_0^{\infty} \left(\sum_{n=1}^{\infty} e^{-nt}\right) t^{s-1}\, dt.$$

And we have a geometric series — we can write

$$\sum_{n=1}^{\infty} u^n = u(1 + u + u^2 + \cdots) = \frac{u}{1-u}$$

(note that $t > 0$, so the geometric series is convergent), so this is equal to

$$\zeta(s) = \frac{1}{\Gamma(s)} \int_0^\infty \frac{e^{-t}}{1 - e^{-t}} t^{s-1}\, dt = \frac{1}{\Gamma(s)} \int_0^\infty \frac{t^{s-1}}{e^t - 1}\, dt.$$

(This was proved by Riemann, among a slew of other things.)

Then Riemann made a ton of observations about the ζ function. Our first goal is to find an analytic continuation, and we'll do that by examining this integral.

First, $\frac{1}{\Gamma(s)}$ is entire, so that factor is great. And then we have this integral; where are its problems? We don't have any problems from t near ∞ — there we have an e^t factor in the denominator, which becomes really big (we don't care what s is). The issue is for t near 0. So we'll split this into two pieces, as

$$\zeta(s) = \frac{1}{\Gamma(s)} \int_0^1 \frac{t^{s-1}}{e^t - 1}\, dt + \int_1^\infty \frac{t^{s-1}}{e^t - 1}\, dt.$$

And the integral on the right is an entire function. So getting the meromorphic continuation really depends on the first part of the integral, which we'll now look at in more detail (since the second piece has no issues).

We'll prove a slightly more general statement. We can think of

$$e^t - 1 = \int_0^t e^{-r}\, dr = t\Phi(t)$$

where $\Phi \in \mathcal{C}^\infty$ is never zero. So then we can write

$$\frac{1}{e^t - 1} = \frac{1}{t}\psi(t),$$

where $\psi = \frac{1}{\Phi}$ is also $\mathcal{C}^\infty$. We can forget about the $\frac{1}{\Gamma(s)}$ factor; then we can write

$$\int_0^1 \frac{t^{s-1}}{e^t - 1}\, dt = \int_0^1 t^{s-2}\psi(t)\, dt,$$

where $\psi \in \mathcal{C}^\infty$.

Lemma 19.12

For any $\mathcal{C}^\infty$ function ψ, the integral $\int_0^1 t^{s-2}\psi(s)\, dt$ is holomorphic for $\mathrm{Re}(s) > 1$, and can be meromorphically continued to all of $\mathbb{C}$ with simple poles at $s = 1, 0, -1, -2, \ldots$.

Once we prove this lemma, we can go back to $\frac{1}{\Gamma}$, which we proved has zeros at $0, -1, -2, \ldots$; this cancels out nearly all the poles, leaving just a single pole at $s = 1$.

Proof of lemma. Let $N \in \mathbb{N}$, and apply the Taylor formula to $\psi(t)$; then we can write

$$\psi(t) = \psi(0) + t\psi'(0) + \cdots + \frac{1}{N!} t^N \psi^{(N)}(0) + t^{N+1} E(t),$$

where $E(t) \in \mathcal{C}^\infty$. We'll pursue the same type of strategy as for the Γ function — if $\mathrm{Re}(s) > 1$ (which is where this initially makes sense), we have

$$\int_0^1 t^{s-2}\psi(t)\, dt = \int_0^1 t^{s-2} \left(\psi(0) + t\psi'(0) + \cdots + \frac{1}{N!} t^n \psi^{(N)}(0)\right) dt + \int_0^1 t^{s+N-1} E(t)\, dt.$$

The last piece has problems when $\text{Re}(s) \le -N$; so this last bit is holomorphic on $\{\text{Re}(s) > -N\}$ (because then we avoid $\frac{1}{t}$). And now we're in business — we can just integrate the other things, and those give us monomial integrals.

So we can write

$$\int_0^1 t^{s-2}\left(\psi(0) + t\psi'(0) + \cdots + \frac{t^N}{N!}\psi^{(N)}(s)\right) dt$$

(for $\text{Re}(s) > 1$) as

$$\psi(0)\left[\frac{1}{s-1}t^{s-1}\right]_0^1 + \psi'(0)\left[\frac{1}{s}t^s\right]_0^1 + \cdots + \frac{\psi^{(N)}(0)}{N!}\left[\frac{1}{s+N-1}t^{s+N-1}\right]_0^1,$$

which simplifies to

$$\frac{\psi(0)}{s-1} + \frac{\psi'(0)}{s} + \cdots + \frac{\psi^{(N)}(0)}{N!}\frac{1}{s+N-1}.$$

And this is a rational function, with poles at $s = 1$, $s = 0$, ..., $s = -N+1$.

So we did this computation; it's a rational function that agrees with our original integral on $\text{Re}(s) > 1$, and this function is defined everywhere.

And now we're in business — we've found an analytic continuation all the way up to $\text{Re}(s) > -N$ (where the $\text{Re}(s) > -N$ bound comes from the remainder term, which was only defined there). This means

$$\zeta(s) = \frac{1}{\Gamma(s)}\left(\text{holomorphic in } \text{Re}(s) > -N + \frac{\psi(0)}{s-1} + \cdots + \frac{\psi^{(N)}(0)}{N!}\frac{1}{s+N-1} + \int_1^\infty \frac{t^{s-1}}{e^t - 1}\, dt\right)$$

(where the last term is holomorphic in $\mathbb{C}$). And finally the zeros of $\frac{1}{\Gamma(s)}$ at $s = 0, -1, -2¡\ldots$ cancel the poles of our rational function at those points. This means every pole gets cancelled except the pole at $s = 1$.

And this gives an extension to $\text{Re}(s) > -N$; no matter what s is, we can find an integer N with $\text{Re}(s) > -N$, and use this to define the meromorphic continuation there. So ζ extends to all of $\mathbb{C}$ with only a simple pole at $s = 1$. □

§19.4 The Riemann hypothesis

The Riemann hypothesis is probably one of the most famous open problems in math.

> **Conjecture 19.13** (Riemann hypothesis) — All zeros of $\zeta(s)$ have the form $s = \frac{1}{2} + it$ for $t \in \mathbb{R}$.

You can show that the zeros have to lie inside the strip $0 < \text{Re}(s) < 1$; and the Riemann hypothesis states that all are on the line $\text{Re}(s) = \frac{1}{2}$.

One of the reasons the ζ function is so important (and why this is such an important problem) is that the ζ function is related to the distribution of the prime numbers, via the Euler formula — we can write $\zeta(s)$ as

$$\zeta(s) = \prod_p \frac{1}{1 - p^{-s}}.$$

And the tongue-in-cheek way of discussing the importance of the Riemman hypothesis is that if it's false, then it'd lead to chaos in number theory (it'd destroy some reasonable things people expect about the distribution of primes).

> **Remark 19.14.** Prof. Gallagher (Prof. Lawrie's teacher when he was an undergraduate) worked for a long time on the Riemann hypothesis, and made significant improvements trying to zero in on this line — he improved the known bounds on where the zeros might lie.

Next class Prof. Lawrie will show us a cool thing (that might really help us understand what analytic continuation means) — we'll see a different way to meromorphically continue the ζ function, and that'll give us a nice formula that'll allow us to compute certain values of it. Of course, the continuation has to agree with the one we've found already; but we'll find anothe rformula extending $\zeta(s)$ to $\mathbb{C}$, and this formula will let us compute specific values — for example $\zeta(-1)$. (We know ζ only has a pole at -1, so this is a perfectly valid thing to do.) This spits out $-\frac{1}{12}$; and that's the sense in which $\sum n = -\frac{1}{12}$.

§20 December 12, 2023

The final exam is on Gallagher 1–20 and the first page of Gallagher 21 (specifically, the theorem on the first page, which we did cover in class), as well as Möbius transformations and the Riemann mapping theorem — in particular, the last 3ish classes covering the Γ and ζ function will not be in the final exam. The format of the final will be similar to the midterm, though maybe longer. (We'll have 3 hours, but Prof. Lawrie intends for the exam to be shorter if we know all the problems.)

Last time, we came to the formula

$$\zeta(s) = \frac{1}{\Gamma(s)} \int_0^\infty \frac{t^{s-1}}{e^t - 1}\, dt.$$

We studied the convergence properties of this integral, in particular near $t = 0$, to see that ζ extends to a meromorphic function on $\mathbb{C}$ with a simple pole at $s = 1$. For $\operatorname{Re}(s) > 1$, we had the original representation

$$\zeta(s) = \sum_{n=1}^\infty \frac{1}{n^s},$$

and when $s \to 1$ we get the harmonic series, which diverges; that explains the pole at $s = 1$. But other than that we have this representation for ζ that is valid on all of $\mathbb{C}$.

(The integral has a pole at $s = 1$; it also has other poles, but those are cancelled by the zeros of $\Gamma(s)^{-1}$, which are at $0, -1, \ldots$.)

Today we'll see another way to view the analytic continuation of ζ.

§20.1 Another analytic continuation

For now, assume $\operatorname{Re}(s) > 1$. We're going to find a different formula for the ζ function which agrees with the previous formula when $\operatorname{Re}(s) > 1$, but is valid on all of $\mathbb{C}$.

Consider the following contour in $\mathbb{C}$, which we call c.

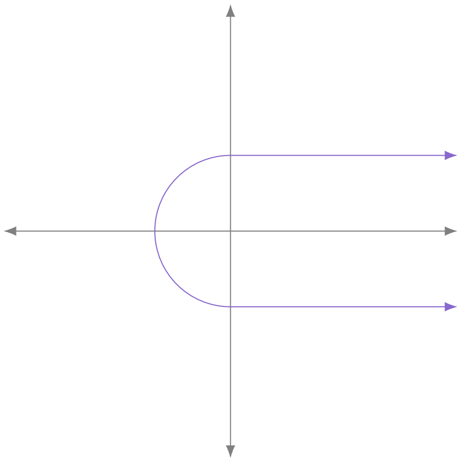

Then we can consider $\int_c \frac{z^{s-1}}{e^z-1}\,dz$. Implicitly we're imagining that the contour extends off to ∞ to the right; we can do this because the integral is convergent along the contour as we go to the right ($|e^z|$ goes to 0 very fast). And the contour avoids 0, so the function is well-defined everywhere on it.

In particular, this function $\int_c \frac{z^{s-1}}{e^z-1}\,dz$ is holomorphic in s for any $s \in \mathbb{C}$.

Now we're going to try to deform this contour — we'll try to get a different formula for ζ involving this contour c. First, we truncate the contour at some R, and then take it all the way down close to the real axis, drawing a tiny disk of radius ε around 0.

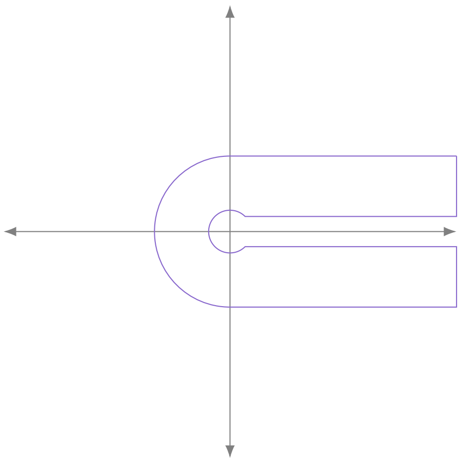

Remark 20.1. For each fixed s, we define z^{s-1} by taking a suitable logarithm — $e^{(s-1)\log z}$.

Let D be the region inside this contour, and ∂D our contour (i.e., the boundary of D). Then by Cauchy theory we have

$$\int_c \frac{e^z}{z^s-1}\,dz = 0$$

(since for fixed s, this is holomorphic as a function of z). Now we can break this into pieces — letting c_R be the truncated version of c, we get an integral over c_R, an integral over the two vertical pieces, an integral over the size-ε loop, and the integrals over the two lines near the real axis (in opposite directions).

The two vertical segments have big real part — $z = R + it$ where t ranges between two bounded values — so we have e^R in the denominator, which means their integrals go to 0 as $R \to \infty$. And the integral over c_R converges to the original integral over c we saw earlier, while the integral over the small circle goes to 0 as $\varepsilon \to 0$. (We assumed that $\operatorname{Re}(s) > 1$, so that we have a removable singularity at $z = 0$.)

On the path to the right on the top, we get $\int_0^\infty \frac{t^{s-1}}{e^t-1}$ (we can imagine parametrizing it as $t \mapsto t + i\varepsilon$ as $\varepsilon \to 0$). On the bottom path, we could parametrize as $t \mapsto t - i\varepsilon$, but it's actually easier to use polar coordinates — we'll use t to denote the radius, and at $x - i\varepsilon$ the angle is $2\pi - \arctan(-\varepsilon/x)$, so we can write $z(t) = te^{2\pi i(1-\delta_t(\varepsilon))}$. But then $z'_\varepsilon(t) \to e^{2\pi i}$ as $\varepsilon \to 0$. So as we let $\varepsilon \to 0$, we get

$$\frac{(e^{2\pi i}t)^{s-1}}{e^{(e^{2\pi i}t)} - 1} \cdot e^{2\pi i}\, dt = e^{2\pi is} \int_0^\infty \frac{t^{s-1}}{e^t - 1}\, dt$$

(as $R \to \infty$ and $\varepsilon \to 0$). So then this means

$$\zeta(s) = \frac{1}{1 - e^{2\pi is}} \frac{1}{\Gamma(s)} \int_c \frac{z^{s-1}}{e^z - 1}\, dz.$$

(The point of the extra factor is whether angles are close to 0 or 2π — when we're approaching 0 from the top we get angles close to 0, and from the bottom we get angles close to 2π.)

So we have a new formula for $\zeta(s)$. And we also had that

$$\frac{1}{\Gamma(s)} = \frac{\sim \pi s}{\pi} \Gamma(1 - s).$$

And we have

$$1 - e^{2\pi is} = -e^{\pi is}(e^{\pi is} - e^{-\pi is}) = -2ie^{\pi is} \sin \pi s.$$

So this means

$$\zeta(s) = -e^{-\pi is} \cdot \frac{1}{2\pi i} \Gamma(1 - s) \int_C \frac{z^{s-1}}{e^z - 1}\, dz.$$

So we get another formula for the ζ function.

And now we can think about where this integral makes sense for all s. And it makes sense for all s.

Now let's try to compute $\zeta(0)$ using this formula. To do so, we deform our contour again. By the same logic as before, the integral around c can be deformed to the following integral (but for now we're not taking $\varepsilon \to 0$).

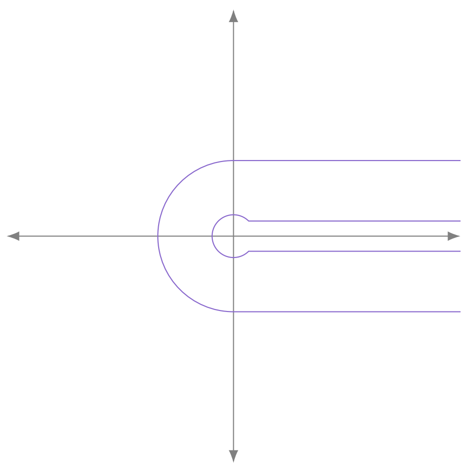

Then we get that the integral over c is the same as the integral over this keyhole of $\frac{z^{s-1}}{e^z-1}$ which splits into three pieces.

Note that the values this function takes on the top and the bottom become the same as we take width to 0, and so they cancel each other out. And so

$$\zeta(0) = -e^{\pi is} \frac{1}{2\pi i} \Gamma(1 - s) \oint_{|z|=\varepsilon} \frac{z^{s-1}}{e^z - 1},$$

where this loop goes clockwise, and we can swap it to go counterclockwise and remove the − sign. Then we get

$$\zeta(0) = \frac{1}{2\pi i} \cdot \oint \frac{z^{-1}}{e^z - 1} \, dz = \operatorname{Res}\left(\frac{1}{z(e^z-1)}, 0\right),$$

using the residue theorem (we have a pole at $z = 0$). We have $e^z - 1 = z + \frac{z^2}{2} + \cdots$, so

$$\frac{1}{z(e^z - 1)} = \frac{1}{z(z + z^2/2 + \cdots)},$$

and we get that the residue is $-1/2$. So this is equal to $-1/2$.

The other thing we promised was $s = -1$.

Claim 20.2 — We have $\zeta(-1) = -1/12$.

More generally, we can compute ζ at any negative integer in the following way. First,

$$\frac{z}{e^z - 1} + \frac{1}{2}z$$

is even. Second, we can define the *Bernoulli numbers* by

$$\frac{z}{e^z - 1} = 1 - \frac{1}{2}z + \sum_{k=1}^{\infty} (-1)^{k+1} B_k \frac{z^{2k}}{(2k)!}$$

(we take the series expansion of this function). And then you can compute

$$\zeta(-n) = \begin{cases} 0 & \text{if } n = 2\ell \\ (-1)^\ell B_\ell/(2\ell)! & \text{if } n = 2\ell - 1. \end{cases}$$

And so you can compute $B_1 = \frac{1}{6}$, so $\zeta(-1) = -\frac{1}{12}$.

Remark 20.3. The argumnt where the two long things become equal is true for all s being an integer, and not otherwise.

§20.2 Looking forwards

Analytic continuation leads to deeper waters. One application of the ζ function is the prime number theorem; there's a book on this we can read if we want.

Another direction you can go is to keep building on analytic continuation. We'll now be pretty loose, and Prof. Lawrie will pretend he is a topologist, which he is not.

Think about the function $\sqrt{\cdot}$ — think about the function ω such that $\omega(z)^2 = z$. So one way to think about this is that $\omega(z) = e^{\frac{1}{2}\log z}$.

You can first define this function that's holomorphic if you take away 0 and a line tending to ∞ — you have to define the slit plane (we'll take the positive real axis out). And now we can look at the values that $\sqrt{}$ takes as we go around.

Let's think about $\mathbb{C} \setminus \mathbb{R}_{\geq 0}$. We can think about $ze^{i\theta}$; then as we take the square root, we get

So if we choose a positive square root on top, then on the bottom of the cut we come to $-\sqrt{z}$. This is why we have to have the cut.

So where we put the cut is arbitrary, but there has to be a cut. This is the largest subset of $\mathbb{C}$ where we can define square root — we have to cut out something like this.

But now we can think about what if I can leave the complex plane, and find an even larger domain. Imagine taking two copies — the square root of z is $+$ on top and $-$ on the bottom of the cut in our first copy (Domain I). Domain II will also be $\mathbb{C} \setminus \mathbb{R}_{\geq 0}$, but here we'll have the opposite — we'll choose the other branch of the square root, so that we have $-\sqrt{z}$ on the top and $+\sqrt{z}$ on the bottom.

And now we can construct a new space by gluing these together — we know $\sqrt{}$ has two branches, and they agree if we flipflop them around. So we construct a space $\hat{X}$ by gluing I and II. We define ω to be $\sqrt{z}$ on I and $-\sqrt{z}$ on II.

What's the topological structure of X? We can think about stereogrpahic projection. If we have a slit plane and put it up onto the Riemann sphere, then the north pole gets excluded (corresponding to ∞), but now the south pole is alos excluded (since we cut out 0), and we've also cut out the positive real axis, which corresponds to a slit running from the north pole to the south pole.

Our second copy has the same structure.

In topology, if you want to understand what the topological structure of a space is, you can continuously deform it. So this sphere with this line cut out is like the same thing is if we take this slit and open it up, to make it a half-sphere. And the other guy — the plane II that we cut out — we can do the same thing there, where we have a slit from the north pole to south pole but with opposite signs. So that also gives a half-sphere like this after we open it up.

And our gluing procedure was to identify the places where the signs agree; and gluing this together gives us a sphere without the north pole and the south pole.

So then X is the sphere minus the north pole and the south pole. sThis is now the domain for the square root function — and this is bigger than the slit complex plane.

And then what you want to do next is, this is the topological structure of the space (a sphere minus two points). And then we want to understand what the function looks like on this space. Now we have to make sure that everything is okay — we had a function that was $\sqrt{z}$ on I and $-\sqrt{z}$ on II, but now we want to understand it as a function on this sphere minus two points, nad we hve to make sure it's well-defined everywhere — because when we put these two halves together, we get a sphere minus a north pole and south pole. And now we want to think about the square root function on this sphere minus two points.

If we pick a point z_0 on the sphere, we want to understand what the function looks like. We do this via a coordinate chart. So at each z_0 of $\widehat{X}$, there exists a neighborhood V of z_0 and a bijection $\Phi \colon V_{z_0} \to D$ where D is an open disk in $\mathbb{C}$, such that $\omega(\Phi_{z_0}^{-1}(z))$ is a holomorphic function on D; and if we have overlapping disks then everything is fine, meaning that $\Phi_{z_1} \circ \Phi_{z_0}^{-1}$ is holomorphic on the overlap $\Phi_{z_0}(V_{z_0} \cap V_{z_1})$ (and the derivative of this thing is nonzero).

These bijections are called coordinate charts. And if we are able to do this, then this gives $\hat{X}$ a structure of a *complex manifold*.

So you can take classes like 950 to learn about manifolds, and also 101. And this is an example of a complex manifold. The distinction is that these coordinate charts need to be holomorphic.

You have to check this carefully for our $\sqrt{}$ function; because you might have taken a neighborhood right where the gluing happened, and now you have half of a chart in I and half a chart in II. And you have to check that the function makes sense. But we did this by construction, since we identified the correct branches of the squrae root.

And you can even extend the domain to include the two points at ∞ — you get a double pole at the north pole and a double zero at teh south pole (which makes sense as $\sqrt{0} = 0$).

Remark 20.4. Why are we doing this? One thing it's relatd to is — we defined square roots (or nth roots, where you have n different choices), and this is a way of reconciling those choices. So it relates to that. We also were talking about anlayitic continuation, where we were trying to analytic continue to a bigger subset of $\mathbb{C}$. Here we're not limiting ourselves to $\mathbb{C}$; we're allowing ourselves to take a bunch of idfferent copies. (For example for $\sqrt{}$ there's two different $\sqrt{}$'s; the question is can we combine them into one domain). SO this goes past what we started talking about with analytic continuation, where we don't just restrict ourselves to one $\mathbb{C}$.

You can also reconcile the different branches of log this way. What we just saw is called the Riemann surface of $\sqrt{}$. The log function has a beautiful-looking Riemann surface, which is a spirally-type thing.

Prof. Lawrie wants to mention one other function we considered in class, elliptic functions, or rather, their inverses. The function $w^2 = z(z-1)(z-\lambda)$ where $\lambda \notin \{0,1\}$ is a polynomial of degree 3 on the right-hand side, and then we can take the square root of this. What do we need to do? Imagine 0, 1, λ in that order on the real axis (it doesn't really matter), and imagine we look at what happens to the values this functiont akes as we go around.

A priori, we have to remove the whole slit $\mathbb{R}_{\geq 0}$. On this line segment $[0,1]$, as we go around, $\sqrt{z} \mapsto -\sqrt{z}$ so the values don't match up. But then $\sqrt{z-1}$ is fine there — it goes to itself — and same with λ. So we get one $-$ and two $+$'s, which means the whole function's sign switches.

How about on $[1,\lambda]$? There $\sqrt{z} \mapsto -\sqrt{z}$, $\sqrt{z-1} \mapsto -\sqrt{z-1}$ also changes sign, but $\sqrt{z-\lambda} \mapsto \sqrt{z-\lambda}$ doesn't change sign. So here the whole function doesn't change sign, which means we actually don't have to remove this slit.

And then all three of them change sign past λ, since all three factors do, so we do have to remove that.

But the point is we don't have to remove the middle slit.

We define plane II where we do the same procedure, removing these two slits. And now we try to do the same procedure and glue them together.

When we look at the Riemann sphere, I looks like the north pole with a slit out of it, and the south pole with a separate slit out of it. And we can open up these two holes, so we get half of a donut. And we can do the same thing with the other II, which gives the other half of a donut. And then you can glue these together, and you get the donut (torus). So the Riemann sphere for this function is the torus.

You can ask which torus; there's another way of thinking about toruses, as you take one of these parallelograms and then wrap it up and identify the sides. The important thing is that you have generators related to λ, so you can actually figure out which torus you get.

The proper setting to develop $\wp$ is in this context; they're supposed to be inverses of each other. You can ask which complex functions live on Riemann surfaces; this turns out to be a deep question, and there's lots of different areas of math where this is relevant.

Remark 20.5. If two functions have homeomorphic Riemann surface, can you say stuff about them?

Maybe homeomorphism isn't the right thing to distinguish, because lots of very different functions have homeomorphic Riemann surfaces.

Prof. Lawrie is more familiar with questions in the other direction — you want to answer a geometric question by loooking at which holomorphic functions live on some Riemann surface. You can generalize to differential forms and stuff.