1 Some very useful theorems

The ultimate goal of this class is to understand the geometric properties of analytic functions, but before we can do that, there are quite a few theorems that we need to prove first. All of these theorems are suprising, beautiful, and insanely useful.

1.1 Isolated roots and unique continuations

Definition. A function $f : \mathbb{C} \to \mathbb{C}$ is said to have *isolated roots* if its roots (the places where it evaluates to 0) are "spaced out". Formally, this means that if $z_0 \in \mathbb{C}$ is a root, namely $f(z_0) = 0$, then there is some $\varepsilon > 0$ such that if $0 < |z - z_0| < \varepsilon$, then $f(z) \neq 0$. In other words, the roots of f are not arbitrarily close to one another.

Having isolated roots is a very special property. For instance, the function $g : \mathbb{C} \to \mathbb{C}$ defined by $g(x + iy) = x - y$ is an extremely nice function, but it does *not* have isolated roots: the entire line $y = x$ is composed of roots of g, so every root has other roots arbitrarily close to it. Nevertheless, we have the following result:

Theorem. *If $f : \mathbb{C} \to \mathbb{C}$ is analytic and not the zero function, then f has isolated roots.*

Proof. Suppose not. Then we have a root z_0 that has other roots arbitrarily close to it. This means that we can find a sequence of points $z_1, z_2, \ldots \in \mathbb{C}$ with $f(z_k) = 0$ and $\lim_{k\to\infty} z_k = z_0$. We expand f as a Taylor series centered at z_0, by writing

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

for some collection of complex numbers a_n. Since f is not the zero function, one of the a_n must be nonzero; let m be the smallest integer such that $a_m \neq 0$. Then the above Taylor series expansion becomes

$$f(z) = \sum_{n=0}^{\infty} a_n (z - z_0)^n = \sum_{n=m}^{\infty} a_n (z - z_0)^n = a_m (z - z_0)^m \left(1 + \sum_{n=1}^{\infty} \frac{a_{n+m}}{a_m} (z - z_0)^n\right)$$

So if we define

$$g(z) = \sum_{n=1}^{\infty} \frac{a_{n+m}}{a_m} z^n$$

then we have written

$$f(z) = a_m (z - z_0)^m (1 + g(z - z_0))$$

Moreover, from the definition of g, we see that

$$\lim_{z \to z_0} g(z - z_0) = 0$$

This, along with the fact that $z_k \to z_0$, implies that for some k large enough, we have that $g(z_k - z_0) \neq -1$. Therefore, plugging in such a k to our formula for f, we get that

$$0 = f(z_k) = a_m (z_k - z_0)^m (1 + g(z_k - z_0)) \neq 0$$

since every term on the right-hand side is nonzero. This is a contradiction, as desired. □

Corollary. *Suppose $f_1, f_2 : \mathbb{C} \to \mathbb{C}$ are analytic functions, and suppose they are equal on some non-discrete set (e.g. a line segment, or a curve, or a disk). Then they are equal everywhere.*

Proof. Consider the function $f = f_1 - f_2$. It is analytic, and we wish to prove that it's the zero function (since that implies that $f_1 = f_2$). For this, observe that $f(z) = 0$ whenever $f_1(z) = f_2(z)$, by definition. Additionally, since we assumed that f_1, f_2 agreed on some non-discrete set, this implies that the roots of f cannot be isolated. So by the previous theorem, we must have that $f = 0$, as desired. □

This result is quite powerful, and in my opinion, very surprising. It tells us that the behavior of an analytic function on some tiny segment actually determines its behavior on the whole complex plane. It also makes the Cauchy Integral Formula seem less impressive—that theorem allows us to determine what an analytic function does inside a region based on what it does on the boundary, but this result is *much* more general than that.

One corollary of this result is a new proof of a theorem that we already know, namely that the conjugation function $f(z) = \overline{z}$ is not analytic. Indeed, this function agrees with the identity function $g(z) = z$ on the non-discrete set $\mathbb{R} \subset \mathbb{C}$. Since g is analytic, if f were analytic, then they'd have to agree everywhere, and they don't; thus, f is not analytic.

1.2 The Argument Principle

Our next result is a very useful consequence of the Residue Theorem, which tells us how to count roots and poles of meromorphic functions.

Theorem (The Argument Principle)**.** *Let $f : \mathbb{C} \to \mathbb{C}$ be meromorphic, and let C be a simple closed curve. Suppose that f is never zero and has no poles on C. Then*

$$\frac{1}{2\pi i} \int_C \frac{f'(z)}{f(z)} \, \mathrm{d}z = R - P$$

where R is the number of roots of f inside C, and P is the number of poles of f inside C, and both are counted with multiplicity.

Proof. First, we observe a very useful property of the quantity f'/f, which is called the *logarithmic derivative* of f (for the reason why, see the homework!). Namely, if f_1, f_2 are two functions, then by the product rule,

$$\frac{(f_1 f_2)'}{f_1 f_2} = \frac{f_1' f_2 + f_1 f_2'}{f_1 f_2} = \frac{f_1'}{f_1} + \frac{f_2'}{f_2}$$

In other words, the logarithmic derivative of a product of functions is the sum of the logarithmic derivatives of the individual functions.

Now, recall the Residue Theorem. It tells us that

$$\frac{1}{2\pi i} \int_C \frac{f'(z)}{f(z)} \, \mathrm{d}z = \sum_{\text{poles } z \text{ of } f'/f \text{ in } C} \operatorname{res}_z \left(\frac{f'}{f} \right)$$

So where does f'/f have poles? Whenever f is zero, we're dividing by zero, so we certainly expect a pole there. Additionally, wherever f has a pole, then f' will also have a pole there, so we expect a pole of f'/f. Finally, no other point will be a pole: at any other point, we will be dividing an analytic function (namely f') by a nonzero analytic function (namely f), so we will get something analytic as the quotient.

So now we just need to understand the residues at these two types of points, namely roots of f and poles of f. Let's start with the roots: suppose z_0 is a root of f with multiplicity m. Recall that this means that we can write

$$f(z) = (z - z_0)^m g(z)$$

for some analytic function g with $g(z_0) \neq 0$. Then using our additivity formula, we see that

$$\frac{f'(z)}{f(z)} = \frac{[(z - z_0)^m]'}{(z - z_0)^m} + \frac{g'(z)}{g(z)} = \frac{m}{z - z_0} + \frac{g'(z)}{g(z)}$$

Since g is analytic and nonzero near z_0, we see that g'/g is also analytic there. So that term contributes nothing to the residue. So the only term that matters is $m/(z - z_0)$, whose residue at z_0 is precisely m. This means that at a root of multiplicity m, f'/f has a simple pole with residue precisely m.

On the other hand, if z_1 is a pole of f of order n, then we can write

$$f(z) = (z - z_1)^{-n} h(z)$$

for some analytic function h with $h(z_1) \neq 0$. Then we get that

$$\frac{f'(z)}{f(z)} = \frac{[(z-z_1)^{-n}]'}{(z-z_1)^{-n}} + \frac{h'(z)}{h(z)} = \frac{-n}{z-z_1} + \frac{h'(z)}{h(z)}$$

which again implies that f'/f has a simple pole of residue $-n$ at z_1.

So when we add up all residue of f'/f inside C, we find that this sum is precisely $R - P$, as desired. □

This theorem is extremely useful in lots of contexts. For one cool application, recall that the Riemann Hypothesis says that a certain function (the Riemann ζ function) has its ("non-trivial") roots on the line $\{\frac{1}{2} + it \mid t \in \mathbb{R}\}$. This has been checked for the 10^{13} smallest roots, and they indeed all lie on this line. But how do people actually check this? It turns out that the way they do so is by numerically integrating ζ'/ζ over a big rectangle near this critical line, which tells them how many roots are in this rectangle by the Argument Principle. Using the intermediate value theorem on a related function defined on the critical line, they can also count how many roots are actually on the line, and then by comparing this number, they can confirm that all the roots are on the line.

1.3 Rouché's Theorem

Rouché's Theorem is another very useful tool, which tells us that when we perturb a function a little bit, then the number of roots in a region is unchanged.

Theorem (Rouché)**.** *Suppose that f, g are analytic functions, C is a simple closed curve, and for every point z on the curve C, we have that*

$$|f(z)| > |g(z)|$$

Then f and $f + g$ have the same number of roots inside C.

Proof. We're going to introduce a "time" parameter $t \in [0, 1]$, and define

$$f_t(z) = f(z) + tg(z)$$

In other words, at time 0, we have that f_0 is just f, whereas at time 1, $f_1 = f + g$. So we can imagine f_t as being a sort of evolution over time of our function as it goes from f to $f + g$.

Since $|f| > |g|$ on C, we necessarily have that f_t has no roots on the curve C. Indeed, if $f_t(z) = 0$ for some $z \in C$, then

$$0 = f(z) + tg(z) \Longrightarrow |f(z)| = t|g(z)| < t|f(z)| \leq |f(z)|$$

which is a contradiction. Therefore, if we let n_t be the number of roots of f_t inside C, then the Argument Principle tells us that

$$n_t = \frac{1}{2\pi i} \int_C \frac{f_t'(z)}{f_t(z)} \, \mathrm{d}z$$

On the other hand, since f_t varies continuously in t and since it's never zero on C, we get that f_t'/f_t also varies continuously in t. Thus, we can conclude that n_t is a continuous function of t. However, n_t is definitely an integer. By the intermediate value theorem, this implies that n_t must be a constant. In particular, $n_0 = n_1$, which is what we wanted to prove. □

1.4 The Open Mapping Theorem

First, we need to recall the definition of an open set:

Definition. A set $U \subseteq \mathbb{C}$ is called *open* if for every $z \in U$, there is some $\varepsilon > 0$ such that if $|w - z| < \varepsilon$, then $w \in U$. In other words, U is open if no point in U is "arbitrarily close" to a point outside U.

Definition. A function $f : \mathbb{C} \to \mathbb{C}$ is called an *open mapping* if for every open set U, $f(U)$ is also open, where

$$f(U) = \{f(z) : z \in U\}$$

Being an open mapping is another extremely special property. For instance, the function $g : \mathbb{C} \to \mathbb{C}$ defined by $g(x + iy) = x - y$ is not an open mapping, since its range is the real line $\mathbb{R} \subset \mathbb{C}$, and no subset of that is open. Nevertheless, we have the following theorem:

Theorem (Open Mapping Theorem). *If $f : \mathbb{C} \to \mathbb{C}$ is analytic and non-constant, then f is an open mapping.*

Proof. Fix some open set U, and let $w_0 \in f(U)$; we wish to show that points sufficiently close to w_0 are also in $f(U)$. Since $w_0 \in f(U)$, we can find some $z_0 \in U$ such that $f(z_0) = w_0$. Since U is open and $z_0 \in U$, there is some $\delta > 0$ such that the disk $\{z : |z - z_0| \leq \delta\}$ is contained in U. Moreover, since the analytic function $f(z) - w_0$ has isolated roots (this is crucially where we use that f is non-constant), we can pick this δ so that $f(z) \neq w_0$ on the circle $C = \{z : |z - z_0| = \delta\}$.

Now, the function $|f(z) - w_0|$ is a continuous real-valued function, and it's never zero on the circle C, so we can find some $\varepsilon > 0$ so that $|f(z) - w_0| \geq \varepsilon$ for all z with $|z - z_0| = \delta$.

Now, fix w so that $|w - w_0| < \varepsilon$; we wish to prove that $w \in f(U)$, which will imply that U is open. For this, define

$$g(z) = f(z) - w \qquad F(z) = f(z) - w_0 \qquad G(z) = w_0 - w$$

Then we have that $g = F + G$. Moreover, by construction, we have that on the circle C,

$$|G(z)| = |w_0 - w| < \varepsilon \qquad |F(z)| = |f(z) - w_0| \geq \varepsilon$$

and thus $|G| < |F|$ on C. So we can apply Rouché's Theorem to conclude that F and $F + G$ have the same number of roots inside C. Since z_0 is inside C and $f(z_0) = w_0$, we see that F has at least one root inside C, so $g = F + G$ does as well. Thus, there is some z inside C with $f(z) = w$. Since all of the interior of C is contained in U, this implies that $w \in f(U)$, as desired. ☐

Note that this open mapping property is very much unique to the world of complex analysis. For instance, the function $h : \mathbb{R} \to \mathbb{R}$ given by $h(x) = x^2$ is basically as nice as we can imagine (e.g. infinitely differentiable), but is *not* an open mapping. For instance, the interval $(-1, 1)$ is open, but $h((-1, 1)) = [0, 1)$, which is not an open set in $\mathbb{R}$.

1.5 The Maximum Modulus Principle

Theorem (Maximum modulus principle). *If f is a non-constant analytic function defined on some open set U, then $|f|$ cannot attain a maximum anywhere in U.*

Again, this is a property very unique to the complex world: the real-valued function $1 - x^2$ achieves its maximum at 0, which is inside many open sets (e.g. the open interval $(-1, 1)$).

Proof. Suppose for contradiction that $z_0 \in U$ were a maximum for $|f|$. Since f is analytic and non-constant, by the Open Mapping Theorem, we know that there is some small disk D centered $f(z_0)$ such that $D \subseteq f(U)$. But in particular, some of the points in D have a larger absolute value than the center of the disk $f(z_0)$, which means that we've found some z_1 such that $|f(z_1)| > |f(z_0)|$. This contradicts the assumed maximality, as desired. ☐

For one application of this, suppose that U is some bounded open set. Then since $|f|$ is continuous on the closure $\overline{U}$, we know that $|f|$ must attain a maximum somewhere in $\overline{U}$, and it can't be in U. Thus, $|f|$ must be maximized on the boundary ∂U. So when trying to find a maximum of an analytic function on a bounded set, it suffices to only consider the boundary.

2 Conformal Geometry

Let's look at some pictures! You can find the pictures I used at `https://goo.gl/52da7S`

The first set of pictures shows us what happens to a grid of lines as they're mapped forward by various analytic functions. As we can see from these pictures, shapes can get extremely distorted, but one very important property is preserved: every blue line meets every red line at right angles.

Of course, this is a lie, as some of the pictures show us: the function $f(z) = z^2$ is analytic, but can't possibly preserve angles at the origin. Indeed, a horizontal line (the real numbers) gets mapped to the non-negative real axis, whereas a vertical line (the imaginary numbers) gets mapped to the non-positive real axis, and these two meet at an angle of $180°$. Indeed, as the picture on Slide 8 shows, things go kind of crazy near the origin. As we will soon see, the big problem that's happening here is that the function $f(z) = z^2$ has its derivative equal to 0 at the origin.

In the next set of pictures, we look at what happens to concentric circles under analytic maps. Again, they can get weirdly deformed, but we see that as they get smaller, they look more and more like actual circles. This property even appears to be true for our weird function $f(z) = z^2$, though this is an illusion: the last picture shows us a better picture of what's actually happening, by plotting the small perturbation $z^2 + 0.1z$. In that picture, we can see that the nice circles we saw for z^2 are actually "wrapping around twice", and are thus not really circles at all, no matter how small they are. Again, this is because all of these circles, no matter how small, contain a point where the derivative is zero.

We can summarize these results as follows:

Analytic functions preserve all "local" geometry up to scaling: angles are preserved, and lengths are preserved up to a single scaling of all of them.

This is called the *conformal* property of analytic maps.

3 Conformal Equivalence

Recall that a key property of analytic functions is that they're *conformal* wherever their derivative is nonzero. This means that they preserve all local geometry: angles are unchanged, and all small lengths are scaled by (almost) the same amount.

Why do we expect this to be true? Recall that over the reals, a differentiable function is one that can be approximated by a linear function at every point: if we zoom in enough on the graph, then it will just look like a line. Similarly, an analytic function is one that can be approximated by a *complex line*, namely something of the form $az + b$, where $a, b \in \mathbb{C}$. The addition of b is just a translation, which shouldn't affect the geometry at all. Moreover, if we write the complex number a as $re^{i\theta}$, then multiplication by a is the same as scaling by r and rotating by θ. Since this approximation of any analytic function by a linear function gets more and more accurate as zoom in more, this precisely gives us the desired property.

Moreover, this explains why everything goes wrong when the derivative vanishes. In that case, $a = 0$, so our "scaling" actually contracts the entire plane into a point. In that case, what actually ends up mattering isn't the first-order linear approximation, but the higher-order approximations, and these nice geometric properties just don't appear.

We won't actually state or prove any of this formally—if you want to see an example of a formal statement and proof, check out the last problem of yesterday's homework—because it's kind of annoying to do 100% formally. Instead, we will move on to an important application of this idea, which is conformal equivalence.

Definition. Let U and V be two open subsets of $\mathbb{C}$. A *conformal equivalence* (or *biholomorphism*) between U and V is a pair of functions $f : U \to V, g : V \to U$ such that both f and g are analytic on their respective domains, and they are inverses, namely

$$f \circ g = \mathrm{id}_V \qquad\qquad g \circ f = \mathrm{id}_U$$

If such a conformal equivalence exists, then we say that U and V are *conformally equivalent*.

We will prove in a moment that if $f : U \to V, g : V \to U$ are a conformal equivalence, then f' and g' never vanish. This implies that a conformal equivalence between U and V gives that U and V have the same local geometry. In other words, conformal equivalence is a very strong notion: it tells us that U and V have the "same shape", though this needs to be interpreted locally.

As promised, we now prove that biholomorphisms have nowhere-vanishing derivatives.

Proposition. *If $f : U \to V, g : V \to U$ are a conformal equivalence, then $f'(z) \neq 0, g'(w) \neq 0$ for all $z \in U, w \in V$.*

Proof. Since $g \circ f = \mathrm{id}_U$, we know that for any $z \in U$,

$$g(f(z)) = z$$

Differentiating this equation gives, by the chain rule,

$$g'(f(z))f'(z) = 1$$

In particular, if $f'(z_0) = 0$ for some z_0, then we must have that $g'(f(z_0))$ is not well-defined, meaning that g must have a pole at $f(z_0)$. However, $f(z_0)$ is some point of V, and we assumed that g is analytic on all of V, and this is a contradiction.

Reversing the roles of f and g similarly proves that $g'(w) \neq 0$ for all $w \in V$. $\square$

So, as above, we get that conformal equivalence is a very restrictive condition, which tells us that two regions have the same local geometry. We will soon quantify just how restrictive this is by classifying all biholomorphisms from the disk to itself. However, it's not quite as restrictive as you might expect.

Example. From now on, let $\mathbb{D}$ denote the open unit disk

$$\mathbb{D} = \{z : |z| < 1\}$$

and let $\mathbb{H}$ denote the open upper half-plane

$$\mathbb{H} = \{z : \mathrm{Im}(z) > 0\}$$

Then despite the fact that $\mathbb{D}$ and $\mathbb{H}$ look very different (e.g. $\mathbb{D}$ is bounded and $\mathbb{H}$ is not), they turn out to be conformally equivalent! We can write down an explicit biholomorphism between them:

$$f(z) = \frac{i-z}{i+z} \qquad\qquad g(w) = i\frac{1-w}{1+w}$$

Then we claim that f maps $\mathbb{H} \to \mathbb{D}$ and g maps $\mathbb{D} \to \mathbb{H}$, and they are inverses.

It's actually easiest to check that they are inverses; indeed,

$$\begin{aligned}
g(f(z)) &= i\frac{1-\frac{i-z}{i+z}}{1+\frac{i-z}{i+z}} = i\frac{(i+z)-(i-z)}{(i+z)+(i-z)} = i\frac{2z}{2i} = z \\
f(g(w)) &= \frac{i-i\frac{1-w}{1+w}}{i+i\frac{1-w}{1+w}} = \frac{(1+w)-(1-w)}{(1+w)+(1-w)} = \frac{2w}{2} = w
\end{aligned}$$

Next, we check that f actually maps $\mathbb{H} \to \mathbb{D}$. Indeed, if $z \in \mathbb{H}$, then it is necessarily closer to i than it is to $-i$. So $|i-z| < |i+z|$, and thus $|f(z)| < 1$, so $f(z) \in \mathbb{D}$. On the other hand, to check that g maps $\mathbb{D} \to \mathbb{H}$, we need to check that $\mathrm{Im}(g(w)) > 0$ for any $w \in \mathbb{D}$. Indeed, if we write $w = a + bi$, then

$$\begin{aligned}
\mathrm{Im}(g(w)) &= \mathrm{Re}\left(\frac{1-w}{1-w}\right) \\
&= \mathrm{Re}\left(\frac{1-a-bi}{1+a+bi}\right) \\
&= \mathrm{Re}\left(\frac{(1-a-bi)(1+a-bi)}{(1+a)^2+b^2}\right) \\
&= \frac{1-a^2-b^2}{(1+a)^2+b^2}
\end{aligned}$$

Then the denominator is always non-negative, and is nonzero when $w \neq -1$, and the numerator is always positive if $|w| = a^2 + b^2 < 1$. So for $w \in \mathbb{D}$, we indeed get that $\mathrm{Im}(g(w)) > 0$. Thus, we indeed have that $\mathbb{H}$ and $\mathbb{D}$ are conformally equivalent. Here is a picture of what f does to $\mathbb{H}$:

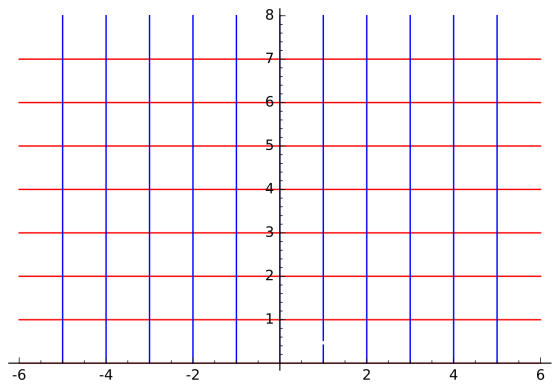

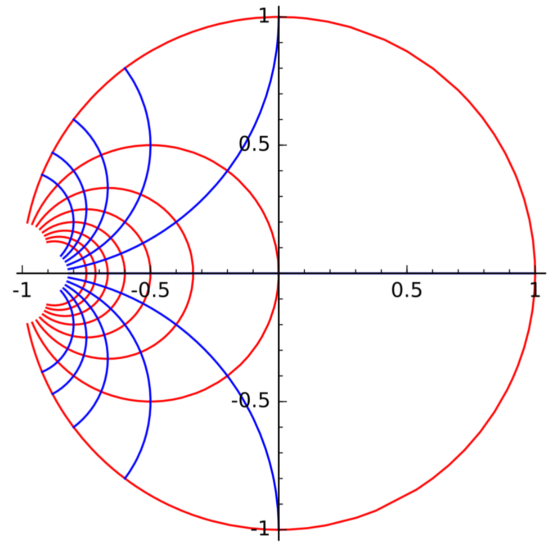

4 Classifying Conformal Equivalences

4.1 The Schwarz Lemma

The conformal equivalence between $\mathbb{H}$ and $\mathbb{D}$ might be surprising, but the intuition that conformal equivalences are rare and special is still correct. To demonstrate this, we will classify *all* biholomorphisms $\mathbb{D} \to \mathbb{D}$. We will be able to just write down a list of them, which agrees with our intuition that they should be rare, even in a nice case like $\mathbb{D}$.

Our first important tool is the Schwarz Lemma:

Lemma (Schwarz)**.** *Let $f : \mathbb{D} \to \mathbb{D}$ be analytic with $f(0) = 0$. Then the following hold:*

i. For any $z \in \mathbb{D}$, $|f(z)| \leq |z|$.

ii. If for some $0 \neq z_0 \in \mathbb{D}$ we have $|f(z_0)| = |z_0|$, then f is a rotation, namely

$$f(z) = e^{i\theta} z$$

for some $\theta \in \mathbb{R}$.

iii. $|f'(0)| \leq 1$, and if $|f'(0)| = 1$, then f is again a rotation.

Proof. First, we claim that $f(z)/z$ has a removable singularity at 0. Indeed, if we write f as a Taylor series, then the constant term will be zero since $f(0) = 0$. So we can write

$$f(z) = a_1 z + a_2 z^2 + a_3 z^3 + \cdots \implies \frac{f(z)}{z} = a_1 + a_2 z + a_3 z^2 + \cdots$$

which shows that we can indeed extend $f(z)/z$ to be analytic at 0. Now, since f is a function $\mathbb{D} \to \mathbb{D}$, we have that $|f(z)| \leq 1$ for all $z \in \mathbb{D}$. Fix some $z \in \mathbb{D}$, and let $r = |z| < 1$. Then we have that

$$\left|\frac{f(z)}{z}\right| = \frac{|f(z)|}{|z|} \leq \frac{1}{r}$$

Now, we apply the Maximum Modulus Principle to the open disk of radius r. Since the value of $|f(z)/z|$ is upper-bounded by $1/r$ on the boundary of that disk, we conclude that

$$\left|\frac{f(z)}{z}\right| \leq \frac{1}{r}$$

for all z with $|z| \leq r$. Now, let $r \to 1$, and we conclude that for all $z \in \mathbb{D}$, we have that

$$\left| \frac{f(z)}{z} \right| \leq 1$$

which gives $|f(z)| \leq |z|$, proving (i).

For (ii), suppose that we have some $z_0 \neq 0$ with $|f(z_0)| = |z_0|$. Then $|f(z)/z|$ must be maximized at z_0, meaning that $f(z)/z$ attains a maximum inside an open set. This contradicts the Maximum Modulus Principle unless $f(z)/z$ is a constant, say $f(z)/z = c$ for some $c \in \mathbb{C}$. But then we have that

$$|z_0| = |f(z_0)| = |cz_0| = |c||z_0|$$

Since $z_0 \neq 0$, then $|z_0| \neq 0$ as well, so we can divide out and conclude that $|c| = 1$. But that means that we can write $c = e^{i\theta}$ for some $\theta \in \mathbb{R}$, and we get that indeed $f(z) = e^{i\theta}z$, so that f is a rotation.

Finally, for (iii), set $g(z) = f(z)/z$. Then by (i), we know that $|g(z)| \leq 1$ for all $z \in \mathbb{D}$. Additionally,

$$f'(0) = \lim_{z \to 0} \frac{f(z) - f(0)}{z - 0} = \lim_{z \to 0} \frac{f(z)}{z} = \lim_{z \to 0} g(z) = g(0)$$

So since $|g(0)| \leq 1$, we get that $|f'(0)| \leq 1$. Moreover, if $|f'(0)| = 1$, then g must attain a maximum at 0, and must thus be a constant by the Maximum Modulus Principle. So if we set $c = g(z)$, then we get that

$$1 = |f'(0)| = |g(0)| = |c|$$

and we can again conclude that $c = e^{i\theta}$, and thus that f is a rotation. $\square$

4.2 The Classification of Conformal Equivalences $\mathbb{D} \to \mathbb{D}$

On the homework, you were introduced to the functions ψ_α (sometimes called Blaschke factors), one for each $\alpha \in \mathbb{D}$, defined by

$$\psi_\alpha(z) = \frac{\alpha - z}{1 - \overline{\alpha} z}$$

You proved that for every $\alpha \in \mathbb{D}$, ψ_α is a conformal equivalence from $\mathbb{D}$ to itself, whose inverse is also given by ψ_α. You also proved that $\psi_\alpha(0) = \alpha, \psi_\alpha(\alpha) = 0$. These functions turn out to be crucial in the classification of all conformal equivalences $\mathbb{D} \to \mathbb{D}$.

Theorem (Classification Theorem)**.** *Suppose that $f : \mathbb{D} \to \mathbb{D}$ is a biholomorphism, namely it's an analytic bijection with an analytic inverse. Then there exist $\theta \in \mathbb{R}, \alpha \in \mathbb{D}$ such that*

$$f(z) = e^{i\theta}\psi_\alpha(z)$$

Thus, the only conformal equivalences $\mathbb{D} \to \mathbb{D}$ are compositions of a rotation and a Blaschke factor.

Proof. Since f is a bijection $\mathbb{D} \to \mathbb{D}$, there is a unique $\alpha \in \mathbb{D}$ such that $f(\alpha) = 0$. Define $g = f \circ \psi_\alpha$, so that

$$g(0) = f(\psi_\alpha(0)) = f(\alpha) = 0$$

So we may apply the Schwarz Lemma to g to conclude that for all $z \in \mathbb{D}$,

$$|g(z)| \leq |z|$$

On the other hand, g is also a biholomorphism, since it's the composition of two biholomorphisms, and its inverse is given by $g^{-1} = \psi_\alpha \circ f^{-1}$. We also have that $g^{-1}(0) = 0$, so we may apply the Schwarz Lemma to g^{-1} to conclude that

$$|g^{-1}(w)| \leq |w|$$

for all $w \in \mathbb{D}$. If we plug in $w = g(z)$ into this inequality, we get that

$$|z| = |g^{-1}(g(z))| \leq |g(z)|$$

Combining this with the above inequality, we get that $|g(z)| = |z|$ for all $z \in \mathbb{D}$. So applying the second part of the Schwarz Lemma, we get that g must be a rotation, namely $g(z) = e^{i\theta}z$ for some $\theta \in \mathbb{R}$. Therefore, using the fact that $\psi_\alpha \circ \psi_\alpha = \mathrm{id}_{\mathbb{D}}$, we get that

$$f = f \circ \psi_\alpha \circ \psi_\alpha = g \circ \psi_\alpha = e^{i\theta}\psi_\alpha$$

as desired. □

One nice further corollary of this classification is the following, which comes from considering the case where $\alpha = 0$ in the above proof:

Corollary. *If $f : \mathbb{D} \to \mathbb{D}$ is a biholomorphism that fixes the origin (namely $f(0) = 0$), then f is a rotation.*

5 Conformal Equivalences between other Regions

Again, we expect conformal equivalences to be very rare: for random open subsets $U, V \subseteq \mathbb{C}$, there is no reason to expect U and V to be conformally equivalent, since that would mean that they must have the same local geometry. Moreover, we already know that they must be rare: we've managed to write down a full list of all conformal equivalences $\mathbb{D} \to \mathbb{D}$. So let's try to quantify this intuition, by coming up with conditions on U and V that guarantee that they're *not* conformally equivalent.

First, do we expect $\mathbb{C}$ and $\mathbb{D}$ to be conformally equivalent? If they were, then we'd have an analytic map $f : \mathbb{C} \to \mathbb{D}$. However, since $\mathbb{D}$ is a bounded set, that would imply that f is a bounded map that is analytic on the entire complex plane. By Liouville's Theorem, such a function must be constant, and in particular can't be a bijection. So we certainly can't get a conformal equivalence between $\mathbb{C}$ and $\mathbb{D}$, or in fact between $\mathbb{C}$ and any bounded set. Moreover, this implies that $\mathbb{C}$ and $\mathbb{H}$ aren't conformally equivalent—if they were, then we could compose with our conformal equivalence between $\mathbb{H}$ and $\mathbb{D}$ and get a conformal equivalence $\mathbb{C} \to \mathbb{D}$, which we just proved is impossible.

Another obstruction to conformal equivalence is topology; for instance, do we expect $\mathbb{D}$ to be conformally equivalent to the annulus $A = \{z : 1 < |z| < 2\}$? Intuitively, the answer must be no, since A "has a hole in it". Indeed, there can't even be a homeomorphism (a much weaker notion than conformal equivalence) between $\mathbb{D}$ and A, since $\mathbb{D}$ and A have different topological structures. This difference is measured by the *fundamental group*, which we won't be discussing in any detail. However, for our purposes, a subset U of $\mathbb{C}$ is called *simply connected* if it "has no holes"; more formally, U is simply connected if every loop in U can be deformed to a constant loop while staying inside U.

The truly astonishing result is that these are the *only* obstructions to being conformally equivalent; any two subsets that don't run into one of these problems will be conformally equivalent. This goes completely against my intuition, which expects actual strong geometric conditions on U and V; as it turns out, no such conditions exist. Here is the formal statement:

Theorem (Riemann Mapping Theorem)**.** *Let U and V be two non-empty open subsets of $\mathbb{C}$. Suppose that neither U nor V is the whole plane $\mathbb{C}$, and suppose that they're both simply connected. Then they are conformally equivalent.*

Example. On the homework, I asked you whether you thought there was a conformal equivalence between a square and a rectangle that is not a square. You probably expected there to not be such an equivalence (that's my intuition), but the Riemann Mapping Theorem actually guarantees that there is one. Nevertheless, your intuition can be salvaged somewhat: one can prove that there is no conformal equivalence between a square and a rectangle that maps vertices to vertices, meaning that a conformal equivalence between these two shapes must stretch the boundary in really weird ways.

A stranger example is that the interior of a Koch snowflake, which is a very weird fractal shape, must be conformally equivalent to the disk. I find this extremely surprising, because the geometry of the Koch snowflake is very complicated near its boundary, so it is really surprising that one can find a way to identify this geometry (locally) with that of the disk.

Despite the Riemann Mapping Theorem, our intuition that conformal equivalences have got to be extremely rare is still more or less correct. For instance, one can prove that two annuli

$$A_1 = \{z : a < |z| < b\} \qquad \text{and} \qquad A_2 = \{z : c < |z| < d\}$$

are not conformally equivalent unless $a/c = b/d$ (in which case we can just rescale one to get the other) or unless $ad = bc$ (in which case we can first invert the annulus by sending $z \mapsto 1/z$ and then rescale).

A slightly more convenient way of phrasing this theorem is as follows:

Theorem. *Let U be a non-empty open subset of $\mathbb{C}$. Suppose U is simply connected and not all of $\mathbb{C}$. Then there is a conformal equivalence between U and $\mathbb{D}$. Moreover, for any $z_0 \in U$, there is a* unique *biholomorphism f such that*

$$f(z_0) = 0 \qquad \textit{and} \qquad f'(z_0) > 0$$

This theorem is certainly equivalent to the previous statement: to get a conformal equivalence between U and V, just compose the biholomorphisms $U \to \mathbb{D}$ and $\mathbb{D} \to V$. This statement is a bit nicer because it also allows us to state a uniqueness result. As it turns out, the uniqueness result is much easier to prove.

Proof of uniqueness. Suppose that $f_1, f_2 : U \to \mathbb{D}$ are biholomorphisms with

$$f_1(z_0) = f_2(z_0) = 0 \qquad\qquad f_1'(z_0) > 0, f_2'(z_0) > 0$$

Then $F = f_1 \circ f_2^{-1}$ is a biholomorphism $\mathbb{D} \to \mathbb{D}$, and

$$F(0) = f_1(f_2^{-1}(0)) = f_1(z_0) = 0$$

So by the classification, we know that F must be a rotation, namely $F(z) = e^{i\theta} z$, and thus $F'(0) = e^{i\theta}$. Moreover, we know that

$$F'(0) = (f_1(f_2^{-1}(0)))' = f_1'(z_0)\frac{1}{f_2'(z_0)} > 0$$

Since $e^{i\theta}$ is not a positive real number unless $\theta = 0$, we conclude that F must be the identity map, so that $f_1 = f_2$. □

Proof of existence. This proof is very hard, and we have no chance of actually doing it. The (very rough) outline of the steps is as follows. Using the fact that U is simply connected and not all of $\mathbb{C}$, it turns out that we can uniquely define a logarithm on it. By using this logarithm, we can define an injective analytic map f on U whose image is bounded. By rescaling and translating, we can make sure that this image is actually a subset of $\mathbb{D}$. Now, we look at the set of all injective analytic maps $U \to \mathbb{D}$, which we now know is non-empty. Using a powerful compactness result from analysis (namely the Arzelà-Ascoli Theorem), together with a result on the "equicontinuity" of families of analytic functions (called Montel's Theorem), we can prove that this family of functions has a "biggest" one, where bigness is measured by $|f'(z_0)|$. Finally, one has to argue that this biggest function is surjective onto $\mathbb{D}$, for which one uses the Schwarz Lemma and the Maximum Modulus Principle; if it weren't surjective, we'd be able to use these tools to build an even "bigger" function. This function is the desired conformal equivalence $U \to \mathbb{D}$. □