Thurston's Geometrisation Conjecture

Hunter Liu

June 2, 2022

Classification theorems and conjectures are present throughout nearly every field of mathematics and are immensely useful for understanding and simplifying a variety of problems. In an algebraic setting, several examples come to mind: the famous classification of simple groups (which hurts to even look at), the classification of finite abelian groups, and the structure theorem for finitely generated modules over a principle ideal domain are all examples of complete descriptions of a certain class of mathematical object. The simplest example is the classification of finite abelian groups, which allows us to express any finite abelian group (of which there are several, perhaps even many) just in terms of cyclic groups.

As (differential) geometers that abuse paretheses, we would be far more interested in classification theorems that involve manifolds and surfaces. The theorem that comes to mind is the classification of surfaces, and it (roughly) states that every compact surface without a boundary is homeomorphic to the connected sum of a sphere, some tori, and possibly a projective plane or a Klein bottle; the number of tori is equal to the genus. This simplifies the study of surfaces in the sense that some computations only need to be done for the torus, sphere, projective plane, and Klein bottle.

While this classification is great for topologists, it fails to capture many pieces of geometric data — the notions of distance, geodesics, and even smoothness are forgone. Naturally, we would like to know if there is a classification of Riemannian manifolds. The answer is yes: Thurston had conjectured this answer in the twentieth century, and it was famously proved by Perelman around the turn of the twenty-first century.

Remark: This paper assumes that the intended audience (intentionally left ambiguous) has a background in differential geometry; hence, some basic definitions will be omitted so the author can make the page limit. Informal remarks will be indicated by the word "remark" in boldface so that the reader knows which paragraphs to omit from the page count. The author has considered making the margins smaller, but doing so makes the paper ugly, a similarly undesirable outcome.

Old People

A major contributor to the development of differential geometry and the study of manifolds was Bertrand Riemann (the guy "Riemannian metrics", "Riemann surfaces", and "Riemann sums" are named after). One surprising theorem he proved is:

Theorem (Riemann Mapping Theorem)**:** Every simply connected subset open $U \subsetneq \mathbb{C}$ is conformal (complex diffeomorphic) to the open unit disc, $\mathbb{D}$.

Remark: This theorem is unintuitive and surprising because the disc is so even and exhibits many symmetries whereas other regions, such as Brazil or the Koch snowflake or a slit-disc like $\mathbb{D} \setminus [0, 1]$ have strange shapes or choking hazards near the boundary. Strengthening this theorem to include the boundary is thus difficult; Carathéodory gave a theorem where such a conformal map could be extended to a homeomorphism between the closures of U and $\mathbb{D}$ when the boundary of U is a piecewise smooth path.

This was proved in the mid-nineteenth century, and this was later generalised further to include any simply connected complex 1-manifolds by Poincare and his peers. This generalisation is known as the uniformisation theorem, and it carries some additional geometric information with it as well:

Theorem (Uniformisation Theorem)**:** Every simply connected complex 1-manifold is conformal to one of the following surfaces:

1. $\mathbb{C}$, the entire complex plane (with the standard Euclidean geometry)
2. $\mathbb{D}$, the open unit disc (with the hyperbolic geometry given by the Poincare disc model)
3. $\mathbb{C} \cup \{\infty\}$, the Riemann sphere (with the spherical geometry)

Notably, conformal maps do not necessarily preserve geometric structure. Double notably, this only talks about simply connected manifolds. This latter point is very important because there are lots of complex 1-manifolds that aren't simply connected, such as the torus. This can really jangle one's spangles, but it's worth bearing in mind that simply connected Riemann surfaces are much nicer to work with than those with holes in them: tools like Cauchy's theorem (the one about contour integrals) or Stokes' theorem much more accessible.

So what the heck? The reader expects this paper to discuss the classification of Riemannian manifolds (or something of the sort), yet this uniformisation theorem misses the vast majority of Riemann surfaces! Furthermore, there's some weird stuff going on with the geometry part — conformal maps aren't isometries, and they can mess with the metrics given to a manifold in a somewhat controlled, but still problematic way.

Slinkies, Shrimps, and Universal Covers

Thinking back to the classification of surfaces mentioned in the introduction, we can see that there are three basic "building blocks" of surfaces: the sphere, torus, and projective planes (Klein bottles are two projective planes glued together). A more compact statement of that classification theorem would then be to only describe the "building blocks", then describe how one can decompose any surface into its building blocks.

While the connected sum is a helpful topological construct, it's quite hard to make it compatible with notions of lengths and angles and tangent spaces and Riemannian metrics and smooth structures and the like. Hence, the uniformisation theorem strays away from this notion of "building blocks"; it's hard (but not impossible (this is called foreshadowing hehe)) to do this deconstruction when geometric structure involved. Rather, a more natural thing to do is go the other way: we'll make our space bigger, but simpler.

Definition (Universal Covers)**:** Let Σ be a topological space. Its universal cover is the set of homotopy classes of paths in Σ, endowed with some messed up topologies that I don't want to get into for the sake of brevity.

Every topological space has a universal cover, and the thing that's great about them is that they're very compatible with geometric structure. Universal covers have the very important properties of being unique (up to homeomorphism) as well as being simply connected. Arguably more important is the fact that they're compatible with metrics — more on that later.

An example of a universal cover would be the real line $\mathbb{R}$ is the universal cover of the unit circle S^1. Physically, one may think of a universal cover as "unwinding" a topological space around its holes; in the case of S^1, this coils around an infinite cylinder and can be "uncurled" into a straight line.

Remark: This is where either figures (which I'm too lazy to embed) or physical demonstrations become very helpful. The "uncoiling" of S^1 into the real line can be visualised by extending a slinky; however, I will slap anyone that tries to uncoil my slinky. It was expensive! A less trivial example would be covering a torus T^2 by the plane $\mathbb{R}^2$, and I believe this is best visualised by shrimp (or another malleable torus-shaped object): the coiled shrimp looks like a torus, and we can "uncoil" the shrimp into an infinitely long shrimplinder. Then, we can turn the shrimplinder into a shrimplane by "rolling it out"; physically, this would probably be best demonstrated by making a butterfly cut, adding salt, pepper, lemon, and butter, then lightly sautéeing until tender.

The construction of universal covers is a great way to identify manifolds with *simply connected* manifolds; in the case of 1 complex dimension (or 2 real dimensions), this tells us that every manifold's universal cover is conformal to a sphere, a plane, or the hyperbolic disc. Furthermore, these three manifolds' geometries give them a constant curvature (left as an exercise for the avid reader)!

Remark: For the avid reader, the universal cover is a special example of something called a *covering space* — a space Σ is covered by M if there exists a continuous surjection $\pi : M \to \Sigma$ such that the preimage of a small enough neighbourhood around any point in Σ is a disjoint collection of neighbourhoods in M.

While we can obtain a simply connected surface by looking at the homotopy classes and playing with the topologies a little bit, one may wonder how to go from the universal cover back to the base space. As a physical example, S^1 (or the slinky) was "uncoiled" into a very long spiral, and to go back, all one has to do is collapse the slinky.

More generally, there's a correspondence between single points on a circle and vertical "stacks" of points on the spiral. Mathematically, one can envision this using trig facts (these haunt my dreams at night): we can *cover* the circle with the real line by using

$$\begin{aligned} \pi \ : \ \mathbb{R} &\longrightarrow S^1 \\ \theta &\longmapsto (\cos\theta, \sin\theta) \end{aligned}$$

This is a surjection, though it's not injective: any two values of θ a multiple of 2π apart will get sent to the same spot on the unit circle. That is, the set $\{\theta + 2\pi n : n \in \mathbb{Z}\}$ is collapsed into a single point on S^1. Assuming the reader has a mild background in group theory, we can envision this as a coset of $\mathbb{R}$ and write the following diffeomorphism:

$$\begin{aligned} \pi \ : \ \mathbb{R}/2\pi\mathbb{Z} &\longrightarrow S^1 \\ \theta + 2\pi\mathbb{Z} &\longmapsto (\cos\theta, \sin\theta) \end{aligned}$$

This example is illustrative of the correspondence between the simply connected universal covers of manifolds and the manifolds themselves. Rephrasing the uniformisation theorem in this light, we can say that every Riemann surface is conformal to a *quotient* of the sphere, the plane, or the hyperbolic disc. There's a special name for this:

Definition (Model Geometries)**:** Let X be a simply connected Riemannian manifold, and let $\mathrm{Isom}(X)$ be its group of isometries (metric-preserving automorphisms). Let M be a Riemannian manifold. M is said to be *modelled by* X if M is isometrically diffeomorphic to X/Γ for some subgroup Γ of $\mathrm{Isom}(X)$.

Rephrasing the uniformisation theorem again, it states that all Riemann surfaces are *modelled by* the sphere under spherical geometry, the plane under Euclidean geometry, or the disc under hyperbolic geometry. The isometries of these geometries are rotations in $\mathbb{R}^3$, compositions of translations and rotations in $\mathbb{R}^2$, and certain Möbius transformations that fix the disc in $\mathbb{R}^2$.

Remark: This definition is bulky, but it formalises the notion of "collapsing" or "folding" a universal cover back into its base space. There are some additional constraints on X and Γ: the metric on X needs to be complete and homogenous, and Γ needs to act freely and discretely on X. These are important for ensuring that the resulting topology and metric in the quotient are well-behaved, but this is a detail that's not crucial to our discussion today (though still worth mentioning).

To reiterate, the uniformisation theorem is a super powerful result in the study of 2-manifolds (or complex 1-manifolds), for it reduces the possible geometries to just, well, more or less three options. Can this be generalised to higher dimensions?

What You've Been Waiting For This Whole Time

During the twentieth century, Thurston had conjectured an answer to this question:

Theorem (Thurston's Conjecture)**:** Every compact orientable closed 3-manifold can be canonically decomposed into submanifolds such that each piece is modelled by one of the eight Thurston geometries.

Remark: The astute reader will notice that this theorem is called a conjecture (or this conjecture is called a theorem?). This is because the proof is very recent: Hamilton established a special case of this conjecture in the 1980s, and Perelman began working on it during the 1990s. This culminated in a full proof in the early 2000s. This proof is so young, in fact, that some resources online predate its existence.

The astuter readest will notice that the "canonical decomposition" is uncontrollably left vague. Thinking back to the classification of surfaces, each surface was broken down into what are called "prime" topological spaces: they cannot be expressed as a nontrivial connected sum of two smaller spaces. It's a similar idea here, but the pieces here are instead either incompressible tori (huh) or Siefert fibred spaces (wha huh).

Before exploring why there's ambiguity in this math paper, we should discuss the eight Thurston geometries. The fact that there are eight of them is already surprising; it's a significant increase from the 2-dimensional case ($p \ll 0.05$).

- Euclidean 3-space: $\mathbb{R}^3$. The isometries are given by $\mathbb{R}^3 \times O_3(\mathbb{R})$, where $O_3(\mathbb{R})$ represents the group of orthogonal 3×3 matrices with real entries. The $\mathbb{R}^3$ component of the isometries represent translations, and the $O_3(\mathbb{R})$ part represents rotations. An example of a compact 3-manifold modelled by Euclidean 3-space is the 3-torus: one may construct this by identifying opposing faces of a cube with each other.
- The 3-sphere S^3. Its isometries are given by $O_4(\mathbb{R})$, the group of orthogonal 4×4 matrices with real entries. This can be easily seen by embedding the 3-sphere in $\mathbb{R}^4$ and using the power of anime to observe that any rotation of $\mathbb{R}^4$ fixes the sphere. I'd include a figure, but the four dimensions are too big to fit in the BruinLearn submission.
- Hyperbolic 3-space, H^3. The term "hyperbolic space" doesn't actually refer to a specific Riemannian manifold; rather, it describes any manifold with a constant curvature of -1. This is highly applicable in physics: it is the space $\mathbb{R}^4$ endowed with the metric $\mathrm{d}s^2 = \mathrm{d}x^2 + \mathrm{d}y^2 + \mathrm{d}z^2 - \mathrm{d}t^2$ restricted to a certain 4-dimensional hyperboloid.
- The product space $H^2 \times \mathbb{R}$, where H^2 is the hyperbolic plane. This is a poorly understood Thurston geometry, and though there are many examples of spaces modelled by $H^2 \times \mathbb{R}$, it remains an area of active research.
- The product space $S^2 \times \mathbb{R}$, where S^2 is the 2-sphere. In contrast to the previous example, this is well-understood for the reason that there are 2 or 3 spaces modelled by this geometry (up to isometric diffeomorphism).
- The universal cover of $\mathrm{SL}_2(\mathbb{R})$, the set of all 2×2 matrices with real entries and determinant equal to 1. The reader has likely noticed the lack of discussion surrounding the isometries of these past few geometries, as well as the increasing WTF-ness of them.
- The "nil space",

  $$\mathrm{Nil} = \left\{ \begin{pmatrix} 1 & x & y \\ 0 & 1 & z \\ 0 & 0 & 1 \end{pmatrix} : x, y, z \in \mathbb{R} \right\}$$

  It's called the nil space because as a group, it is nilpotent.
- The "sol space",

  $$\mathrm{Sol} = \left\{ \begin{pmatrix} e^z & 0 & x \\ 0 & e^{-z} & y \\ 0 & 0 & 1 \end{pmatrix} : x, y, z \in \mathbb{R} \right\}$$

  Similar to the nil space, this is called the sol space because it is a solvable group.

Remark: The sane reader will question the inclusion of matrices in Thurston's geometries. These groups of matrices actually carry much more structure than their algebraic group structure. They are called Lie groups, which, beyond having group structure, admit the structure of a smooth manifold. In the nil space, one can image projecting the x, y, and z entries onto the coordinate axes. Their metrics and isometries are omitted for the reader's and author's sake (but mostly for brevity).

The proof employs a technique called Ricci flow, which involves slowly changing the metric according to the differential equation

$$\frac{\partial}{\partial t} g_{ij}(t) = -2 \sum_{p,q} R_{ipjq} g^{pq}$$

Here, R_{ipjq} is the covariant Riemannian curvature tensor. In essence, this differential equation attempts to homogenise the Riemannian metric by either "smoothing out" regions with bumpy and irregular geometries or by "pinching off" areas with high curvature.

This "pinching off" posed a problem to Hamilton: it would cause singularities in the metric in finite time, thus breaking the smooth structure of the manifold. Hence, Hamilton's work could only apply to special cases of 3-manifolds. However, when Perelman began working on the problem, he showed that one could work around these singluarities by performing surgery on them. After applying sufficient topical analgesics to the manifold, Perelman just "cut around" the singularities. This cutting action on the singularities produces the "canonical submanifolds" mentioned in the conjecture, and it does not involve any arbitrary choices.

This concludes the discussion of Thurston's conjecture. An important corollary of the conjecture is Poincaré's conjecture — that every 3-manifold is homotopy equivalent to the 3-sphere if and only if they are homeomorphic. Some closing remarks:

Remark: One of the difficulties of extending the uniformisation theorem to 3 dimensions is that it relied heavily on techniques in complex and harmonic analysis, which is not available here because 3-manifolds don't usually admit complex structures. However, it has been shown that analysing Ricci flows can be used to give an independent proof of the uniformisation theorem without leaning on complex analysis. Applications of Ricci flow to higher dimensional manifolds appears to require heavier assumptions on the manifold's geometry.

Remark: [4] provides an excellent set of computer renders of several of the Thurston geometries. In addition, they provide a creepy monkey named Suzanne and disembodied hands living in these 3-manifolds.

References

[1] Michael Boileau. *Geometrization of 3-Manifolds with Symmetries*. 2004.

[2] Ángel Prieto de la Cruz. "The Geometrisation Conjecture of 3-Manifolds". University of Barcelona, 2019.

[3] Jeremy Gray. "On the History of the Riemann Mapping Theorem". In: *Rendiconti del Circolo Matematico di Palermo* (1994), pp. 47–94.

[4] Tiago Novello et al. *How to see the eight Thurston geometries*. Sept. 2021.