A COMPARISON OF THE CLASSIFICATION OF SURFACES

CARLOS AZPURUA

Abstract. In this expository paper, we discuss the classification of Euclidean surfaces to introduce the Killing-Hopf theorem. We then introduce hyperbolic geometry and its isometries, so that we may consider a difference in the classification of Euclidean and hyperbolic surfaces. We end with a discussion on Riemann surfaces and briefly introduce the Uniformization theorem.

Contents

1. Surfaces of local geometries and Killing-Hopf

The goal of this section is to answer the question of which surfaces locally look like one of the "three basic two-dimensional geometries-" Euclidean, spherical, and hyperbolic space. By euclidean space, we of course just mean $\mathbb{R}^2$. We will not go into detail about spherical space, but for our purposes we can think of it as just the 2-sphere, with an inherited metric from $\mathbb{R}^3$. We will formally introduce hyperbolic space in section 2.

Definition 1.1. A euclidean surface is a set S with a function $d_S\colon S \times S \to \mathbb{R}$ that is locally euclidean. That is for all $x \in S$, there exists an $\epsilon > 0$ where the neighborhood $N_\epsilon(x) = \{x' \in A \mid d_S(x, x') < \epsilon\}$ is isometric to a euclidean disk, i.e. there exists a bijection f between $N_\epsilon(x)$ and a disk $D \subset \mathbb{R}^2$ where for all $y, z \in N_\epsilon(x)$, we have that $d_S(y, z) = d(f(y), f(z))$.

The above definition can be generalized to spherical and hyperbolic space by instead requiring an isometry to a disk of the respective space at each point. When we refer to a surface, we mean a set that is locally like either euclidean, spherical or hyperbolic space in the sense of definition 1.1, unless we specifically mention which kind of surface it is. For sections 1 and 2 this kind of surface will be the main object of interest. The following closely related definition will seem very abstract at first, but we will see how it arises naturally using an example from Euclidean space.

Definition 1.2. Let Γ be a group acting on a metric space T. Define the quotient of T by Γ, T/Γ, to be the set of all sets of the form $\Gamma p = \{g(p) \mid g \in \Gamma\}$ where $p \in T$.

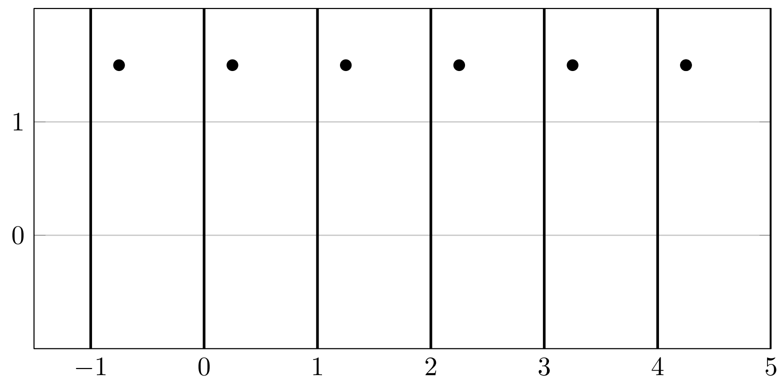

FIGURE 1. The Γ Orbit of the point $(\frac{5}{4}, 1)$

That is, $T/\Gamma = \{\Gamma p \mid p \in T\}$. We call Γp the $\Gamma-$orbit of p. Define $\pi\colon T \to T/\Gamma$ by $\pi(p) = \Gamma p$.

It can easily be seen that by the properties of a group that this procedure partitions T into equivalence classes, namely the sets Γp. However, this is not enough to guarantee that T/Γ is a surface, or that it inherits the geometrical properties of whatever T is, so we will need to impose more conditions, so of course it will be better to work with an example at this point.

Example 1.3. *The Cylinder, C, as a euclidean surface*

What motivates us to check that the cylinder is euclidean is that it can be made by joining a pair of opposite sides of a sheet of paper (which can be thought of as a piece of the plane). More formally, this would be like taking two parallel lines, say $x = 0$ and $x = 1$, and identifying every point in the region bounded by these lines with a point in C. Yet this would mean we would have to identify points on $x = 0$ and $x = 1$ with the same collection of points in the cylinder, because it would otherwise not geometrically "look" like a cylinder. Points on $x = 0$ and $x = 1$ would then be distinguished from the ones in the interior of the region. This conflicts with the rotational symmetry of the cylinder, so it would be best to adjust our model to treat all points of the cylinder equally, which we can do by instead using all the points of $\mathbb{R}^2$. Seeing that C has a this nice "rectangular" representation, we will use this representation to tile the plane. We will identify a point on C with a collection of points of $\mathbb{R}^2$. If a point (x, y) is in the collection of points P which corresponds to a point in C, then points of the form $(x + n, y)$ where $n \in \mathbb{Z}$ should also be in P. So a point of C is a set of the form $\{(x + n, y) \mid n \in \mathbb{Z}\}$.

This is exactly a partition of $\mathbb{R}^2$ into equivalence classes by the group of integer translations in the x-direction. Using the language from definition 1.2, if Γ is this group of translations, a point of $\mathbb{R}^2/\Gamma$ (that is, a point of C) is a set of the form $\{g(p) \mid g \in \Gamma\}$ for some $p \in \mathbb{R}^2$. The Γ orbit of $(\frac{5}{4}, 1)$ is given above in figure 1. To give C a notion of distance, recall that points in C are collections of points in $\mathbb{R}^2$ and that for two points $Q, R \in C$, we let

$$d_C(Q, R) = \min\{d(q, r) \mid q \in Q, r \in R\} \tag{1.0.1}$$

where d is the euclidean distance in $\mathbb{R}^2$. If we fix some $q' \in Q$, the right hand side above is equal to

$$\min\{d(q', r) \mid r \in R\} \tag{1.0.2}$$

because each $q \in Q$ has the same set of distances to the points in R. This can be seen as follows: if some $q_1 \in Q$ is some distance from some $r_1 \in R$, then for some $q_2 \in Q$ there exists $g \in \Gamma$ with $g(q_1) = g_2$, and because g is an isometry, $d(q_1, r_1) = d(g(q_1), g(r_1)) = d(q_2, g(r_1))$. Yet $g(r_1) \in R$ by definition, so expression (1.0.2) is the same regardless of the choice of $q' \in Q$ and is equal to (1.0.1). It is easy to see from (1.0.2) that the minimum does exist, because there is a nearest $r \in R$ to q'. Thus the metric d_C is well defined. It can be seen that for any point in C, there exists a neighborhood (neighborhoods of radius less than 1/2) where d_C is equal to the Euclidean metric. The metric d_C being defined and euclidean isn't the case for every isometry group, but we can generalize by seeing that this in part due to some nice properties of the translation group, which we will now introduce.

Definition 1.4. A group Γ acting on a metric space T is *discontinuous* if no $S \in T/\Gamma$ has a limit point in T. That is, there is no point p in T and $S \in T/\Gamma$ such that for any neighborhood U around T, we have that $(U/\{p\}) \cap S$ is non empty. Moreover, Γ is *fixed point free* if for all $g \in \Gamma$ with $g \neq id.$, for all $p \in T$, $g(p) \neq p$.

These two properties will be sufficient to guarantee that when an isometry group of euclidean, spherical, or hyperbolic space has them, its quotient of its respective space will have the local geometry of that space. Discontinuity helps guarantee that the metric in (1.0.1) is well defined by ensuring that a minimum exists. If some Γ orbit R had a limit point, q then if for example $q \notin R$, the set $\{d(q, r) \mid r \in R\}$ would not contain its infimum, and we would have two different Γ orbits with distance 0 from each other. There are also isometry groups whose quotients have a defined notion of distance, i.e. that of eq. (1.0.1), but are not everywhere locally euclidean, which is why we include the fixed-point free condition. For example, if Γ is the group generated by the rotation $r_{\pi/2}$ about the origin of $\mathbb{R}^2$, then any neighborhood around the origin contains multiple representatives from each $\Gamma-$orbit. This is problematic because then a circle of radius ϵ centered on the origin has circumference $\epsilon\pi/2$ in $\mathbb{R}^2/\Gamma$, as opposed to the euclidean value of $2\pi\epsilon$.

Theorem 1.5. *Let T be a metric space and Γ be a group of isometries of T which are discontinuous and fixed point free. Then each point $p \in T$ has a neighborhood U in which each point belongs to a different $\Gamma-$orbit, or in other words, π as defined in definition 1.2 is injective on U.*

Proof. Suppose that Γ is fixed point free and discontinuous. Let $p \in T$ and let $\pi(p) = S$. For the sake of contradiction, suppose that for every neighborhood U around p, there distinct points, q and r, in U, with $\pi(q) = \pi(r)$. We will show that p is a limit point of S, a contradiction. Let U be a neighborhood of p of radius δ. Let the neighborhood around U with radius $\frac{\delta}{2}$ be called V. Then there exist distinct points q, r in V with $\pi(q) = \pi(r)$. Then there exists some $g \in \Gamma$ with $g(q) = r$ and $g \neq id.$. By the triangle inequality, we have that

$$d(p, g(p)) < d(p, r) + d(r, g(p)).$$

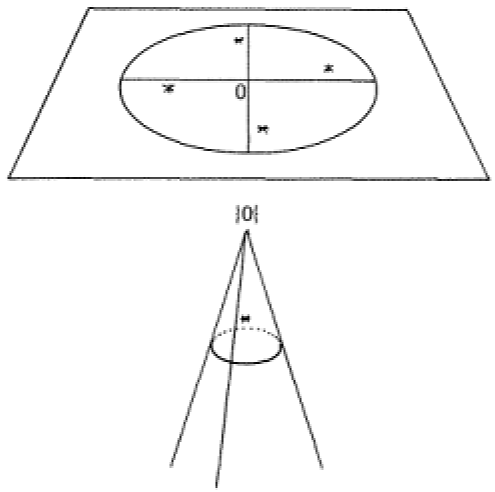

FIGURE 2. Quotient of $\mathbb{R}^2$ by $r_{\pi/2}$ (image from Stillwell Section 2.4)

Because g is an isometry, $d(p, q) = d(g(p), g(q)) = d(g(p), r)$, so by this and the fact that $r, q \in U$, we have that

$$d(p, g(p)) < d(p, r) + d(p, q) < \frac{\delta}{2} + \frac{\delta}{2} = \delta,$$

so $g(p) \in U$. Since Γ is fixed point free, $p \neq g(p)$, so p is a limit point of S, a contradiction, so π is injective on some neighborhood of p.

□

From here on we restrict T to be one of our three spaces of interest (euclidean, spherical, and hyperbolic). Before we discuss the converse of theorem 1.5, we discuss a fact that we will use without rigorous proof. Specifically, that for any isometry of the three spaces, if the isometry is the identity on the neighborhood of a point, then it is the identity on the whole space. This can be seen by observing that for any of theses spaces, circles (sets which are a set of points with fixed distance from a chosen point), behave nicely in that 3 or more circles who all intersect, can only do so at exactly one point. Then, if an isometry is an identity on a neighborhood of a point, any point outside the neighborhood must map to itself, because if we consider any circle through this point with center in the neighborhood (which there are infinitely many of), these circles must map to themselves, and thus the point outside the neighborhood must be fixed because it is the only point in common with all these circles.

Theorem 1.6. *Let T be either euclidean, hyperbolic, or spherical space, and Γ be a group of isometries of T such that each point $p \in T$ has a neighborhood U in which each point belongs to a different $\Gamma-$orbit. Then Γ is discontinuous and fixed-point free.*

Proof. Suppose for the sake of contradiction that Γ is not discontinuous. Then there is some $p \in T$ that is a limit point of its $\Gamma-$orbit, $\pi(p) = S$. Then every neighborhood around p contains a point $q \neq p$ in S, a contradiction that there exists a neighborhood around p in which each point belongs to a different $\Gamma-$orbit. For the sake of contradiction, suppose that Γ is not fixed point free. Let p be a fixed point of some $g \in \Gamma$ that is not the identity. We know that g is not the identity on any neighborhood of p, because otherwise it would be the identity on the whole space. So for any neighborhood U of size δ around p, there exists some $r \in U$ that does not map to itself under g. Thus we have that $d(p, g(r)) = d(g(p), g(r)) = d(p, r) < \delta$, meaning that $g(r) \in U$, a contradiction that there exists a neighborhood around p where each point maps to a different $\Gamma-$orbit, so Γ is fixed point free. □

Theorem 1.7. *The quotient of* $\mathbb{R}^2$ *by a fixed point free, discontinuous group* Γ *is a euclidean surface*

Proof. We equip $\mathbb{R}^2/\Gamma$ with the distance metric (1.0.1), which again for some $P, Q \in T/\Gamma$ is equivalent to $\min\{d(q, r) \mid r \in R\}$ where $q \in Q$. The set $\{d(q, r) \mid r \in R\}$ is bounded below, so it has an infimum, d. Suppose it doesnt contain its infimum, then for any $\epsilon > 0$, there are infinitely many $x \in R$ with $d < d(q, x) < d + \epsilon$. However, the set of such x is bounded and infinite, so it must have a limit point by Bolzano-Weierstrass. This contradicts the continuity of Γ, so d must be in $\{d(q, r) \mid r \in R\}$, and d is of course the minimum, so this metric is well defined on $\mathbb{R}^2/\Gamma$.
Now let $\Gamma p \in \mathbb{R}^2/\Gamma$. Then by theorem 1.5 there exists a neighborhood U around p with radius ϵ where every point belongs to a different $\Gamma - orbit$. We claim that on the neighborhood V around p of radius $\frac{\epsilon}{3}$, that $\pi \colon \mathbb{R}^2 \to \mathbb{R}^2/\Gamma$ is an isometry. Let $q, r \in V$. By definition, $d(q, r) \geq d(\Gamma q, \Gamma r)$. Suppose for the sake of contradiction that $d(q, r) > d(\Gamma q, \Gamma r)$. Then there is some $r' \in \Gamma r$ with $r' \neq r$ such that $d(q, r') < d(q, r)$. By two applications of the triangle inequality,

$$d(p, r') \leq d(p, q) + d(q, r') < \frac{\epsilon}{3} + d(q, r) < \frac{\epsilon}{3} + d(q, p) + d(p, r) < \epsilon.$$

So we have that $d(p, r') < \epsilon$, so $r' \in U$ contradicting that r is the only member of its quotient group in U. So $d(q, r) = d(\Gamma q, \Gamma r)$, so π is an isometry, and $\mathbb{R}^2/\Gamma$ is a euclidean surface. □

This characterization of fixed-point free, discontinuous isometry groups of our three spaces of interest is nice because we now know that a quotient of one of theses spaces by an isometry group not satisfying these conditions cannot be a surface with the local properties we desire. That is, if an isometry group isn't fixed point free and discontinuous, then the quotient by this group of its respective space will fail to have the locally geometry of that space, meaning there will be a point in the quotient space where no neighborhood is isometric to the respective space (the quotiented space). One can use the classifications of isometries of euclidean space in fact, to simply list all possible objects that result from quotienting $\mathbb{R}^2$ by a fixed point free, discontinuous group. Because rotations always have fixed point, the only possible generators for Γ are glide-reflections and translations. In fact, it can be shown that fixed-point free discontinuous isometry groups of $\mathbb{R}^2$ are generated by only one or two elements. The previous fact allows us to work case by case on the possible generators, resulting in the following theorem.

Theorem 1.8. *The quotient of* $\mathbb{R}^2$ *by a fixed point free, discontinuous group* G *is either a cylinder, twisted cylinder, torus, or Klein bottle.*

Surface $\mathbb{R}^2/\Gamma$	Generators of Γ
Cylinder	Single translation
Twisted Cylinder	Single Glide reflection
Torus	Two translations (that are not parallel)
Klein bottle	One translation and glide reflection

Another advantage of the use of quotient groups is that we have a model for objects like the Klein bottle, which cannot be embedded in 3-D euclidean space without self-intersection. Now we will discuss some theorems about surfaces themselves. It should be noted that not all surfaces with the local geometric structure of definition 1.1 are a quotient of a discontinuous fixed point free group (like the union of two disjoint parallel planes), but sufficiently "nice" surfaces are.

Definition 1.9. A surface S is *connected* if for any two $A, B \in S$, there exists a finite sequence of points starting with A and ending in B such that consecutive points lie together in a euclidean disk of S. We refer to this sequence of points as a *polygonal path* from A to B. S is *complete* if any line segment in S can be continued indefinitely.

Remark 1.10. When we refer to a line segment on a euclidean surface, we mean a polygonal path whose sides are entirely contained in euclidean disks where successive sides meet at an angle of π.

Connectedness excludes the parallel plane example, and completeness excludes for example, the plane minus a point, which is a euclidean surface, but is geometrically inconvenient because every line segment cannot be continued indefinitely. We now introduce a more general definition for the notion of a surface, so that we can discuss an important classification theorem.

Definition 1.11. A *Riemannian Manifold* is a real smooth manifold M with a smooth inner product structure g, that is, a family of functions defined at each point p from the Cartesian product of the tangent space T_pM at p with itself to $\mathbb{R}$. That is, g_p takes in two tangent vectors at p and returns a scalar (just like a dot product). By the inner product being smooth, we mean that if X and Y are vector fields on M, then the function defined by $p \mapsto g_p(X|_p, Y|_p)$ is smooth. We call this family of functions g a *Riemannian metric tensor* on M.

Note that having a metric tensor induces an ordinary metric on the space, as we can use the metric tensor to integrate along a curve and get a length. A metric tensor in fact provides more structure than a regular distance metric because it allows us to measure tangent vectors, angles between them, and thus define an inner product, structure which will give rise to the geometry of the space. For example curvature, a quantity that will appear in the following theorem, is something we can define based off of this structure.

Theorem 1.12. ***(General Killing-Hopf)*** *Any complete, connected Riemannian Manifold of constant curvature is the quotient of either euclidean, hyperbolic, or spherical space by a discontinuous, fixed point free group of isometries.*

In the case of euclidean surfaces, this translates to the fact that any complete, connected euclidean surface is of the form $\mathbb{R}^2/\Gamma$ for some fixed point free, discontinuous group Γ. Moreover, we have that these surfaces are exactly those described in theorem 1.8, if the surface is not $\mathbb{R}^2$ itself. We shall see that the classification of hyperbolic surfaces however, is not as straight forward. Though we have Killing-Hopf for hyperbolic surfaces, i.e. that those surfaces are all the quotient of hyperbolic space by some fixed point free discontinuous group of isometries Γ, the classification of all such Γ is more difficult because there is another type of hyperbolic isometry that can be included in Γ, *limit rotations*.

2. Hyperbolic Space and its Isometries

The purpose of this section is to in brief, introduce hyperbolic geometry, and explain how its difference in isometries lead to a different in the approach of the classification of hyperbolic surfaces. Hyperbolic geometry can be introduced through a discussion on curvature, and its distinction from spherical and euclidean geometry in that it is a surface of constant negative curvature (where as the other two have positive, and 0 curvature, respectively). However, this is outside of scope of this paper so we begin with some models of hyperbolic space and a discussion of its metric metric.
From here on, we denote the open unit disk by $\mathbb{D}^2 = \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 < 1\}$, and the open upper half plane by $\mathbb{H}^2 = \{(x, y) \in \mathbb{R}^2 \mid y > 0\}$. We refer to $\mathbb{D}^2$ as the *conformal disk model* and to $\mathbb{H}^2$ as the *half-plane model*. Both of these serve as models of hyperbolic space, and there is no "difference" between each space equipped with the proper metric. The difference arises in working with the models, where things that might seem obvious in one model, might not seem obvious in the other model. They can be equipped with a metric tensor, as in definition 1.11. In the case of $\mathbb{H}^2$, we have the following:

Definition 2.1. The *Poincaré metric* is a metric tensor on $\mathbb{H}^2$ is given by

$$ds = \frac{\sqrt{dx^2 + dy^2}}{y}.$$

A consequence of this metric is that as we approach the x-axis, points which are near each other vertically in the euclidean sense, become further and further apart, while the x-axis is at a distance of infinity. One thing that still "looks" euclidean, however, is angle. We can use the ratios of the sides of infinitesimal right triangles to define angle, and observe that the Poincaré metric is equal to the euclidean distance metric divided by what is in the case of an infinitesimal triangle, a constant, y, so the ratios in hyperbolic space are equal to the those in euclidean space. It will be much easier to express hyperbolic isometries in terms of complex functions, in which case the above metric is simply

$$ds = \frac{|dz|}{\mathrm{Im}(z)}.$$

To see that $\mathbb{H}^2$ and $\mathbb{D}^2$ are models of the same thing, we define a bijection $J\colon \mathbb{H}^2 \to \mathbb{D}^2$ by inverting the complex plane in the circle centered at $-i$ with radius $\sqrt{2}$, and then reflecting over the x-axis. We can express J as a complex function like so:

$$J(z) = \frac{iz + 1}{z + i}.$$

The description of J as a composition of an inversion and reflection is convenient because we then know that angles are also preserved in $\mathbb{D}^2$, hence the name conformal disk model. While we could build up to a metric "from scratch" on $\mathbb{D}^2$,

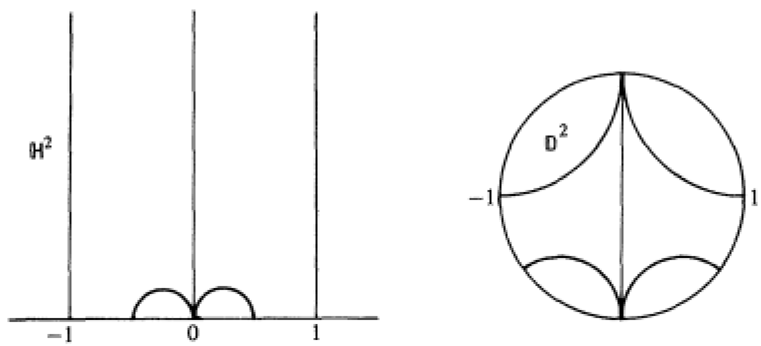

FIGURE 3. Our two models and some lines (figure from Stillwell section 4.2)

since we already know it is a model of hyperbolic space we can use the bijection we already have to define a distance metric on $\mathbb{D}^2$ by taking the distance between $w_1, w_2 \in \mathbb{D}^2$ to be the distance of their preimages $J^{-1}(w_1), J^{-2}(w_2)$ in $\mathbb{H}^2$. For $w = J(z)$, this results in that

$$ds = \frac{dz}{\text{Im}(z)} = \frac{2dw}{1 - |w|^2}.$$

Note that as $|w|$ approaches 1, the distance of $|w|$ from the origin approaches infinity. We should then verify that J does send $\partial\mathbb{H}^2$, i.e. points on the x-axis to the boundary of the unit circle. Let $a \in \mathbb{R}$, then we have that

$$|J(a + 0i)| = |\frac{ai + 1}{a + i}| = |\frac{(ai + 1)(a - i)}{a^2 + 1}| = |\frac{2a + (a^2 - 1)i}{a^2 + 1}| = 1.$$

So the boundary of $\mathbb{D}^2$ plays the role of the x-axis plays in the half-plane model in that it is both a boundary, and a collection of points (that aren't actually in $\mathbb{D}^2$) that are "off at infinity". It will soon be relevant to understand what lines look like in hyperbolic space, but we should first establish what we mean by a line.

Definition 2.2. A line is a set of points equidistant from two other points.

In the $\mathbb{H}^2$ model there are two possibilities, euclidean circles which are orthogonal to $\partial\mathbb{H}^2$, and vertical lines. In the $\mathbb{D}^2$ model there are also two cases, arcs of euclidean circles orthogonal to $\partial\mathbb{D}^2$, and diameters of the disk.

We now have the foundation necessary to discuss and classify the isometries of hyperbolic space. It is obvious that rotating around the origin in the $\mathbb{D}^2$ model leaves its metric invariant, because $|w|$ is unchanged, this gives the so-called $\mathbb{D}^2-$rotations. An obvious type of isometry on $\mathbb{H}^2$ is of the form $t_\alpha(z) = \alpha + z$ where $\alpha \in \mathbb{R}$. This is strictly a horizontal translation of the plane. This kind of isometry is called a *limit rotation* because in $\mathbb{H}^2$ it permutes vertical lines, though it sends horizontal lines (which are not actually lines in $\mathbb{H}^2$) to themselves. Though this kind of isometry does not fix any points in hyperbolic space, when considered in $\mathbb{D}^2$, it fixes i, a point on $\partial\mathbb{D}^2$. Note that although $J^{-1}(i)$ is empty, in some sense it fixes a point of $\partial\mathbb{H}^2$, the point at "infinity" of this line, which would map to i if we extended J to

$\hat{\mathbb{C}}$, the complex projective plane.
Another obvious type of isometry on $\mathbb{H}^2$ is called a translation, and is of the form $d_p(z) = pz$ for $p > 0$. It is a translation more-so than the previous kind of isometry in the sense that it sends the y-axis to itself, so it has an axis of translation which is actually an $\mathbb{H}^2$ line.

Now we consider reflections. There is an obvious candidate in $\mathbb{H}^2$, reflection over vertical lines. However, there is another kind of reflection, corresponding to another kind of $\mathbb{H}^2$ line, euclidean inversions in circles perpendicular to $\partial\mathbb{H}^2$. To see that this is indeed an isometry, we will need to first work in $\mathbb{D}^2$. Consider reflection over the real line, $r(w) = \overline{w}$, which is an isometry of $\mathbb{D}^2$ because it leaves $|w|$ unchanged. Yet the real line in the $\mathbb{D}^2$ model corresponds to the unit circle in in $\mathbb{H}^2$. Moreover, one can verify that the conjugate by J^{-1} of $r(w)$ is inversion in the unit circle. That is,

$$J^{-1}rJ(z) = \frac{1}{\overline{z}},$$

but not only do we now have that inversion in the unit circle is an isometry of $\mathbb{H}^2$, by the previous two types of isometries, inversion in any circle with center on the x-axis and any radius is an isometry. This makes up all the $\mathbb{H}^2$ reflections, so we can discuss the final type of isometry then the glide reflection. This is the product of a reflection with a translation whose axis is the line of reflection. This is either the composition of the reflection of a vertical line with a $\mathbb{H}^2$ translation, that is, dilation, or it is simply an inversion in a semi-circle in $\mathbb{H}^2$, the other type of reflection we have discussed. This kind of isometry has one invariant line and fixes two points, but its fixed points are not actually in hyperbolic space, i.e. they are in $\partial\mathbb{H}^2$ or $\partial\mathbb{D}^2$, depending on the model you are working with.

Theorem 2.3. *Every hyperbolic isometry is either a rotation, a limit rotation, a translation, or a glide-reflection. Moreover, the only isometries that are not fixed point free are rotations.*

Thus the only kind of hyperbolic isometry which fixes a point in hyperbolic space itself are rotations. This complicates things for the classification of all possible complete, connected hyperbolic surfaces, because the resulting possibilities for fixed point free discontinuous groups are quite vast.

3. Hyperbolic Surfaces

At this point, we give up the approach of theorem 1.8, where we were able to list all possible fixed-point free discontinuous groups Γ because of the properties of Euclidean space, and then use Killing-Hopf to establish that the quotient of the plane by these groups are indeed the only possible Euclidean surfaces. For hyperbolic space, it will simply be easier to find the surfaces directly, so we introduce a new object called a *hyperbolic polygon*, in order to construct hyperbolic surfaces.

Definition 3.1. A set $\Pi \subset \mathbb{H}^2$ is called a *hyperbolic polygon* if there exists a finite consecutive sequence of $\mathbb{H}^2$ line segments and segments of $\partial\mathbb{H}^2$ that bound Π and are in Π. We call the line segments in $\mathbb{H}^2$ *proper* edges and the segments of $\partial\mathbb{H}^2$ *improper*. We similarly call vertices of these line segments proper if they lie in $\mathbb{H}^2$, while those in $\partial\mathbb{H}^2$ are called improper.

The idea is to paste edges of the polygon together, in a similar way to how in the representation of a cylinder as a quotient of $\mathbb{R}^2$, we pasted together the lines $x = 0$ and $x = 1$ as in the paper analogy.

Definition 3.2. An *edge pairing* of a hyperbolic polygon Π is a partition of the proper edges into pairs $\{e, e'\}$ of equal length (where length may be infinite), with an $\mathbb{H}^2$ isometry $g_{e,e'} \colon e \to e'$ for each pair. Points $p \in e$ and $p' \in e'$ are said to be *identified* if $g_{e,e'}(p) = p'$.

In the case that p is a vertex, it is possible that p is identified with some p', which is identified with some $p'' \neq p$. We say that p and p'' are also identified, and we call a set of vertices containing an entire such chain $\{v_1, \dots, v_k\}$ a *vertex cycle*. Note that this partitions the vertices of Π into equivalence classes, which are the vertex cycles.

Definition 3.3. The *identification space* S_Π of an edge pairing of a polygon Π is a set whose points are

(1) the interior points z of Π
(2) pairs $\{p, p'\}$ of interior points of proper edges that are identified
(3) the vertex cycles of the proper vertices of Π.

With the following definition, we finish make the idea of pasting the edges of the hyperbolic polygon rigorous.

Definition 3.4. If S_Π is the identification space of some hyperbolic polygon Π, and $A, B \in S_\Pi$ we define a polygonal path P from A to B in the identification space as a finite sequence of k paths $\{P_n\}$ in Π with the following properties

(1) P_1 starts at A and p_k ends at b
(2) For all $1 \leq i \leq k$, if P_i ends at u, and P_{i+1} starts at v, then v and u are identified.

Moreover, we define the length of P to be $|P| = \sum_{n=1}^{k}(|P_n|)$, that is, the sum of the lengths of the individual paths.

Theorem 3.5. *For a hyperbolic polygon* Π*, if the angles of each vertex cycle sum to* 2π*, its identification space* S_Π *is a hyperbolic surface. In particular we equip the identification space with the metric*

$$d_s(A, B) = \inf\{|P| \mid P \textit{ is a polygonal path from } A \textit{ to } B\}$$

At this point, for some identification space S_Π we cannot yet use Killing-Hopf to guarantee a corresponding group Γ such that $\mathbb{H}^2/\Gamma$ is isometric to S_Π. This is because Killing-Hopf requires that our surface be complete, which we do not yet know. The following theorem provides a criteria that will guarantee us this; so long as the sides of Π are finite, then S_Π is complete.

Theorem 3.6. *If* Π *is compact, then* S_Π *is complete.*

Thus by Killing-Hopf we have that any compact polygon Π has a corresponding group Γ where $\mathbb{H}^2/\Gamma$ is isometric to S_Π. A natural question that arises is whether the converse is true, whether for every quotient group, (and thus every complete, connected, hyperbolic surface), there exists a corresponding S_Π. It turns out that this is the case for any compact hyperbolic surface.

Theorem 3.7. *For any compact hyperbolic surface* $\mathbb{H}^2/\Gamma$*, there exists a hyperbolic polygon* Π *whose identification space is* $\mathbb{H}^2/\Gamma$*.*

Proof. For some $p \in \mathbb{H}^2$, define the *Dirichlet region* $D(p)$ with center p to be

$$D(p) = \{q \in \mathbb{H}^2 \mid d_{\mathbb{H}^2}(q,p) \leq d_{\mathbb{H}^2}(q,g(p)) \text{ for all } g \in \Gamma\},$$

That is, $D(p)$ is the collection of points that are closer to p than any other member of p's Γ orbit, $\Gamma p = \{g(p) \mid g \in \Gamma\}$. Note that $D(p)$ contains at least one element of every Γ orbit, and its interior contains at most one element from $\mathbb{H}^2/\Gamma$, hinting that this is indeed the hyperbolic polygon we are looking for.
We can express $D(p)$ as the intersection of the collection of closed half planes of the form

$$H_g(p) = \{q \in \mathbb{H}^2 \mid d_{\mathbb{H}^2}(q,p) \leq d_{\mathbb{H}^2}(q,g(p))\},$$

that is, $D(p) = \bigcap_{g\in\Gamma} H_g(p)$, so $D(p)$ is also closed and convex (because half-planes are convex), and has some boundary $\partial D(c)$. This boundary contains at most one segment from each of the lines bounding the above half planes. That is, the lines

$$L_g(p) = \{q \in \mathbb{H}^2 \mid d_{\mathbb{H}^2}(q,p) = d_{\mathbb{H}^2}(q,g(p))\}.$$

Recall that from definition 2.2, $L_g(p)$ is indeed a line, as it is the set of points equidistant from p and $g(p)$. We wish to show that for only finitely many g, $L_g(p) \cap \partial D(p) \neq \emptyset$, because then we will have a hyperbolic polygon.
Because $\mathbb{H}^2/\Gamma$ is compact, $D(p)$ is compact. So there exists a hyperbolic disk U of radius r with $D(p) \subset U$. If a line $L_g(p)$ passes within a distance r of p, i.e., there exists a point $z \in L_g(p)$ such that $d_{\mathbb{H}^2}(z,p) = d_{\mathbb{H}^2}(z,g(p))$ and $d(z,p) < r$, then by the triangle inequality, $d_{\mathbb{H}^2}(p,g(p)) < 2r$. Then infinitely many lines $L_g(c)$ cannot intersect with $\partial D(p)$, because otherwise there will be infinitely many points of the Γp within $2r$ of p, meaning that Γp will have a limit point, a contradiction that Γ is discontinuous. □

Notice that the above proof doesn't actually depend on any properties of hyperbolic geometry. In fact, it is the case that for any compact surface $\mathbb{R}^2/\Gamma$, $\mathbb{S}^2/\Gamma$ (the sphere), or of course $\mathbb{H}^2/\Gamma$, there is a corresponding polygon, with the exception of $\mathbb{S}^2$ itself.

4. Uniformization Theorem, Riemann Surfaces

In this section we introduce another type surface, the Riemann surface (which is different from the Riemannian Manifold!). We will introduce another classification theorem, and see how the language of quotient spaces is useful for this new object.

Definition 4.1. A *Riemann surface* is a connected complex manifold of complex dimension 1. In particular, the transition maps between two overlapping charts are required to be holomorphic.

Definition 4.2. A *conformal* transformation is a map that preserves angles locally. We say that two sets are *conformally equivalent* if there exists a conformal transformation from one to the other (and thus vice versa).

Theorem 4.3. *(**Uniformization Theorem**) A simply connected Riemann surface is conformally equivalent to one of the following*

(1) $\mathbb{C}$, the complex plane
(2) $\hat{\mathbb{C}}$, the complex projective line (i.e. complex plane plus a point at infinity)
(3) $\mathbb{D}^2 \subset \mathbb{C}$, the open unit disk in the complex plane

The three above objects are of course familiar to us. The the first corresponds with euclidean space, the second with spherical space, and the last with hyperbolic space. It turns out that for Riemann surfaces, these three objects are the only possible universal covers, meaning that in some sense, these are the objects we should use to construct Riemann surfaces. In the following theorem, we apply the methodology of section 1, where we attempt to construct surfaces from the quotient of a group of automorphism of our spaces of interest.

Theorem 4.4. *Let Γ be a group of conformal transformations on one of $\mathbb{C}$, $\mathbb{P}^1$, or $\mathbb{D}^2$. If Γ acts discontinuously and is fixed point free, then the quotient of the respective space by Γ is a Riemann surface.*

Proof. Let T be on of $\mathbb{C}$, $\mathbb{P}^1$ or $\mathbb{D}^2$ and Γ be a conformal group action on T. Let $s \in T/\Gamma$. Choose some $p \in \pi^{-1}(s) = s$. By theorem 1.5 there exists a neighborhood U around p such that every point of U belongs to a different $\Gamma-$orbit. We claim that $\pi\colon (U) \to \pi(U)$ is a homeomorphism. It is of course a bijection. But we can also equip T/Γ with the same metric discussed in (1.0.1) without issue, so π is in fact and isometry and so is continuous. These will be the charts for our manifold, so we now wish to show that the transition maps are holomorphic. Let $(U_1, \varphi_1), (U_2, \varphi_2)$ be two charts of T/Γ with $U_1 \cap U_2$ nonempty (where $U_1, U_2 \subset T/\Gamma$). Because $\pi(\varphi_2(U_1 \cap U_2)) = \pi(\varphi_1(U_1 \cap U_2))$, for any $p \in U_1 \cap U_2$, there is a conformal map $g \in \Gamma$ such that $g(\varphi_1(p)) = \varphi_2(p))$. In fact, $g(\varphi_1(U_1 \cap U_2) = \varphi_1(U_1 \cap U_2)$, that is, g is independent of the choice of p, so $\varphi_1 \circ \varphi_2^{-1}\colon \varphi_2(U_1 \cap U_2) \to \varphi_1(U_1 \cap U_2)$ is holomorphic.

□

Though the above statement is weaker than Killing-Hopf, in that we only have that quotients are surfaces (and not vice versa), it still motivates us to ask a certain question: what are the automorpohisms of our spaces of interest? And specifically, which subgroups of the conformal group are the fixed point-free discontinuous ones? The answer to the first question is well known, and is presented in the theorem that follows.

Theorem 4.5. .

(1) *Every automorphism of $\mathbb{C}$ is of the form $f(z) = az + b$ where $a, b \in \mathbb{C}$ and $a \neq 0$.*
(2) *Every automorphism of $\hat{\mathbb{C}}$ is of the form $f(z) = \frac{az+b}{cz+d}$ where $a, b, c, d \in \mathbb{C}$ and $ad - bc \neq 0$*
(3) *Every automorphism of $\mathbb{D}^2$ is of the form $f(z) = \frac{az+b}{\bar{b}z+\bar{a}}$ where $a, b \in \mathbb{C}$ and $|a|^2 - |b|^2 = 1$*

Remark 4.6. Every automorphism of $\hat{\mathbb{C}}$ has at least one fixed point, so it is not possible construct a surface in this way from $\hat{\mathbb{C}}$. This corresponds with the fact that the only compact Riemann surface with universal cover $\hat{\mathbb{C}}$ are those of genus 0, so the sphere itself.

Acknowledgements

I owe a great deal of thanks to my mentor Randy Van Why, who with great patience guided me through the process of writing this paper. I would also like to thank Peter May for organizing the REU, and the faculty who volunteered to give talks this year.

References

[1] J. Stillwell, *Geometry of Surfaces,* Springer-Verlag, New York, 1992
[2] J.H. Hubbard, *Teichmller Theory And Applications To Geometry, Topology, And Dynamics. Volume 1 Teichmller Theory*, Matrix Editions, Ithaca, 2006
[3] J. Jost, *Compact Reimann surfaces*, Sringer-Verlag, Berlin, 2006