

Topology and cognition

Roberto Casati

CNRS Institut Nicod, Paris

Topological properties are among the most general spatial properties. Spatial representation as registered in vision and language appears to display sensitivity to them.

0. Introduction

We can see that a smiling face turns smoothly into a sad face, we can appreciate the difference between stretching an object and tearing it apart, we can distinguish between an object's being inside or outside another, we see the difference between a pretzel and a donut or between a broken and a whole glass, we can understand that in a sense the letter B has the same configuration as the digit 8 and that both differ from the digit 2. In all these cases, we explicitly countenance the *topological* structure of the scene we perceive or think about. Some important topological properties and relations are those of *continuity, interior, exterior, tangential part, boundary*, and the presence and number of *holes*. These properties – and their cognates – have attracted the interest of cognitive scientists as they may also implicitly figure in descriptions or explanations of cognitive performances. In a visual scene, the difference between an area corresponding to a figure and an area corresponding to the background is related to the fact that the figural area is assigned the boundary. Sensitivity to subtle topological differences has been alleged to explain our understanding of spatial prepositions such as 'across' or 'between'. Some psychological theories explain children's understanding of physical objects in terms of sensitivity to topological properties.

This constellation of facts suggests that a consideration of topological properties is required to account for some of the contents of spatial representation. However, there is at present no available systematic study of topology in cognition. This entry shall concentrate on some aspects of current research on cognition where topological notions and properties play a key role, whether as subject matter of explicit judgement, or as hidden components of psychological explanation. Although the visual and the linguistic domains will be the main focus, it should be mentioned that topology is countenanced in the study of other faculties. For instance, haptic continuity has been described and in the domain of action it has been surmised that reaching is guided by the hypothesis that objects move on continuous paths. In the auditory domain one can find loose analogies for the notion of contact (such as collisions, that may be heard).

Various uses of topology in cognitive science do not figure here, such as the study of the topology of neural networks, or the topology of color and other quality spaces; these are topological instruments for studying cognition, not objects of cognition. It is hoped that this entry will contribute to a clarification of the uses of the term ‘topology’ in cognitive science.

1. What is topology?

Topology as a branch of mathematics focuses on basic abstract spatial properties. How abstract? From a topological point of view, a square and an ellipse are the *same* shape. In a first approximation, topological properties are those that are not altered by “distortion” and “stretching”. In general, we can classify the geometric properties of a shape by noting whether they are or are not preserved in some classes of transformations. Copying or rotating a square will preserve all of its shape properties such as the length of its sides or the width of its angles. Stretching a quadrilateral may alter the width of its angles and the length of its sides; accordingly, these are *not* topological properties of the quadrilateral. The projection that takes a square onto a larger one is a transformation that preserves the width of the angles but not the length of the sides. The transformation that carries a square onto an ellipse seems to destroy all geometrical properties of the square. The square is no longer “recognisable” and still, some of its properties are preserved in the ellipse. Both the square and the ellipse are closed curves. If point a is between points b and c on the square, the projection of point a will be between the projections of points b and c on the ellipse. The order of points is not altered through distortion and stretching, and counts as a topological property of the square.

Cognitive science deals with relatively simple and intuitive topological properties of objects, such as being in one piece or having holes. The scope of the mathematical theory of topology encompasses situations that are not as easily graspable as the ones suggested by the foregoing examples, but the basic idea is simple. A topological characterisation of a set S of objects – intuitively, points – is a way of using subsets of S in order to describe relations among points in S . Subsets of S allow one to separate points of S . The set of these subsets, called a topology T , must be such as to contain: S itself, the empty set \emptyset , the union of any collection of sets in T , and the intersection of any finite collection of sets in T . A given set S can have different topologies which correspond to different collections of subsets of S , each with a different discriminatory power. A discrete *topology* (one that contains all subsets of S) is able to discriminate all points in S . A *topological space* is a set S with an associated topology T .

Countenancing individual points and sets of points permits to define neighborhoods of a point, which lead to a notion of closeness (a generalization of the metric notion of distance) in terms of which continuous functions and continuity are defined. This in turn makes the notion of a topological space – defined in the relatively sparingly terms above – a tool for classifying remarkable properties of figures. A topological space is *connected* whenever in its topology no two disjoint non-empty sets exist such that their union is the space. A topological space may be *path-connected*. A space S is path-connected whenever for any two points a, b in S there is a continuous path whose extremes are a and b . In general there will be many paths connecting any two points of a path-connected surface. These paths can be grouped in homotopy classes, members of which can be

transformed into one another via a continuous mapping. If only one homotopy class exists, the object *is simply connected*. For instance, the sphere is simply connected. An intuitive way of putting this is to consider a particular class of paths, those whose ends coincide, i.e., loops. Many loops will pass through a given point on a sphere, but to see that they are homotopic we need only picture them all as capable of shrinking to a point. On the other hand, on a toroidal (donut-shaped) surface not all loops through a given point will be such as to shrink to a point; neither would all loops on a plane on which a finite portion has been removed. In general, surfaces can be classified in terms of the homotopy classes of the loops we can draw on them. The presence of distinct homotopy classes on a space indicates the presence of discontinuities that intuitively prevent loops from shrinking. These discontinuities are most typically holes and handles.

2. Relevance of Topology to Cognition

Topology enters cognitive science as the object of a special, indeed fundamental, case of spatial representation. Topology related spatial representation can be isolated within the class of *non-immersed* spatial representations, as opposed to immersed representations such as ‘to the right of’, whose interpretation depends upon a viewpoint embedded in the representation (be it object- or view-centered). For instance ‘is one meter long’ or ‘is inside the box’ are non-immersed representations as their applicability is not affected by changes in point of view. This ball is or is not inside this box no matter how we look at these two objects, whereas this ball is to the right of this box only from a given viewpoint.

Within the class of immersed spatial representations one can further isolate the representations that are size and shape related from those that are not. We can perceive and judge about spatial facts without necessarily having to assess the size and distance of the objects involved. One can for instance judge that John is between Paul and Mary without accessing any specific metric information concerning the relative sizes and distances of these three people. Typically, topological representation abstracts from size and shape. Moreover, the output of a judgement about the topology of a situation is of a yes-or-no type, whereas the output of a metric judgement is continuous.

It appears from some of the opening examples that topological properties appear to be required to describe some of the contents of spatial representation. But do we possess an implicit knowledge of topology? At what level is this knowledge accessed, and how is it used? Are topological relations computed online at some level?

We can distinguish here two types of evidence: Direct evidence, from personal or sub-personal assessment of topological facts, and indirect evidence, stemming from an analysis of data that although do not concern an explicit topological representation, seem to require some access to a topological representation or processing.

Direct evidence for topological analysis

As to personal, direct evidence, consider an object such as the Möbius strip (a strip whose ends have been twisted once and glued together). If one is asked to predict the result of cutting it along its median line (one which runs parallel to its borders), one will probably

infer that *two* strips are to be expected. (As a matter of fact, one gets again a Möbius strip, this time doubly twisted). The *border* of a Möbius strip is, upon reflection, weird. It is as if one reasoned from the principle that if there is a border, then the two sides separated by the border are such that one cannot reach one of them from the other without crossing the border. Once you are on one side, it seems, the border bounds you to that side. But the Möbius strip is exactly what contravenes this principle.

Exposure to some technical topology does not, by itself, correct one's propensity to misjudge some topological facts. Those who understand and master simple topological tasks, such as the detection of the topological equivalence between a sphere and a cube, or of a donut and a cup of coffee, may be unwilling to accept that a cylindrical surface (such as the portion of a tube) and a disc with an internal portion removed (such as a CD disk) are topologically equivalent. Popular topology books emphasize the oddity of the topological equivalence among objects that look utterly different. The case resembles that of intuitive physics. It has been found that subjects' understanding of Newtonian physics is blocked by their implicit and automatic reliance on an intuitive theory that is at odds with Newtonian physics. Analogously, there may exist an *intuitive topology* that prevents us from appreciating actual topological equivalencies (or non-equivalencies). Intuitive topological classification may be the by-product of independent criteria of classification that emphasize some other spatial aspects of objects - such as the presence of holes, conceived as objects in themselves - (Casati and Varzi, 1999), or else be the by-product of higher-order physical principles that dictate plausible shapes for physical objects.

Sub-personal direct evidence presents methodological problems, related to the extreme generality of topological properties. Experimental designs may be not sufficiently fine-tuned to detect sensitivity to topological structure. For instance, in experiments about numerosity, a test display with two black spots may follow a habituation setting in which a single black spot is present. Sensitivity to the change can be related to a sensitivity to numerosity (implying a sensitivity to disconnection, hence to a topological property of the setting) but it may as well reflect a sensitivity to the overall quantity of blackened space or to the overall length of the perimeter of the two spots. The experimental setting should control for some of these factors, but the trouble with topological properties is that they are so basic and ubiquitous that almost any factor interferes with them.

[Here figure number IV, provided after number I-III]

There is evidence that perceptual priming can be size-, location- and orientation-independent (Biederman and Cooper 1990, 1992), showing that recognition processes can abstract from these spatial features of objects. Does abstraction go all the way down to the most basic, topological features? Experiments so far have provided little evidence. Chen (1982, 1990) has advocated that the extraction of topological properties is one of the primary functions of the visual system. This claim may be supported by the experimental finding that subjects are better at discriminating pairs of topologically distinct figures than pairs of topologically similar but metrically distinct figures, and by the occurrence of illusory conjunctions in displays that involve holes. However, it has been pointed out (Rubin and Kanwisher 1985) that other factors (such as differences in luminous flux and in the overall perimeter of the figures) could account for most of these results.


Indirect evidence for topological analysis

If sensitivity to topological structure is difficult to assess directly in relation to configurational properties of perceptual displays, it is nevertheless possible to postulate that some mechanism for evaluating the topology of the display is activated in visual perception.

As a matter of fact, some cognitive abilities do not *prima facie* relate to topology, but turn out to presuppose sensitivity to the topological parsing of a scene. We shall consider (A) the assessment, at personal level, of object unity, (B) the sub-personal segregation of perceptual units (C) part-whole parsing, (D) infants' mastering of the notion of a physical object, and (E) representational advantages. Further evidence is related to language, and will be discussed in section 3.

(A) Personal assessment of object unity

The display in fig. 1 could be described in two mutually incompatible ways:



- (1) There are *two* black objects, one on the left and one on the right
- (2) There is *one* black object, partly on the left and partly on the right.

Although the second description sounds utterly artificial, there are no obvious reasons for ruling it out. The source of the awkwardness seems to lie in the fact that in judging what counts as one object in the display we rely on background assumptions concerning some topological properties of the display itself. There are two *disconnected* black patches, each of which is *maximally self-connected*. This suggests that we use connectedness as a criterion of unity, and that the visual system is able to compute it. If topology biases the choice, this in turn entails that cognition takes topology into account.

However, a general rule such as: “Always interpret topological self-connectedness as necessary and sufficient as sufficient for unity” does not generally hold. Two juxtaposed objects may count as two even though they are connected, and some kinds of causal unit (such as two objects in coherent motion) may not require connection.

(B) Segmentation of the visual field

Gestalt psychologists have proposed an explanation for perceptual organization (e.g. figure-ground articulation) of the visual field in terms of our tendency to group perceptual units according to principles such as similarity or relative proximity. Perceptual units are discrete elements that are put together and considered as parts of superordinate units in virtue of some organizing factors. However, this account takes for granted the formation of the very perceptual units that get grouped, the ‘atoms’ of the visual field. Palmer and Rock (1994) claimed that these units are regions of the visual

field that are *uniformly connected*. Uniform connectedness is not a principle of grouping insofar as it does not presuppose units but creates them and as a consequence must operate prior to any grouping activity.

In order to see the application of the principle we may ask why, in the context of figure II, the left blob, say, is treated as a unit, whereas none of the (2-5) is, and whereas (6) the complement of the total black area is treated as a unit. (Note that dashed white lines do not belong to the image, they are used to pick out the relevant part of the image, furthermore, in (5) the putative unit is the total black area and in (6) its complement, the total gray area.




Fig. II

Connectedness is not sufficient to exclude (3): the dashed area is connected but it is not a perceptual unit. Nor is an unit the “conjunctive object” made up of the two black spots in (5), implying that uniformity is not sufficient either. But now uniform connectedness, the condition resulting from pairing uniformity and connectedness, is not sufficient to exclude (2) and (4), which are uniform and connected *parts* of perceptual units. One must reinforce the condition and add maximality: an image is partitioned into *maximal* self-connected regions of uniform quality. Note that the complement of the total black area is maximally uniform connected and hence counts as a unit. Observe that the questions of both object unity and unit formation are subtle. Consider two regions that are tangent on a point. Are they to be taken as one region or two? Shape may in some case influence judgements about unity. Moreover, shape can dictate that although a certain region is labelled as unitary, it can still have clearly demarcated parts (if we connected the two black regions with a black line, we would obtain a unitary object that would still have salient parts.)

(C) Part-whole parsing

Parsing is the labelling of some perceptual items as salient parts of other items. (Many parts of a perceptual unit are not at all salient.) As such, it is subsequent to unit segregation. It may follow topological principles, albeit topological properties and relations are cognitively intertwined in a complex way with part-whole relations. Unfruitful attempts to reduce topological connectedness to part-whole relations could only show once more the all-present bias towards topological unity. One may think that the two boxes in fig. 1 are disconnected because any “bridging” object that has parts in common with them will have parts in common with the complement. But this is the case only if the bridging object is a connected unit. If one accepts disconnected units to begin with, one may end up with a disconnected bridging object that is exactly composed of the two boxes.

Part-whole relations hold between entities that may not be labelled as salient parts, such as a square and its top half. Surely, a disconnected object such as a broken glass has as many salient parts as are the self-connected units composing it. But a self-connected object may still have salient parts that may be individuated by the geometry of the object, typically by concave discontinuities on an object’s contour (Hofmann and Richards 1984). A V-junction is a topologically unitary object and is topologically equivalent to a U, however the former, but not the latter, is parsed as the conjunction of two elements. Similarly, the topological identity of a Y-junction and a T-junction can be overruled by their configurational differences, so that we do make a difference and consider the Y as composed of three, the T of two bars.

(D) Physical objects in infant cognition

Infants appear to divide up the world into objects according to some basic principles (Spelke 1990). In particular, objects are assumed to move along *continuous* paths, and to interact if and only if they are in *contact*. These principles presuppose that infant cognition takes into account topological features such as contact and continuity. These principles need not require an explicit perceptual representation. A hidden connection may be inferred from a given display (for instance, if there is an occluder in view.)

(E) Same-object advantages

Subjects are faster and more reliable at assessing local changes within a unitary object than they are across disconnected objects. This same-object advantage means that topological unity is countenanced at the sub-personal level in addressing the task.

3. Topology in language

Linguistic data provide a further body of indirect evidence for topological evaluation. English contains a vast number of geometry related terms for spatially describing objects (‘square’, ‘circle’), and properties (‘circular’). In principle, we can describe countless shapes by simply building a predicate that mentions an object having that shape (‘star-shaped’). Some linguistic uses, though, presuppose an appreciation of very abstract geometric properties. For instance, the correct use of the preposition ‘in’ largely abstracts

from the shape and size of the objects that stand in the corresponding relation. A fish can be in a pool and in the same sense of ‘in’ a submarine can be in the ocean.

A broad distinction between two general classes of linguistic items, the open or lexical class and the closed or grammatical class, is believed to reflect two levels of cognitive generality (Talmy 2000). The membership of the open class is subject to great variability, as opposed to the rather fixed membership of the closed class. Grammatically significant items (prepositions, suffixes, prefixes, etc.) cluster in the closed class. For instance, only a few prepositions get added to or disappear from English over time (‘betwixt’). A commonly held hypothesis is that closed class items tend to capture general facts of cognitive import. This would in turn constitute an explanation of the stability of the class over time. In particular, only certain concepts are expressed by items in the closed class, and these concepts have a cognitive structuring role. Language does not have grammatical expressions for size, shape or color. There is no preposition that can be used to convey the idea that an object is square or red. By contrast, the presence of a plane so curved as to define a portion of space is conveyed by uses of the preposition ‘in’, and these uses abstract from size (‘in the thimble’, ‘in the volcano’) and shape (‘in the well’, ‘in the trench’).

An important set of these “structuring” concepts concern space and spatial properties and relations that are shape- and size-neutral, as in the example of ‘in’ above. Similarly, in the sentence ‘I walked through the woods’, the preposition ‘through’ expresses a concept that is indifferent to the shape and size of the path followed. These very abstract concepts come close to the abstract notions of topology, so much that it has been suggested (Talmy 2000) that there could exist a language-based topology whose notions are akin to those of mathematical topology. The notions that would be part of this language based topology are, among others, those of point, linear extent, locatedness, withinness, region, side, partition, singularity, plurality, sameness, difference, adjacency of points. It can be argued that topology as a mathematical theory is a normative refinement of these notions. Note that not all topological relations are lexicalized. There are words for three-place places relations such as ‘between’, but not for four-place relations.

4. Theories of topological representation

Topological classification

Topology as a mathematical theory provides a classification of objects’ shapes. But cognition does not appear to match closely the classification provided by mathematical topology. Matching can occur for very simple surfaces, but with somewhat more complicated surfaces the respective classifications diverge, as is indirectly proven by the counter-intuitiveness of many topological equivalencies. So far no systematic study exists about the actual classifications endorsed by cognition. A hypothesis is that shapes are classified in terms of the number and position of holes (tunnels, indentations) they have, whereby holes are counted as parts of the object or as objects of their own.

Schematic spatial representation

Linguistic representation appears to be relatively independent from size, shape and oftentimes from the type of object represented. It has been surmised that the language

system accesses a particular level of spatial representation, the level of geometric schematization, which is also interfaced with the visual system (see Herskovits 1997 for a general overview). As this level objects are assumed to be represented by simplified geometric schemata that may or may not take into account the objects' metric properties. The study of the structure of prepositions can be used as a heuristics for the study of the level of schematization. How does language access this level? According to Herskovits (1997), prepositions do not refer to schemata, but express constraints on allowable schemata. It is an open question how a proposition such as 'the ball is in the box' is evaluated. Possibly the evaluation relies on an online computation performed at the schematization level that answers a "command" from the linguistic system.

Online computing

If topological facts are computed at some level in spatial representation, the problem arises of finding the mechanisms that do the computing (Ullman 1996 for an overview). Various algorithms have been proposed. For instance, a simple routine such as the one necessary to find out whether two marks lie on the same line or on different lines may explore the line starting from one mark (and then restart from the mark going in the opposite direction) and stop at an endpoint or at the other mark. In order to determine whether a given point is within or without a certain closed region, a "filling in" algorithm can colour the field starting from the point until the boundaries of the shape are met. Algorithms of this sort risk to be too powerful in that they may deliver definite answers where the visual system does not output a definite judgement. For instance, we may not be able to decide whether point a is inside or outside the closed curve or a point b is on the same line as point c in fig. III, whereas a filling in algorithm and, respectively, a line-exploring algorithm will be adamant in delivering an answer.



Fig. III

From the fact that topological properties are basic and ubiquitous it does not follow that their determination is accomplished by simple devices. As a matter of fact, it has been claimed (Minsky and Papert 1969), that evaluating topological properties may not be a simple task at all. Minsky and Papert discuss the limitations of perceptrons, devices capable of assessing whether a certain configuration satisfies a certain predicate by computing the linear (purely additive) results of collections of weighted predicates, each of which is true or false locally (say, a non-maximal cell assembly on a retinal sheet).

They proved that some spatial predicates, such as ‘is convex’, can be easily computed by perceptrons, but others, such as ‘is connected’ are stubbornly resilient to computation for most classes of perceptrons. Connectedness may be computed only by fairly complex systems.

Glossary

Connection. The relation holding between two regions that share a point (either overlap or are in touch.)

Schema: A simplified representation that abstracts from some of the properties of what it represents.

Topological equivalence. The relation holding between two regions that share a common topological structure (intuitively: regions whose shapes can be transformed into each other by elastic deformation).

Unit . A portion of the perceptual field that is bestowed some representational advantage. Basic units are connected, uniform and maximal regions.

Bibliographic references

Biederman I, Cooper EC (1990) Evidence for complete translational and reflectional invariance in visual object priming. *Perception* **20**: 585-593.

Biederman I, Cooper EC (1992) Size Invariance in Visual Object Priming. *Journal of Experimental Psychology: Human Perception and Performance* **18**: 121-133.

Casati R, Varzi AC (1994) *Holes*. Cambridge, Mass.: MIT Press.

Casati R, Varzi AC (1999) *Parts and Places*. Cambridge, Mass.: MIT Press.

Chen L (1982) Topological structure in visual perception. *Science* **218**: 699-700.

Chen L (1990) Holes and wholes: A reply to Rubin and Kanwisher. *Perception and Psychophysics* **1990**: 47-53.

Halligan PW, Marshall JC (1993) When two is one: a case study of spatial parsing in visual neglect. *Perception* **22**: 301-312.

Herskovits A (1997) Language, Spatial Cognition, and Vision. In: Stock O (ed) *Spatial and Temporal Reasoning*, pp. 155-202. Dordrecht: Kluwer.

Hoffman DD, Richards WA (1984) Parts of Recognition. *Cognition* **18**: 65-96.

Minsky ML, Papert SL (1969) *Perceptrons*. (1988, Expanded Edition.) Cambridge, Mass: MIT Press.

Palmer S, Rock I (1994) Rethinking perceptual organization: The role of uniform connectedness. *Psychonomic Bulletin and Review* **1**: 29-55.

Piaget J, Inhelder B (1967) *The child's conception of space*. New York: W. W. Norton.

Rubin JM, Kanwisher N (1985) Topological perception: Holes in an experiment. *Perception and Psychophysics* **37**: 179-180.

Spelke ES (1990) Principles of object segregation. *Cognitive Science* **14**: 29-56.

Talmy L (2000) *Toward a cognitive semantics*. Cambridge, Mass.: MIT Press.

Ullman S (1996) *High-Level Vision*. Cambridge, Mass.: MIT Press.

Vandeloise C (1994) Methodology and analyses of the preposition 'in'. *Cognitive Linguistics* **5**: 157-184.

Further reading

Hilbert D, Cohn-Vossen S, (1932, 1952 trad.) *Geometry and the Imagination*. New York: Chelsea Publishing Company. [A classical study of spatial imagination in mathematics.]

Mendelson B (1990) *Introduction to Topology*. New York: Dover. [Technical but accessible introduction to mathematical topology.]

Scholl B.J. (2001) Objects and Attention: The State of The Art. *Cognition*, **80**: 1-46. [The state of the art on various types of object advantages.]