Lecture 3: Structured surfaces

September 9, 2014

The goal for today is to define what a Riemann surface is. We will also start to give some examples.

For future applications, as well as to isolate some structural features, we'll actually give a more broad definition, namely that of a $\mathcal{G}$-structured surface. The first thing to say is what kind of thing $\mathcal{G}$ is.

Definition 0.1. *An* admissible collection of gluing maps *is a set $\mathcal{G}$ of bijections $U \overset{\varphi}{\leftrightarrow} V$ between open subsets U and V of the plane $\mathbb{R}^2$, satisfying the following requirements:*

1. *If φ and ψ are two composable bijections both of which lie in $\mathcal{G}$, then their composition also lies in $\mathcal{G}$.*
2. *If φ lies in $\mathcal{G}$, then so does its inverse.*
3. *If $U \overset{\varphi}{\leftrightarrow} V$ lies in $\mathcal{G}$ and $U' \subset U$ is an open subset of U, then $V' = \varphi(U')$ is open in V, and the induced bijection $U' \overset{\varphi}{\leftrightarrow} V'$ also lies in $\mathcal{G}$.*
4. *If $U \overset{\varphi}{\leftrightarrow} V$ is a bijection between open subsets of $\mathbb{R}^2$, and if, for every point $p \in U$ there is an open subset $U' \subset U$ containing p such that $V' = \varphi(U')$ is open in V and $U' \overset{\varphi}{\leftrightarrow} V'$ lies in $\mathcal{G}$, then $U \overset{\varphi}{\leftrightarrow} V$ itself also lies in $\mathcal{G}$.*

The idea is that we will think of building surface by gluing together open subsets of $\mathbb{R}^2$. These elements of $\mathcal{G}$ are then the things we will allow ourselves to glue along. By the way, you should feel free to ignore the last two axioms at first blush. They're there to ensure "locality".

Here are some examples:

1. We can take $\mathcal{G}$ to be the collection of all *homeomorphisms* between open subsets of $\mathbb{R}^2$. Recall that these are the continuous bijections having a continuous inverse: equivalently, they are the bijections $U \overset{\varphi}{\leftrightarrow} V$ such that a subset $U' \subset U$ is open if and only if its image V' under φ is open in V. This is the maximal possible collection: every collection satisfying the above axiomatics is contained in this one.
2. We can take $\mathcal{G}$ to be the collection of all *diffeomorphisms* between open subsets of $\mathbb{R}^2$. Recall that these are the differentiable bijections with differentiable inverse.
3. We can take $\mathcal{G}$ to be the collection of all *biholomorphisms* between open subsets of $\mathbb{R}^2 = \mathbb{C}$. Recall that these are the holomorphic bijections with holomorphic inverse.

Note that class 3 is contained in class 2 is contained in class 1. But the next three classes, while all contained in class 3 (though this is not necessarily obvious), will not be in any way pairwise contained in one another. They are also much much smaller classes than the previous ones. Grossly speaking, the previous classes are infinite dimensional, but the next ones will be finite dimensional.

- We can take $\mathcal{G}$ to be the collection of *Euclidean isometries* between open subsets of $\mathbb{R}^2$. These are those bijections $U \overset{\varphi}{\leftrightarrow} V$ which can be obtained by composing translations and rotations. (Thus they preserve all distances, areas, oriented angles, etc.: the entire "metric geometry" of the plane.)

- We can take $\mathcal{G}$ to be the collection of *Hyperbolic isometries* between open subsets of $\mathbb{D} := D(0,1) \subset \mathbb{R}^2$. We will be more precise about what this means later. Suffice it to say that, while $\mathbb{D}$ certainly does inherit the Euclidean geometry from $\mathbb{R}^2$, it can also be equipped with a different geometry, the Hyperbolic one; and the hyperbolic isometries are those bijections between open subsets of $\mathbb{D}$ which preserve all the "metric geometry" derived from this hyperbolic geometry.

  In this geometry, the geodesics (paths of least resistance) are those segments of circles or lines in $\mathbb{D}$ which meet the boundary $\partial\mathbb{D}$ at right angles. Furthermore, compared to the Euclidean distance, the hyperbolic distance stretches greatly as you approach the boundary, which is in fact infinitely far away. Actually, in a precise sense there's more room in hyperbolic geometry than in Euclidean geometry, even though in our model the unit disk sits inside the plane.

- We can take $\mathcal{G}$ to be the collection of *Spherical isometries* between open subsets of $S^2 = \{v \in \mathbb{R}^3 : |v| = 1\}$. These are those bijections which are given by rotation around some axis in $\mathbb{R}^3$. Truth be told, as stated this does not fit our precise framework, because it's based on S^2 instead of open subsets of $\mathbb{R}^2$. We could make it fit our chosen framework by transporting between the two using the operations of stereographic projection away from a point on S^2.

Now we can make the main definition.

Definition 0.2. *Let $\mathcal{G}$ be an admissible collection of gluing maps. A $\mathcal{G}$-surface is a set X together with a collection of so-called* defining charts, *which are bijections $U \overset{\varphi}{\leftrightarrow} V$ with U a subset of X and V an open subset of $\mathbb{R}^2$, subject to the following conditions:*

1. *For every point $x \in X$, there is a defining chart $U \overset{\varphi}{\leftrightarrow} V$ with $x \in U$. [Covering condition.]*

2. *If $U \overset{\varphi}{\leftrightarrow} V$ and $U' \overset{\psi}{\leftrightarrow} V'$ are two defining charts, then the sets $\varphi(U \cap U')$ and $\psi(U \cap U')$ are open in V and V' respectively, and the induced bijection $\varphi(U \cap U') \leftrightarrow \psi(U \cap U')$ (given by $\psi \circ \varphi^{-1}$) lies in $\mathcal{G}$. [Compatibility condition.]*

The idea is this: we imagine the $\mathcal{G}$-surface X as being obtained by gluing the various sets $V \subset \mathbb{R}^2$ along the maps $\psi \circ \varphi^{-1}$, which are called the *transition maps*. Generally speaking, if there is some notion or structure which makes sense for open subsets of $\mathbb{R}^2$, and if that notion or structure is preserved by the bijections in $\mathcal{G}$, then this notion or structure makes sense for $\mathcal{G}$-surfaces. For example, when $\mathcal{G}$ is the collection of Euclidean isometries, we obtain the notion of a *Euclidean surface*. On a Euclidean surface, notions such as straight lines and angles make sense. We check them on any chart; the result is independent of the chart because the transition maps are Euclidean isometries.

The next order of business is to define the notion of $\mathcal{G}$-isomorphism between $\mathcal{G}$-surfaces. The motivation is the following. Consider a $\mathcal{G}$-surface X, meaning the set X together with its defining charts. We could modify the system of defining charts in an innocuous manner: for example, given one chart $U \overset{\varphi}{\leftrightarrow} V$, we can add another by simply restricting φ^{-1} to any open subset of V. This shouldn't "really" change the surface, but it doesn't give us literally the same surface according to the definition. Or, in a slightly different direction, we could change the set X by a bijection, i.e. we could just relabel its elements. This also shouldn't really change the surface. The most natural notion of isomorphism which accommodates these two examples is the following.

Definition 0.3. *Let X and X' be two $\mathcal{G}$-structured surfaces. (The defining charts are implicit.) An* isomorphism *between X and X' is a bijection $X \overset{f}{\leftrightarrow} X'$ satisfying the following condition:*

- *If $U \overset{\varphi}{\leftrightarrow} V$ is a defining chart for X and $U' \overset{\varphi'}{\leftrightarrow} V'$ is a defining chart for X', then $\varphi(U \cap f^{-1}(U')$ and $\varphi'(f(U) \cap U')$ are open in V and V' respectively, and the induced bijection*

$$\varphi(U \cap f^{-1}(U')) \leftrightarrow \varphi'(f(U) \cap U')$$

lies in $\mathcal{G}$.

If f were to be the identity map, this condition would be the same one as used in the compatibility condition in the definition of a $\mathcal{G}$-surface. Thus, a $\mathcal{G}$-isomorphism can be thought of as a combination of a relabeling of the set X, together with a new set of choice of charts which is compatible with the previous one, in that the transition maps lies in $\mathcal{G}$. Here are some basic properties of isomorphisms:

1. The identity map from any X to itself is always a $\mathcal{G}$-isomorphism.
2. The composition of two composable $\mathcal{G}$-isomorphisms is also a $\mathcal{G}$-isomorphism.
3. The inverse of a $\mathcal{G}$-isomorphism is also a $\mathcal{G}$-isomorphism.

The proof follows easily from the corresponding properties of the class $\mathcal{G}$. These properties are reminiscent of the identity, transitivity, and symmetry properties of equality; in total, they ensure that it is formally reasonable to think of two $\mathcal{G}$-isomorphic surfaces as being "the same". In general, we will only be interested in properties of $\mathcal{G}$-surfaces which are invariant under isomorphism. For example, the number of defining charts is not an invariant, but for Euclidean surfaces (and hyperbolic surfaces and spherical surfaces), the notions of angle and distance are invariant.

While the set of defining charts is not invariant, we can define a new notion of chart which is invariant. Namely, we can define a *chart* on X to be a bijection $U \overset{\varphi}{\leftrightarrow} V$, with U a subset of X and V an open subset of $\mathbb{R}^2$, which is compatible with the defining charts in the same sense as always (c.f. condition 2 in the definition of $\mathcal{G}$-structured surface). The identity map gives a $\mathcal{G}$-isomorphism between X equipped with this abstract notion of chart to X equipped with the old defining charts; furthermore the collection of these abstract charts is in a precise sense the "biggest possible" collection which defines the same structure on X.

Besides being invariant, this definition is useful for the following reason: while we may want to specify our surface X in some "minimal" way — using only a few charts, say — nonetheless for many purposes it's convenient to have a lot of freedom in choosing the chart, e.g. for the purposes of making calculations or investigating structure. Choosing charts is somewhat like picking coordinates in physics.

There's only one more bit of abstraction to discuss before we get to examples.

Definition 0.4. *Let X be a $\mathcal{G}$-surface. We say X is* Hausdorff *if for any two distinct points $x \neq x'$ on X, there are charts $U \overset{\varphi}{\leftrightarrow} V$ and $U' \overset{\varphi}{\leftrightarrow} V'$ on X with $x \in U$ and $x' \in U'$ and $U \cap U' = \emptyset$.*

Thus, in the Hausdorff case, any two distinct points can be separated by charts. This is a reasonable and useful condition to impose. Note that the above definition refers not to the defining charts, but only to the invariant notion of chart introduced above. For example, we need to have very small charts to make this hold for nearby points.

Now we can define what a Riemann surface is and start to give examples.

Definition 0.5. *A Riemann surface is a Hausdorff $\mathcal{G}$-surface, where $\mathcal{G}$ is the set of biholomorphisms between open subsets of $\mathbb{C}$.*

The first example is essentially tautological. If we equip $\mathbb{C}$ with the single defining chart given by the identity map $\mathbb{C} \leftrightarrow \mathbb{C}$, then we get a Riemann surface structure on $\mathbb{C}$. In this case the abstract notion of chart unravels to the following: a chart on $X = \mathbb{C}$ is a biholomorphism $U \leftrightarrow V$ where $U \subset X$ and $V \subset \mathbb{C}$.

More generally, any open subset of $\mathbb{C}$ becomes a Riemann surface in essentially the same way. In fact, if X is any Riemann surface, we can define the notion of open subset of X (it should be open in every chart); and any open subset of a Riemann surface is also a Riemann surface, by restricting the charts to the open subset. A particular example is the following: removing finitely many points from a Riemann surface always results in an open subset, hence another Riemann surface.

We only have time for one more example: the *Riemann sphere*, denoted $\mathbb{P}^1$. This is defined as follows. As a set, we take $\mathbb{P}^1 = \mathbb{C} \cup \{\infty\}$, the set of complex numbers together with a formally-added point labeled "∞". On this set we take two defining charts:

1. The identity map from $\mathbb{P}^1 \setminus \{\infty\}$ to $\mathbb{C}$. (This is the "chart around 0".)

2. The map $z \mapsto 1/z$ from $\mathbb{P}^1 \setminus \{0\}$ to $\mathbb{C}$, where we set $1/\infty = 0$. (This is the "chart around ∞".)

The transition map between these charts is the map $z \mapsto 1/z$ from $\mathbb{C} \setminus \{0\}$ to itself; this is a biholomorphism (it is holomorphic, and it is its own inverse), so we indeed have a $\mathcal{G}$-structured surface, for $\mathcal{G}$ the collection of biholomorphisms. The Hausdorff condition is also satisfied, since any two distinct points are either 0 and ∞, in which case they are separated by the defining charts, or they lie in one or the other of the charts, where they can be separated by ordinary open disks in $\mathbb{C}$. Thus $\mathbb{P}^1$ is a Riemann surface.

The motivation behind $\mathbb{P}^1$ is the following. Say for example that we're studying a holomorphic function f defined on all of $\mathbb{C}$. For many purposes, it's convenient to study not just the behavior of f near specific points of $\mathbb{C}$, but also its behavior "at ∞", meaning what happens to $f(z)$ when $|z| \to \infty$. The above charts capture the idea that the behavior of f "at ∞" should be equivalent to the behavior of the new holomorphic function $g(z) = f(1/z)$ at 0. By this definition, we can reduce the nebulous-sounding study of holomorphic functions at (or perhaps "near") ∞ to the much more tangible study of holomorphic functions at or near 0. For example, a polynomial of degree n can be said to have a pole of order n at ∞. (Please feel free to plug $1/z$ into your favorite polynomial to verify this.)