Doughnuts

Achille C. Varzi
Department of Philosophy, Columbia University
New York, NY 10027

(Published in Monima Chadha and Ajay K. Raina (eds.), *Basic Objects: Case Studies in Theoretical Primitives*, Shimla: Indian Institute of Advanced Study, 2001, pp. 41–51.)

Abstract. In classical topology the only thing that matters in a doughnut is the edible stuff. Here I review some good reasons for questioning this attitude and focusing on the hole as well. By studying the topology of the hole one can learn interesting things about the morphology of the doughnut (its shape), and by studying the morphology of the hole in turn one can learn a lot about the doughnut's embedding in three-dimensional space. There is a price, though— one must be serious abou reifying voids.

Introduction

A doughnut always comes with a hole. If you think you can come up with an exception, then *that* would simply not be a doughnut. It would not be a doughnut by definition. Holeless doughnuts are like round squares or unmarried husbands—conceptual nonsense.

Does it follow, then, that when you buy a doughnut, you really buy two things—the edible stuff *plus* the little chunk of void in the middle? In a sense, it does follow. You cannot just take the doughnut and leave the hole at the grocery store. In another sense, however, one might want to resist this answer and insist that the edible stuff is all there is—the hole is nothing at all. In this sense, buying a doughnut is simply buying an object with a fixed shape, just like buying a pizza: there is some room for flexibility (size, rigidity, thickness) but the shape is pritty much fixed. The pizza is *flat*; the doughnut is *perforated*.

As it turns out, these two ways of answering the question correspond to two entirely different ways of representing three-dimensional spatial structures. Traditional wisdom, combined with a form of *horror vacui*, tends to favor the second answer. But, surprising as it might seem, the first answer has its own advantages.

Horror Vacui and Object Topology

Let us first take a closer look at the traditional view: holes are merely a *façon de parler*. On this view, when one says “There is a hole in my doughnut” one really means “My doughnut is perforated”, where ‘… is perforated’ is just an ordinary shape predicate like ‘… is flat’ or ‘… is pyramidical’. It is a perfectly innocuous shape predicate that may truly be predicated of a material object (a doughnut) without any suggestion that perforation is due to the presence of something immaterial (a hole).

If we put it this way, the view expresses a philosophical position.[1] But unlike many philosophical positions, this one can be implemented most thoroughly. In fact, the whole of topology may be regarded as a way of providing a detailed account of perforation along these lines. Here is how.

Topology is a sort of rubber geometry. It is concerned with the way the shape of an object can be transformed into another by pure elastic deformation. This means you can stretch your object and distort it, but you are not allowed to connect what was disconnected (e.g., by pasting two surfaces or two parts of the same surface) or to disconnect what was connected (by making a cut). For instance, a cube can be transformed into a round ball in this way: just imagine it is made of plasticine and gradually smooth out its edges and corners. Now, suppose your doughnut is also made of plasticine. You can deform it into various shapes. You can even transform it into a ball with a handle (Figure 1). However, you cannot get rid of the handle without cutting or pasting somewhere. And that’s what its being perforated amounts to, topologically speaking. In topology, to say that an object is perforated is to say that one cannot transform it into a spherical object (a ball) by mere elastic deformation.

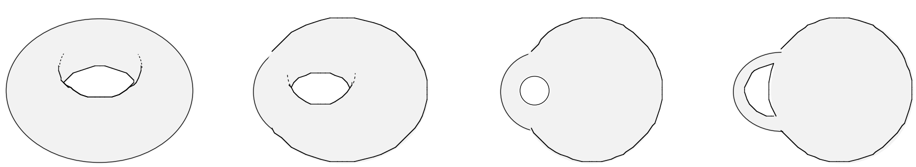

Figure 1: Deforming a doughnut into a sphere with a handle.

---

[1] See David and Stephanie Lewis’s dialogue ‘Holes’ (*Australasian Journal of Philosophy* 48, 1970, 206–12).

Now, there are various ways of making this more precise, depending on how we define a spherical object. For instance, a customary definition for topology is this: an object is spherical (or can be deformed into a sphere) if you cannot draw a circle or a closed curve on its surface without dividing the entire surface into two disjoint regions, namely, the part inside the circle and the part outside. In some cases the opposition between "inside" and "outside" may be inappropriate (think of the equator separating the globe into two hemispheres), but the concept of *division* still applies: there are points on the surface that cannot be connected by a continuous path without intersecting the circle (you cannot drive from Paris to Cape Town without crossing the equator). If we use this definition it is clear that all sorts of objects qualify as spherical: a cube, a champagne glass, even a baroque chandelier may pass the test. However, a doughnut does not. On a doughnut, there are various ways one can draw a closed curve without dividing the surface into two disjoint regions (Figure 2). So a doughnut is not spherical.

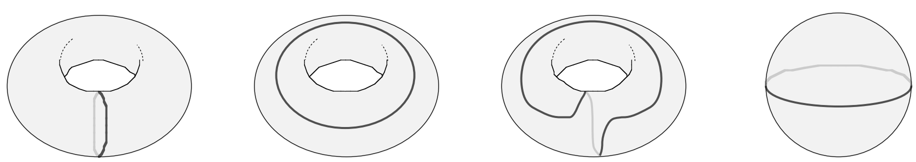

Figure 2. There are many ways of drawing a closed curve on the surface of a doughnut without dividing it into two separate regions; not so with the surface of a sphere.

Equivalently, one can say that the property that topologically distinguishes a doughnut from a sphere is this: any circle or closed curve on a sphere can be shrunk to a single point by elastic deformation; on a torus (the topologist's word for the surface of a doughnut) this is not the case: the curves in Figure 2 cannot be reduced to a point without "cutting" through the surface. Or again, the property in question can be characterized as follows: if two circles on a sphere intersect, they intersect in *two* points (this means "intersect" in the sense of going right through, not just touching); but on a torus two circles may intersect in just one point: consider for instance the circles in the two left diagrams of Figure 2 and imagine drawing them on the surface of the same doughnut.

All of these characterizations are basically equivalent. And they all serve the same purpose: they provide a means for distinguishing an object with a hole and an object without holes *exclusively in terms of the properties of the objects*, in fact, of their surfaces. No reference to the hole is necessary. One can even imagine a population of two-dimensional beings—the Flatlanders—living on the surfaces of our objects and trying to find out whether their two-dimensional world is spherical or toroidal. (They know they don't live on a plane, for they have discovered that if they go on a trip heading Eastward, say, they eventually get back to the same point from the West.) They cannot conceive of a torus as the surface of something *with a hole*, for holes are three-dimensional. Nevertheless topology gives them a key to the solution. Go on two round trips, and make sure to lay out a trail of red thread behind you. First go Eastbound and come back from the West. Then go Northbound and come back from the South. If the two threads intersect only at the starting point, Flatland is a torus; otherwise it is a sphere.[2]

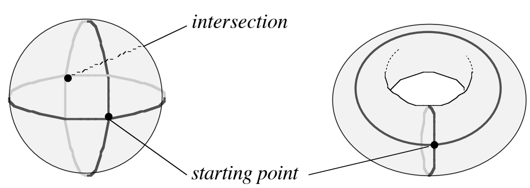

Figure 3. How Flatlanders can figure out the topology of their two-dimensional world.

When the Hole Matters

In a sense, then, topology can be described as a kind of geometry concerned primarily with the identification of holes. But topology is really not concerned with the holes: in topology the only thing that matters in a doughnut is the edible stuff. The hole is a mere *façon de parler*.

It is in this sense that topology gives support to what I have called the "traditional view". Topology provides a framework for maintaining that the

---

[2] The wonders of Flatland are described in Edwin A. Abbott's 1882 classic, *Flatland: A Romance of Many Dimensions* (reprinted by Dover, New York, 1952). More discussion of how Flatlanders can succed in establishing the shape of their world may be found in the first chapters of Jeffrey R. Weeks' book *The Shape of Space. How to Visualize Surfaces and Three-Dimensional Manifolds* (Dekker, New York, 1985).

locution ‘there is a hole in ...’ is neither more nor less than a shape predicate, albeit of a very special sort. In fact, this would be a modest achievement if one could not do the same with other hole-based locutions, such as ‘there are seven holes in ...’ or ‘there are more holes in ... than in ...’. But topology can do that easily. The topologist can count any number of holes without ever mentioning them. For instance, one can count the maximum number of disjoint circles or closed curves that can be drawn on the surface of the object without separating it into disjoint regions. On a regular doughnut that’s one; on a doughnut with two holes, that number is two; and on an object with *n* holes that number is *n*. That number, incidentally, is called the genus of the object’s surface. And the fundamental theorem of topology asserts that the surfaces of all three-dimensional objects can be classified exclusively (and completely) by their genus.

It is precisely this notion of genus that shows the full strength of the topological method for dealing with holes. A possession of any number of holes is reflected in the object’s genus. The question we must ask is: is it reflected in the right way? Is reference to holes made fully redundant by an analysis of the object’s genus? Unfortunately, the answer is not quite in the affirmative. Some, perhaps most hole-statements can be handled in this way. But there are cases where the topology of the object delivers the wrong answer, or it delivers an incomplete answer. And in these cases direct reference to the holes seems necessary. Let us then turn to this side of the story.

One obvious example is that the method is incapable of distinguishing between, say, straight and knotted holes. Reference to the object’s genus is good for establishing the “intrinsic” topology of an object—the topology as it can be figured out by a population of Flatlanders. However, the object’s extrinsic topology, i.e. the way the object is embedded in three-dimensional space, lies beyond that. One simply cannot tell what the embedding is like

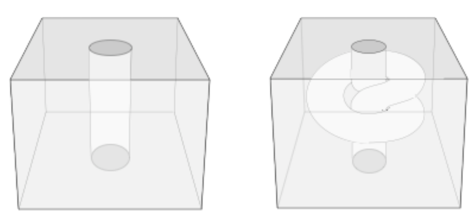

Figure 4. A straight hole, and a knotted one

by focusing only on the object. Indeed one generally cannot tell if the hole one is walking through is knotted or not. Nevertheless there is a significant difference between the two cases and one would like to be able to account for that difference.

Note that this is not a limit of topology as such. In fact, one way of capturing the difference between the two cases would be to consider the topology of the objects' complements. Clearly, the complementary topology of an ordinary doughnut and that of a doughnut with a knotted hole are different. But this shift—from the object to its complement—is crucial. *If* we were only talking about regions of space (as topologists often do), then all is fine: there is no significant difference between a region and its complement, and no reason to restrict oneself to one or the other. But if we are talking about *objects*, then the shift to complementary topology is quite significant. An object's complement is, after all, just as immaterial as a hole. In fact the hole is part of the complement. So the complementary topology of the object is, to some extent, the topology of the hole. The expressive power of the topologist's language is safe. But it doesn't save us from explicit reference to the immaterial. We must keep one eye upon the doughnut, and one upon the hole.

Here is another, more interesting example. Consider the four objects depicted in Figure 5. They are all of genus two, and indeed they can all be trasformed into one another by mere elastic deformation, without cutting or pasting. (Check that!) However, we do want to make distinctions here. For instance, we do want to say that the object on the left (*a*) has two holes, whereas one on the right (*d*) has only one hole (curiously shaped). The topology of the object simply delivers the wrong answer here: the objects may well be equivalent; the holes are not.

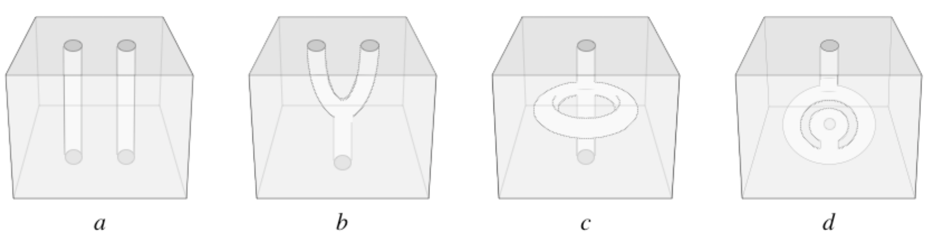

Figure 5. Four objects with the same genus but different holes.

One could say that this outcome is simply a sign of the counterintuitiveness of certain topological equivalences (the same counterintuitiveness, perhaps, that underlies the equivalence of a cube and a baroque chandelier). Once again, however, the problem does not lie in the conceptual apparatus of topology per se. It lies in the application. It lies in the idea that the only topology that matters here is that of the object's surface. If we look at the topology of the hole's surface instead, we get a completely different picture—indeed one that makes all the correct distinctions. By the surface of a hole I really mean that part of the surface of the object which envelops the hole, and which can only be individuated by reference to the hole. In a straight perforation, that superficial part is a cylinder: its normalized topological figure is a sphere with two punctures. In the case of a Y-shaped hole, it is a sphere with three punctures. And in the cases corresponding to Figure 5*c* and 5*d*, the surface of the hole is not a sphere but a torus with two punctures and a bitorus with one puncture, respectively.

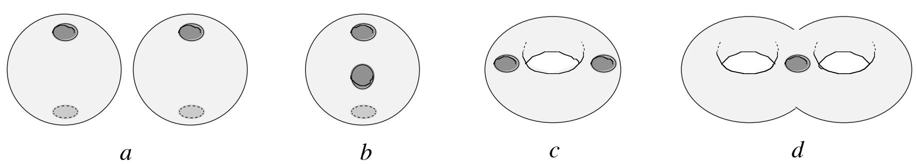

Figure 6. The holes in the objects of Figure 5 have surfaces of different topological kind.

Note that a puncture is not a real hole but a hole of lower dimension—what topologists also call an edge or boundary. The fundamental theorem mentioned above can be formulated more fully using this vocabulary: orientable surfaces are completely characterized up to equivalence by their genus and by the number of their boundaries. Now, the surfaces of material objects don't have any boundaries. But the surfaces of holes (in the sense explained above) do. And that makes all the difference.

The idea of looking directly at the topology of the hole is also helpful in seeing the family resemblance between the type of holes considered so far—perforations—and other kinds of hole. There are holes that are purely superficial, like a hole in a golfing green, or the nostrils of a baby-doll; and there are holes that are entirely hidden in the interior of their material hosts, like a cavity inside a wheel of Swiss cheese. These are all part of the big family of holes and indeed they all give rise to the same puzzle that we started with: you cannot buy the doll without the nostrils; and you cannot buy the cheese and leave the inner cavities at the grocery store. Now, if we look only at the objects hosting these holes, we need to come up with something else than the genus to describe the relevant geometric features. For instance, the presence of an inner cavity is reflected in the fact that the object (cheese) has two separate surfaces: the one that binds it on the outside (the crust) and the one that binds it on the inside (where the cavity is). And the presence of an external hollow is reflected in the fact that the surface presents an abrupt change in its curvature pattern, from positive (convex) to negative (concave). These may all well be effective ways of describing these and similar situations, but they would introduce a disturbing asymmetry among the various cases. By contrast, a hole-based perspective handles all cases in a uniform way. Indeed if we look at the surfaces of these holes we get exactly those patterns that were missing in the case of perforating holes: inner cavities yield unbounded surfaces (spheres, toruses, bitoruses, and so on); superficial hollows yield surfaces with one boundary. And then there are the mixed cases. *Holes come in various species, but they are all species of the same genus*.

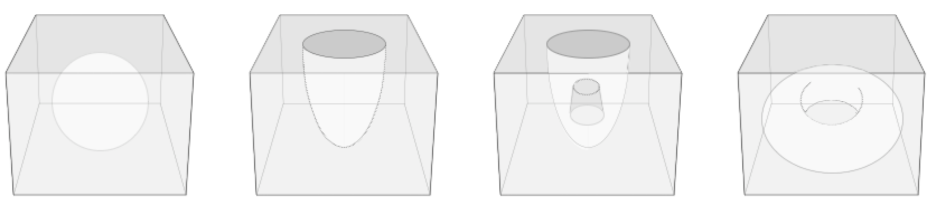

a *b* *c* *d*

Figure 7. Holes come in various species, besides perforations.

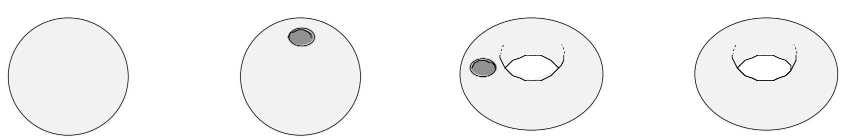

a *b* *c* *d*

Figure 8. The surfaces associated with the holes in Figure 7.

The Matter and the Void

Does it follow, then, that we should desert the topology of the objects and switch to the topology of the holes instead? Not at all. One has to be very careful at this point, for in some cases the surface of the hole may itself be deceptive. Just take cases such as 7*a* and 7*d*. In such cases we have the same problem we had before, except that the difficulty now concerns the inner surface—the surface of the hole. What if the cavity in 7*a* had a dent? What if the torus in 7*d* had a knotted hole instead? What if our wheel of Swiss cheese had cavities whose surfaces are like those of the objects in Figure 5? What if it had an extravagant cavity like the one in Figure 9 below? In all these cases, the same arguments apply: one point of view (whether doughnut-based or hole-based) is not enough, and some form of complementary reasoning seems required.

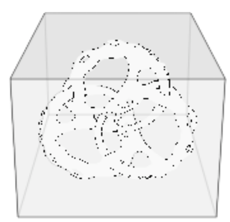

Figure 9. A very complex case: Can you describe the object without referring to the hole? And can you describe the hole without referring to the object?

We thus come to the moral of our little exercise. The interplay between holes and their hosts—between void and matter—can be very complex indeed. And the only way to address it properly and in a systematic way is to grant equal dignity to *both* characters: the void and the matter. Now, this means we must be very serious about reifying holes. We must be very serious about treating them as fully fledged entities, on a par with the material objects that surround them. And as with all cases of void reification, this may carry in its wake an array of difficult philosophical conundrums. (What exactly is the relationship between a hole and its material host? How come you cannot separate one from the other? Can you cut the doughnut without killing the hole? Can you eliminate the hole without destroying the doughnut? What if you fill it up—does the hole go out of existence? Does it squeeze to the side to leave room for the filling? Can you get rid of one half of the hole and keep the other?)[3]

One might protest that this conclusion is still unwarranted. All we have seen is that topology and perhaps geometry at large provide an account of the locution "There is a hole in ..." that does not fully support the view that holes are merely a *façon de parler*. But this falls short of establishing that holes are not, in fact, *façons de parler*. One could still try to do away with such nothings by relying on richer representation systems. For instance, one can resort to a description of the doughnut that combines topology and geometry with other properties, such as its kinematic properties. Eventually, and more generally, one could paraphrase every sentence of the form "There is a hole with such and such characteristics in that object" by means of a point-by-point description of the object in question, including a thorough account of all sorts of properties that are exemplified *at each point*. That should do. That should allow one to stick to the solid matter and avoid any commitment to its immaterial intrusions. But there is an obvious reply to such a strategy: a point-by-point paraphrase is simply too powerful a tool. You can use it to get rid of the hole; but you can also use it to get rid of the doughnut. You could just as well paraphrase every sentence about a doughnut by means of a thorough point-by-point description of the region of space that it occupies at each instant of time, combined with a complete account of the properties (of material constitution, color, texture, electric charge, etc.) that are exemplified at each point of that region. That would hardly be compatible with the idea that doughnuts are not *façons de parler*. But that is the unavoidable boomerang effect of such an eliminative strategy. For this is the dilemma of every eliminative strategy: if successful, it ends up eliminating *everything* just in order to eliminate *nothings*.

---

[3] A first foray into the philosophical territory defined by such questions can be found in Roberto Casati and Achille Varzi, *Holes and Other Superficialities* (MIT Press, Cambridge MA and London, 1994).