Seifert surfaces and genera of knots

Michael Landry

The subject of *topology* encompasses a huge swath of modern mathematics, and yet many undergraduates have never heard of it. One reason for this is that there is a relatively high barrier to entry in terms of definitions and background, so it tends to be taught in graduate and advanced undergraduate courses.

One of the most immediately accessible parts of topology is a field called *knot theory*, which is the study of circles tangled up in 3-dimensional space. Most of us have some experience tying knots (shoelaces, for example), and thus have some intuition that can help us when thinking about knot theory. I hope this short article interest piques your interest in topology!

If you want more details, I encourage you to look at Colin Adam's excellent text [1], which was the primary resource I used in preparing this.

1 Knots

Whatever the definition of a knot is, it should rigorously model the following situation: you have a string which is tangled up in some way, and you attach the 2 loose ends to each other. This could be by gluing them, taping them, or by tying a knot, although the last method is linguistically problematic for obvious reasons. The result is a circular string which you can play with, and which we will call a *knot*. Maybe if you fiddle around with it for long enough, you can untangle it to a circle which lies flat on the table—no cutting allowed though—in this case, the knot you created is what we call the *unknot*. Otherwise, your string is knotted in some intrinsic sense. This knotting phenomenon is what knot theory explores.

If you do not like thinking of your mathematical objects being made out of string and tape, fine. Being more explicit, but still a little vague, think of a knot as an equivalence classes of embeddings of the circle in three-dimensional space, where 2 embeddings are equivalent if one can be continuously deformed to the other such that each intermediate map is injective.

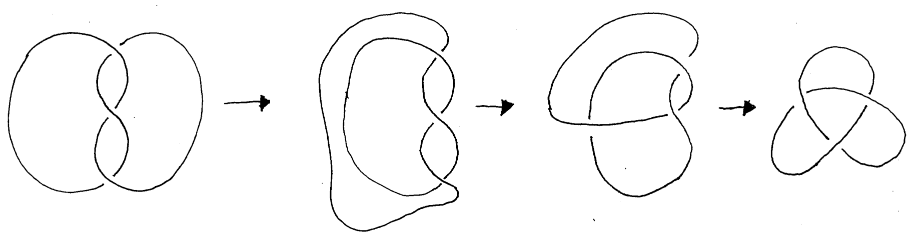

Figure 1: These are all the same knot, as the deformations above show.

2 Connected sums of knots

You're used to thinking about ways of combining two numbers to get another number (addition, multiplication, exponentiation,...). We now define a way to combine two knots and get another knot! It's called *connected sum*. The connected sum of 2 knots, K_1 and K_2, denoted $K_1\#K_2$, is obtained in the following way: draw the two knots next to each other, and remove a small segment from an outer arc of each knot. Now glue the 2 remaining parts together as in Figure 2. It is not clear that this operation is well-defined, and in fact

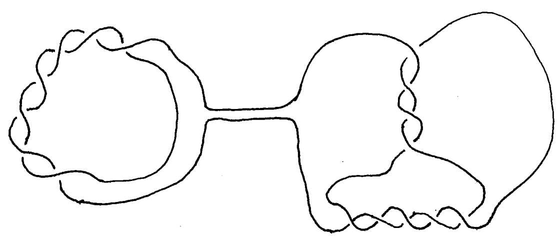

Figure 2: A connected sum of 2 knots.

it is not unless we define it on *oriented* knots, where the gluing is determined by requiring that it preserve orientation. You might be a little nervous that we had to make a choice of which small segment to remove. This is not an issue as long as the segments we choose do not traverse any crossings in our drawings of the knots. Can you prove this?

What can we say about #? Here are some observations:

(a) # is commutative.

(b) # is associative.

(c) if K is any knot and U is the unknot, then $K\#U = K$.

One next natural question to ask is whether # has inverses. That is, does there exist a knot K^{-1} such that $K\#K^{-1} = U$? Could a knot like the one shown in Figure 2 be the unknot? It turns out that the answer is no, and the rest of this article will be devoted to proving this result.

3 The classification of compact, connected, orientable surfaces

We need to make a few definitions.

A *surface* is a two-dimensional manifold. If you are not comfortable with that description, think of a surface as a space that looks like a two-dimensional plane if you zoom in really close to any point. A *surface with boundary* is a space that looks like either a two-dimensional plane or the upper half of a two-dimensional plane if you zoom in on any point. The *boundary* of such a surface is composed of the "upper half plane" points. A surface is *connected* if any 2 points in the surface can be connected by a path which is also in the surface.

A surface is *compact* if it doesn't "go off to infinity" anywhere. One way of phrasing this rigorously is to say S is compact if every continuous function $S \to \mathbb{R}$ has a maximum value somewhere in S. For example, the unit disk

$$D = \{(x, y) \in \mathbb{R}^2 \mid (x^2 + y^2 \leq 1)\}$$

is compact, but the punctured disk

$$D' = \{(x, y) \in \mathbb{R}^2 \mid 0 < (x^2 + y^2 \leq 1)\}$$

is not compact because the function $D' \to \mathbb{R}$ defined by $(x, y) \mapsto \frac{1}{x^2+y^2}$ is unbounded. Alternatively, you can think of compact surfaces as the ones that can be constructed by gluing together finitely many triangles.

Define an *orientable surface* to be a surface (possibly with boundary) which is 2-sided, in the sense that if a tiny person is standing on the surface at a point P and takes a walk which stays away from the boundary, it is impossible for them to return to P with their head pointing in the opposite direction.

Remark 1. Orientability is usually defined in a different, more complicated way. But for surfaces embedded in an orientable 3-manifold, orientability and 2-sidedness are equivalent conditions. However, there are examples of three-dimensional manifolds which are nonorientable and admit two-sided embeddings of nonorientable surfaces (see [5], for example). Since we are picturing our surfaces as sitting in $\mathbb{R}^3$, which is orientable (note that I have not defined orientability for three-dimensional manifolds!), we are good to go.

All those definitions pay off in the following classification theorem.

Theorem 2. *Every orientable, connected, compact surface with boundary is topologically equivalent (or homeomorphic, if you know that word) to one of the surfaces shown in Figure 3.*

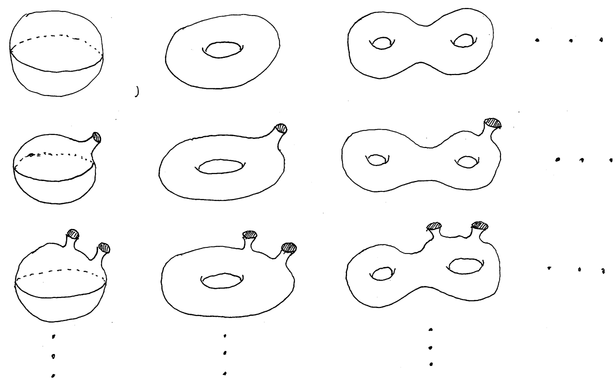

Figure 3: Classification of orientable, connected, compact surfaces with boundary.

That is, any such surface is uniquely determined by its number of boundary components and its number of "holes," which we call the *genus* of the surface (the plural of genus is *genera*). For a proof of Theorem 2, see [4].

4 Seifert surfaces

A Seifert surface for a knot K is a compact, orientable, connected surface with boundary equal to K. By the previous section, any such surface must be topologically equivalent to one of the surfaces in the second row of Figure 3 since it has only one boundary component. Therefore any Seifert surface has a well-defined genus. The genus of a knot K, denoted $g(K)$, is the minimum genus of all Seifert surfaces for K.

Note that the definition above does not make sense unless we know that any knot possesses at least one Seifert surface. But this fact is not that hard to prove. In fact, Seifert himself gave a nice algorithm for producing a Seifert surface given a drawing of a knot, which is illustrated in Figure 4.

1. Pick an orientation for your knot.

2. Start tracing around the knot in your chosen direction. When you get to a crossing turn right or left, in whichever direction agrees with the orientation of the knot. When you get back to where you started, start somewhere else in the drawing. Do this until you've traced over all the arcs in the drawing.

3. In the previous step you created a collection of circles, called Seifert circles. Attach a disk to each Seifert circle, so that each Seifert circle now bounds a disk. In the case of concentric circles, vary the height of the disks so that the innermost Seifert circles are the highest.

4. At the site of each of the old crossings, attach the 2 neighboring disks by gluing in a rectangular band with a half twist, the direction of the twist being determined by the original crossing.

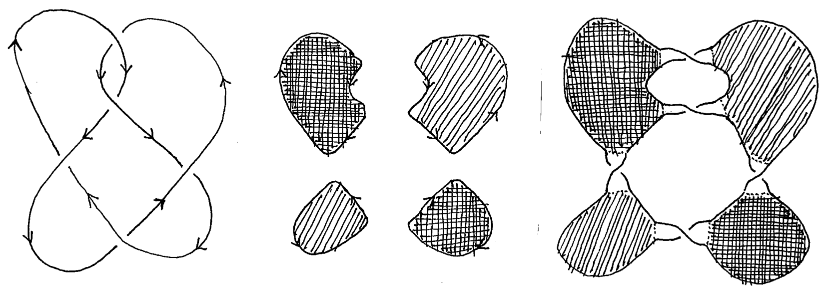

Figure 4

The resulting surface clearly has boundary the original knot. What is not immediately clear is that the resulting surface is orientable. But this is easy enough to see, and I leave it as an exercise for you. Here's a hint: if the boundary of a disk has counterclockwise orientation, paint its top red and its bottom blue. Do the opposite for disks whose boundaries have clockwise orientation. Why does the way we glued in the bands allows us to extend the red and blue paint over the entire surface such that the two colors never meet?

Example 3. Consider the Seifert surface for the trefoil knot shown Figure 5. By cutting and pasting as shown in the figure, we can see that this surface is homeomorphic to torus minus a disk. You can also see this using Euler characteristic, but cutting and pasting feels more fun to me.

This shows that the genus of the trefoil knot is ≤ 1. If the trefoil knot had genus 0, then it would bound a disk, but this is impossible (why?). Therefore we have determined that the genus of the trefoil knot is 1.

Example 4. Let's kick it up a notch. Consider the Seifert surface shown in Figure 6. Well, according to our definition this is not strictly speaking a Seifert surface because its boundary is not a knot, it is a *link*. A link is just a collection of circles embedded in 3 space, up to the same equivalence relation that we used to define knots. We broaden our definition of Seifert surface to include those whose boundaries are links. This particular link is called the Borromean rings, and has the neat property that no two of the three circles are linked to each other, even though the link as a whole cannot be undone.

We can obtain the surface in Figure 6 by gluing together 3 hexagons in the way prescribed by Figure 7, which shows that the surface is topologically equivalent to a torus minus 3 disks. Therefore the genus of the Borromean rings is ≤ 1.

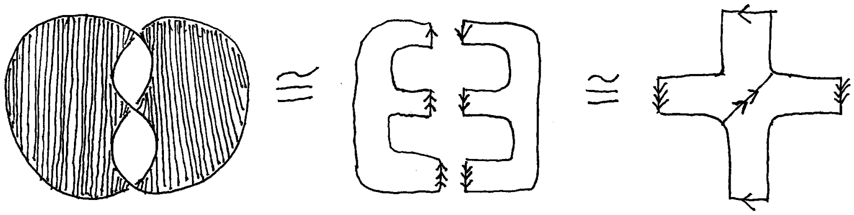

Figure 5: Left: A Seifert surface for the trefoil knot. Center: by cutting the lefthand surface open along a vertical plane intersecting the knot's crossings, we see that the surface is obtained by gluing together these two "3-shaped" surfaces in the way shown. Right: can you see why this surface is the same as a torus minus a disk?

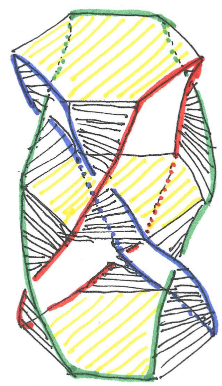

Figure 6: A Seifert surface for the Borromean rings.

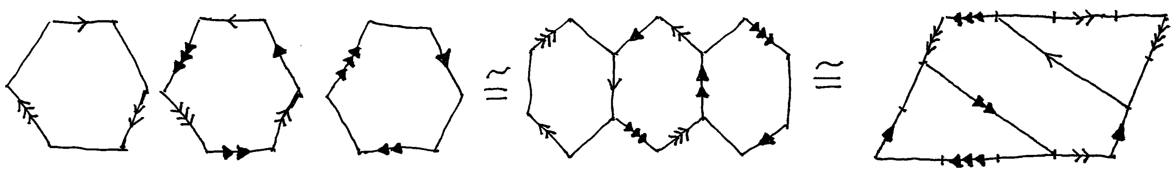

Figure 7

Note that "topologically equivalent" does not mean that we can continuously and injectively deform the surface in Figure 6 (à la our knot deformations) to look like the $(4, 2)$ entry of the table in Figure 3. Indeed, such a deformation would unlink the Borromean rings, which is impossible. We will not prove that, although it is not hard—the 3 component unlink has a property called *tricolorability* which the Borromean rings do not have.

One last thing about this example—we have shown that the genus of the Borromean rings is at most one, but we have not shown that the Borromean rings do not bound a thrice-punctured sphere, so we cannot conclude that their genus is actually 1. The genus is in fact 1, but the only proof I am aware of uses the Thurston norm on homology, which falls beyond the scope of this talk.

5 Main Theorem

It is time to finally get around to resolving that question about "inverses" of knots from Section 2. Conveniently, this follows immediately from the following theorem describing the relationship between genus and connected sum.

Theorem 5. *Let K_1 and K_2 be knots. Then $g(K_1 \# K_2) = g(K_1) + g(K_2)$.*

Proof. First, observe that $g(K_1 \# K_2) \leq g(K_1) + g(K_2)$ because you can attach minimal genus Seifert surfaces for K_1 and K_2 by a rectangular band to get a Seifert surface for $K_1 \# K_2$ whose genus is the sum of the original two.

Now we prove that the reverse inequality holds. Take a Seifert surface Σ for $K_1 \# K_2$ with minimal genus. By assumption, there exists a sphere S, punctured twice by $K_1 \# K_2$, which separates K_1 from K_2 (see Figure 8).

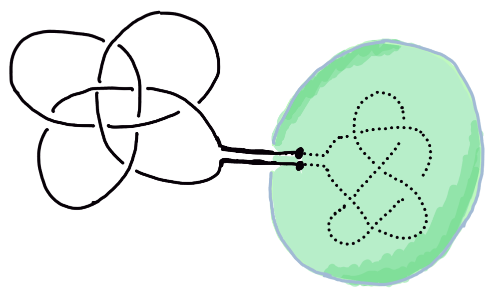

Figure 8: Given a connected sum $K_1 \# K_2$, there is always a sphere separating K_1 from K_2 and intersecting $K_1 \# K_2$ exactly twice. How might a Seifert surface for $K_1 \# K_2$ intersect this sphere?

By perturbing the sphere ever so slightly, we can preserve this separation property and guarantee that Σ intersects S in arcs, all but one of which are circles (the last one is a path P between the 2 punctures). The fact that we can do this follows from the theory of *transversality*, which you can read about in [3] if you are interested.

If there are no circles on the surface of S, we are done—by cutting along P we do not change the genus and we obtain Seifert surfaces for K_1 and K_2, so $g(K_1 \# K_2) \geq g(K_1) + g(K_2)$.

Otherwise there is an innermost circle on S, and by doubling S inside the circle we create a new Seifert surface which intersects S in one less circle.

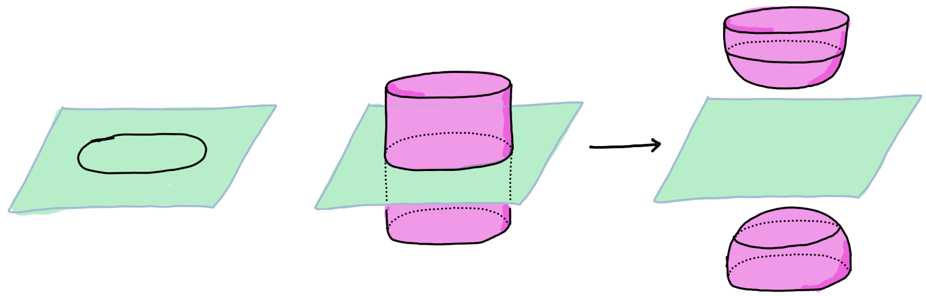

Figure 9: Our Seifert surface Σ intersects the sphere S (green) in one arc and some collection of circles. Look at an innermost circle (left). Near to the innermost circle, Σ (pink) is as shown the center picture. We can cut Σ along the circle and cap it off as shown on the right to get a new surface that intersects S in one less circle.

If the resulting surface is disconnected, discard the component which does not connect to the knot. This surgery does not increase the genus, because it increases Euler characteristic by 2 without changing the number of boundary components. As an aside, this even shows that we always create a sphere component by this surgery, since $g(\Sigma)$ is minimal.

By repeating this type of surgery, we eventually reduce the number of intersecting circles to 0, so we are finished. □

Corollary 6. *If either K_1 or K_2 is nontrivial, then $K_1 \# K_2$ is not the unknot.*

Proof. As we remarked earlier, the only knot whose genus is 0 is the unknot. □

References

[1] Adams, Colin. *The Knot Book.* American Mathematical Society, 2001.

[2] Candel, Alberto and Conlon, Lawrence. *Foliations II.* American Mathematical Society, 2003.

[3] J. Lee. *Introduction to Smooth Manifolds.* Second edition. Springer, 2013.

[4] Massey, William. *Algebraic Topology: An Introduction.* Harcourt, Brace & World, Inc. 1967.

[5] Weeks, Jeffrey. *The Shape of Space.* Chapman & Hall, 2001.