What is a Surface?

For our purposes, a surface refers to a geometric object which obeys the following conditions.

- If you pick a point on a surface and look at all of the points close to it, then they should form something which looks like a disk. (More precisely, every point on the surface has a neighbourhood which is homeomorphic to a disk.)
- A surface must be finite in the sense that you can put it inside your house. (More precisely, a surface should be bounded.)
- A surface must be in one piece, so that you can't have some of it inside your house, some of it outside your house, and be able to close the door. (More precisely, a surface should be connected.)
- A surface should not have boundaries. (More precisely, a surface should not have edges which you can walk off.)

We've already seen some examples of surfaces, such as the sphere and the torus. You should convince yourself that they obey all of the properties that a surface should obey. Note that the plane is not a surface, because it obeys all of the required properties except for the finiteness condition. A disk is not a surface because it obeys all of the required properties except for the boundary condition. And a sphere sitting next to a torus is not a surface because it obeys all of the required properties except for the connectedness condition.

Our goal is to classify all surfaces — in other words, hunt them all down. This is going to be a two step process — first, we have to be able to list all of the surfaces possible, and second, we have to make sure that no two surfaces in our list are homeomorphic to each other. As a start, we have the following examples of surfaces — the sphere which has no holes, the torus which has one hole, a surface which has two holes, and so on. The technical name for the number of holes on a surface is the *genus*.

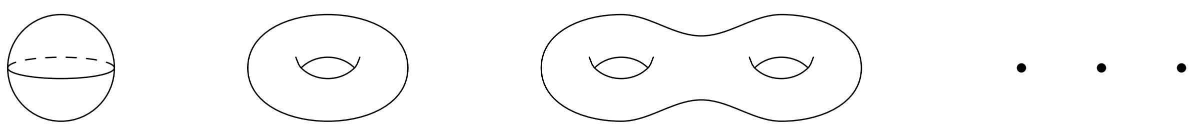

genus = 0 genus = 1 genus = 2

Pac-Man World

Let's change topic a little and look at Pac-Man world, the world in which Pac-Man lives.[1] Pac-Man lives in a world which appears to us as a rectangle on the computer screen.

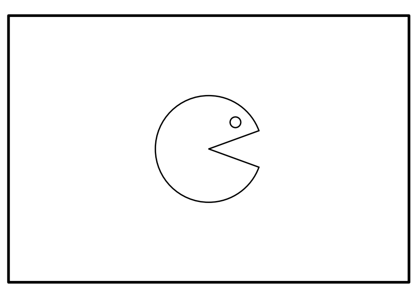

[1]For those of you who are too young to know, Pac-Man is a classic arcade game in which you control Pac-Man through a maze so that he eats lots of pac-dots. It's universally considered one of the classics of the medium, virtually synonymous with video games, and an icon of 1980s popular culture. To be gender neutral, I should also mention Ms Pac-Man, a spin-off arcade game involving a female version of Pac-Man who looks exactly the same as Pac-Man, but with a bow on her head.

But his world is much more interesting than just any old rectangle, because when he moves off the right edge of the screen, he reemerges on the left edge of the screen.

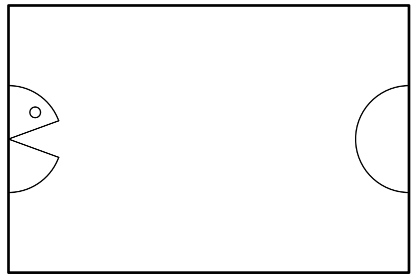

And if you try to move Pac-Man off the bottom edge of the screen, he'll reemerge on the top edge of the screen.

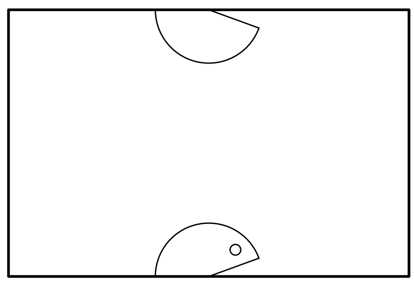

In fact, even if Pac-Man is very naughty and you decide to put him in the corner, you can't hide him because he'll now appear at every corner of the screen.

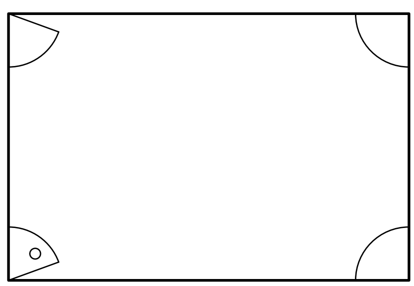

What this tells us is that Pac-Man world is, in fact, an example of a surface. To see this, all you need to do is check that the four properties required of a surface are true. It should be clear that no matter where Pac-Man stands, there is a little disk around him. This is particularly easy to see when he is in the middle of the screen. But even if he is on the right edge of the screen, there is a patch around him which will appear as a semicircle on the right edge and a semicircle on the left edge — these together make up a disk. Of course, this very same argument applies no matter which edge Pac-Man decides to stand on. The most tricky case is when Pac-Man is in the corner of the screen, but we've already seen that he will be surrounded by a disk which will appear to us as four quarter circles.

Pac-Man world also happens to be finite, because it appears on our finite computer screen. It has no boundary because Pac-Man can never run off the edge of the world. And it is certainly connected, because Pac-Man can visit every part of it just by running around eating pac-dots.

So the obvious question to ask is what surface does Pac-Man live on? Is it one we've seen before, or is it something brand new? To see what the answer is, we use the following strategy. Note that a point on the bottom edge of the screen is the same as a point on the top edge of the screen, even though it seems to appear in two different places. So the obvious thing to do is simply glue these two points together — in fact, we'll end up gluing the whole bottom edge to the whole top edge. And similarly, we'll end up gluing the whole left edge to the whole right edge. So we can think of Pac-Man world as a rectangular piece of fabric which comes to us with gluing instructions, indicated by arrows. Edges with matching arrows get glued together with the arrow telling you how the edges should line up with each other.

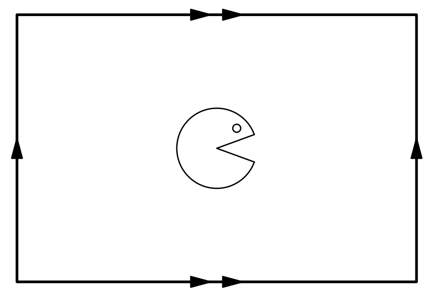

The diagrams below show what happens first when you glue together the top and bottom edges, and the second when you glue together the remaining two edges. The result, as you can plainly see, is simply the torus. Pac-Man lives on a torus.[2]

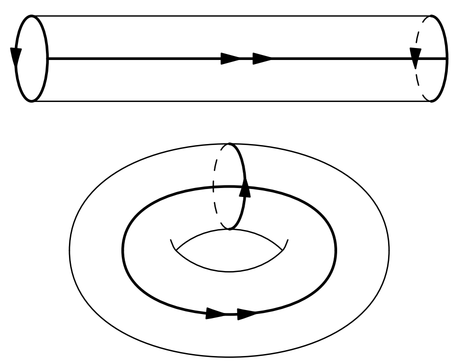

Connected Sums

In many branches of mathematics, it's useful to take a reductionist approach. For example, when we studied Euclidean geometry, we reduced all theorems to just a handful of axioms. In group theory, you can slice and dice groups into pieces called simple groups. When talking about surfaces, there is a very useful way to construct a new surface out of two old ones which will help us apply the reductionist approach. To find the connected sum of two surfaces, you simply cut out a little hole from each and glue the two surfaces together along these holes. Probably the best thing to do is look at the following sequence of diagrams to see exactly how this works. We denote the connected sum of two surfaces A and B as $A \# B$.

[2] This is all a bit of a lie, because Pac-Man doesn't exactly work this way, but I hope you understand the point. Actually, it's possible to take many games on rectangles and turn them into games on tori by gluing the edges of the board together. For example, toroidal chess is a variant of chess which uses such a board. Of course, the pieces have to start in different locations and the rules have to be altered very slightly.

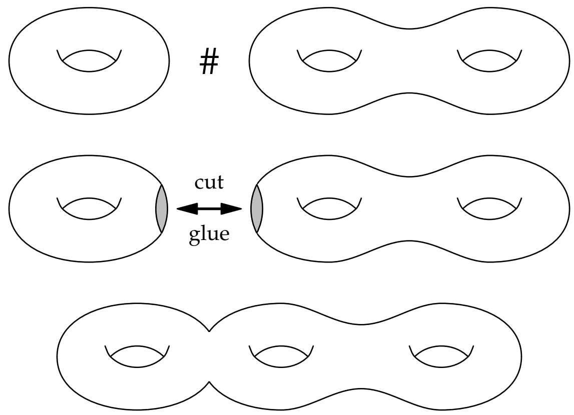

Of course, if you have a complicated surface, then you should be able to cut it up into smaller pieces whose connected sum is the original surface. The obvious question we'd like to answer is the following — is every surface the connected sum of just a few surfaces?

One interesting observation is the fact that if you take the connected sum of a surface and a sphere, you always get the original surface back again. So, in some sense, the sphere is to connected summation as zero is to addition or one is to multiplication or the identity is to group composition, and so on.

We're now going to ask a very subtle question — are there different ways to produce the connected sum of two surfaces? In other words, if I have two surfaces A and B and you have two surfaces A and B, and we both try to form the connected sum $A \# B$, then will the results always be homeomorphic? The answer is that they definitely will be homeomorphic, although there are two issues at play here.

- First, we may choose different locations to cut our holes on A and B. It's not too hard to see that this makes no difference, because you can slide the holes around the surfaces and this corresponds to bending, stretching, warping, morphing or deforming the resulting shape. In other words, sliding holes around before forming the connected sum will give surfaces which are homeomorphic to each other.
- Second, and a lot harder to see, is the fact that there are always two ways to glue two holes together. For example, suppose that you have a very thin wall with a small hole in it and a tube like what you might find on a vacuum cleaner. If someone tells you to plug the vacuum cleaner into the hole, then you could do it from two different sides of the wall. We won't say too much about why these give the same result when you form the connected sum. However, it follows from the fact that you can turn a shape inside out merely by deforming it, especially if you do it in four or more dimensions.

The Möbius Strip

Here is something that you can do in the comfort of your own home, and all you need is some paper, some scissors, some glue and three dimensions. Cut a long strip of paper and glue the ends together — but before you glue the ends together, give the strip a half twist. The resulting object is called a Möbius strip. Note that it isn't a surface by our definition, since it has a boundary. In fact, one of the interesting things about a Möbius strip is that it has only one edge, unlike its relative the cylinder, which is formed by gluing the ends of a long strip of paper together without any twists at all. Furthermore, a Möbius strip only has one side, unlike the cylinder. This means that if you start painting one side of the Möbius strip and come back to where you started painting, then the whole Möbius strip will be painted over. Don't just take my word for it — try this at home for yourself.

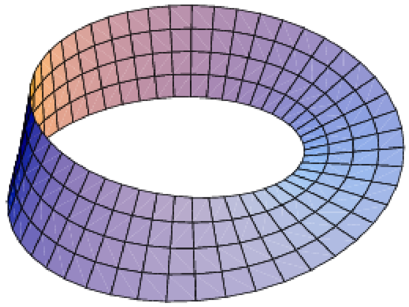

The following woodcut by M. C. Escher entitled *Möbius Strip II* shows some ants walking around a Möbius strip. The Möbius strip also appears on the universal recycling symbol. In fact, the Möbius strip appears in various places, such as the shape of some conveyor belts, since this allows the belt to wear evenly on both sides, rather than on just one. They also appear in the design of computer printer and typewriter ribbons.

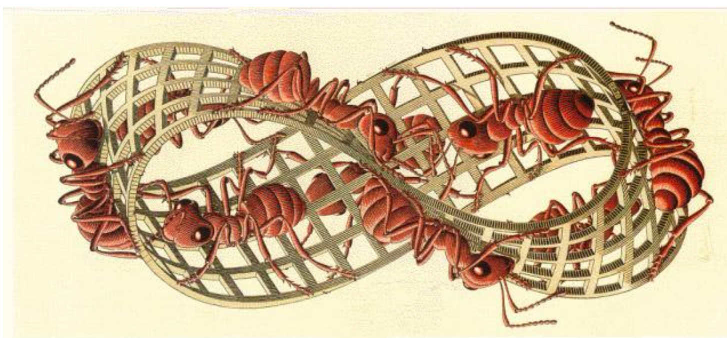

Unfortunately, as interesting as Möbius strips are, they aren't surfaces — and we're interested in surfaces. Fortunately, we can easily make a surface out of a Möbius strip simply by observing that it has precisely one edge. Now take a patch in the shape of a disk which also has only one edge. And then glue the two edges together, so that there are no spare edges remaining. Try as you might, this just isn't possible to do in three dimensions. However, that's just fine, since we don't always have to draw something which looks like a surface for it to be a surface. For example, Pac-Man world didn't look a surface and yet we discovered that it was one, the torus in fact.

One way to think of this surface obtained from patching a Möbius strip with a disk is as follows. Suppose that all along the edge of the Möbius strip are numbers — say from 1 to 100 — and that all along the edge of the disk are numbers — also from 1 to 100. Suppose that you are an ant walking around the Möbius strip and, as soon as you hit the edge of the strip — say at the number 73 — you suddenly get teleported over to the number 73 on the disk and so you just continue walking along the disk. If you try and leave the disk, you get teleported back to the edge of the Möbius strip — precisely where you get teleported back corresponds to where you left the disk. This is now obviously a surface because every point on the edge of the Möbius strip or the patch has a little disk around it made up of one semicircle on the strip and one semicircle on the patch.

The resulting surface is called the *projective plane* and it's definitely not homeomorphic to any of the surfaces we've previously encountered. This is because all of the surfaces we've seen before have two sides whereas this new beast only has one. We say that a surface with two sides is *orientable* and that a surface with only one side is *non-orientable*. It seems that the world of surfaces isn't quite so simple after all.

Orientability

Remember that we're trying to classify surfaces, and that classification is always a two step process. First, we need to be able to list out all of the possible surfaces that exist and second, we need to make sure that we don't have two things in our list which are homeomorphic. We've already mentioned that to show that two surfaces are homeomorphic is often an easy task — you simply find the homeomorphism between them. On the other hand, to show that two surfaces are not homeomorphic is often a difficult task — you need tricks. We'll see that when it comes to surfaces, there are really only two tricks that you need.

The first trick is to consider the orientability of a surface — in other words, whether it's one-sided or two-sided. All of the simple examples of surfaces that we could think of, such as the sphere or the torus, seemed to have two sides so we called them *orientable*. On the other hand, we now have the projective plane, a surface which results from gluing the one edge of a Möbius strip to the one edge of a disk. This surface has only one side and surfaces like this are called *non-orientable*. As another example of a non-orientable surface, you could also take two copies of the Möbius strip and glue them together along their edges. The resulting surface is known as the *Klein bottle*.

Another way to think of orientability is to consider whether or not there is a notion of clockwise and counterclockwise. If you draw a counterclockwise arrow on an orientable surface and move it around, it will always appear to be a counterclockwise arrow. On the other hand, consider a counterclockwise arrow on the Möbius strip — if you take it for a walk all the way around the strip, then it will come back looking like a clockwise arrow. From this definition of orientability, it's easy to prove the following fact.

Proposition. *If S_1 and S_2 are orientable surfaces, then the connected sum $S_1 \# S_2$ is orientable. If S_1 is non-orientable and S_2 is any surface, then the connected sum $S_1 \# S_2$ is non-orientable.*

Example.

- The sphere is orientable, as is any connected sum of tori.
- The projective plane is non-orientable, as is any connected sum of projective planes.

Euler Characteristic

Euler's formula states that in any polyhedron, the number of vertices, edges and faces satisfies the equation $V - E + F = 2$. Actually, we shouldn't treat this as a statement about polyhedra, but a statement about maps on the sphere. We define a *map* to be a way to divide a surface into vertices, edges and faces, where the edges no longer have to be straight lines, just any old curves on the surface. What Euler's formula is really telling us is that for every single map you can draw on the sphere — because all polyhedra are homeomorphic to the sphere — the number of vertices, edges and faces satisfies the equation $V - E + F = 2$. We define the number $V - E + F$ to be the *Euler characteristic* of a surface, so what Euler's formula tells us is that the Euler characteristic of the sphere is 2.

Using the same strategy as for the sphere, we can calculate the Euler characteristic of any surface — but this begs the following question. How do we know that the Euler characteristic is an intrinsic property of the surface and doesn't depend on the map that we draw? In other words, if I draw a map on a surface and you draw a map on the same surface, will they always have the same value of $V - E + F$? This question is answered by the following result.

Proposition. *Two maps on the same surface will always have the same value of $V - E + F$.*

Proof. We proved this fact for planar graphs by deleting the number of vertices, edges and faces and showing that $V - E + F$ remained the same after each deletion. Unfortunately, this proof wasn't very precise and you would have difficulty repeating it for a general surface. This is because we reduced planar graphs to a single vertex and a single face on the sphere, but there is no analogous "simplest map" on other surfaces.

So instead, we're going to see that adding vertices, edges and faces doesn't change the value of $V - E + F$. For example, if we simply take an edge and break it into two by adding a vertex in the middle, then the number of vertices has increased by one, the number of edges has increased by one, while the number of faces is the same. So $V - E + F$ is the same. Similarly if we divide a face into two by an edge, or if we decide to add a new edge which has one endpoint a vertex of degree one. There are various cases to check, but the crux of the argument is that you can add vertices, edges and faces to your map without changing the value of $V - E + F$.

Now I simply take my map and you take your map and we know that the union of these two maps — that is, the map obtained by overlapping them on the surface — can be obtained from mine by adding vertices, edges and faces. By our previous argument, this means that $V - E + F$ for my map is the same as $V - E + F$ for the union of our maps. And by the same reasoning, $V - E + F$ for your map is the same as $V - E + F$ for the union of our maps. In short, our maps have the same value of $V - E + F$. □

This result means that the Euler characteristic of a surface S — usually denoted by $\chi(S)$ — is an intrinsic number associated to S and doesn't depend on which map you decide to draw on S. For example, Euler's formula asserts that $\chi(\text{sphere}) = 2$ and you can calculate that $\chi(\text{torus}) = 0$. A more difficult example is to calculate the Euler characteristic of the projective plane — it's equal to $\chi(\text{projective plane}) = 1$. We'll see some techniques later for computing this and essentially the Euler characteristic of any given surface.

Proposition. *If S_1 and S_2 are two surfaces, then the Euler characteristic of their connected sum is*

$$\chi(S_1 \# S_2) = \chi(S_1) + \chi(S_2) - 2.$$

Proof. Suppose that we have a map on S_1 which contains at least one triangular face and a map on S_2 which contains at least one triangular face. Since we are allowed to remove any hole from the surfaces in order to create the connected sum, let's decide to remove the triangular faces. After gluing the three edges from one hole to the three edges from the other, we can see what happened to the total number of vertices, edges and faces. Clearly, the number of vertices decreased by three, because the three around one hole and the three around the other hole were glued together to result in three vertices. Similarly, the number of edges decreased by three, because the three around one hole and the three around the other hole were glued together to result in three edges. And finally, the number of faces simply decreased by two, since we removed two triangular faces to make the holes and never added any back in. Therefore, the value of $V - E + F$ simply decreased by two and we have the equation $\chi(S_1 \# S_2) = \chi(S_1) + \chi(S_2) - 2$, as desired. □

This theorem allows us to calculate various Euler characteristics using the following result, which you should try to prove.

Corollary. *The Euler characteristic of the connected sum of g tori is $2 - 2g$ and the Euler characteristic of the connected sum of n projective planes is $2 - n$.*

Poincaré

Jules Henri Poincaré was a French mathematician, theoretical physicist and philosopher who lived from 1854 to 1912. He is often described as The Last Universalist, since he excelled in every branch of mathematics known in his day. Poincaré is considered one of the founders of topology, his work on the famous three-body problem led to the birth of chaos theory, and he is also renowned for introducing the modern principle of relativity, although Einstein seems to have received most of the credit.

At school, Poincaré was a pretty clever student, achieving the top marks in almost every topic he studied. Poincaré went on to study mathematics and mining engineering at university. Although he was soon offered a post as a lecturer in mathematics, he never fully abandoned his mining career for mathematics. Poincaré went on to become chief engineer of the Corps de Mines, held various chairs at the Sorbonne, became president of the French Academy of Sciences and was elected to the Académie française.

One of the most famous problems in mathematics — the Poincaré conjecture — was formulated by him. To understand just how important this problem is in mathematics, note that when Stephen Smale proved the conjecture in dimensions five and up in 1961, he was awarded a Fields Medal, the highest honour in mathematics. When Michael Freedman proved the conjecture in dimension four in 1982, he was awarded a Fields Medal as well. And when Grigori Perelman proved the conjecture in dimension three in 2002, he was also awarded a Fields Medal.[3]

Poincaré's work habits have been compared to a bee flying from flower to flower. He didn't care about being rigorous and disliked logic. He believed that logic was not a way to invent but a way to structure ideas and that logic limits ideas. A psychologist by the name of Toulouse wrote the following about Poincaré.

- He worked for short periods of time, doing mathematical research for only four hours a day.
- His normal work habit was to solve a problem completely in his head, before committing the completed solution to paper.
- His ability to visualise what he heard proved particularly useful when he attended lectures, since his eyesight was so poor.
- He was physically clumsy and artistically inept.
- He was always in a rush and disliked going back for changes or corrections.

His method of thinking is well summarised by the following quotation. "Accustomed to neglecting details and to looking only at mountain tops, he went from one peak to another with surprising rapidity, and the facts he discovered, clustering around their center, were instantly and automatically pigeonholed in his memory."

[3] Interestingly, Perelman declined the Fields Medal and seems to no longer be working in mathematics. In fact, as of 2003, he gave up his job at the Steklov Institute and is currently unemployed, living with his mother in Saint Petersburg.