

Roberto Casati - Achille Varzi

## Basic Issues in Spatial Representation

This paper has been published in M. De Glas-Z. Pawlak (eds) *Proceedings of the 2nd World Conference on the Fundamentals of Artificial Intelligence, Paris, Argkor, 1995* (pp. 63-72). Every quotation has to be done from that source. This material is only for private circulation.

**Abstract** - This is a preliminary report on a long term project on the foundations of spatial representation. In perspective, our aim is to unfold a general framework where the major results and open problems in this field can be set up and addressed in a uniform fashion. Here, our concern is mainly with (i) analysing some notions that can be (or have been) used as primitives for this purpose (points, places, regions, and bodies in the first place), and (ii) assessing the degree and significance in which the choice of any such primitive may condition the resulting representation system. We are also interested in investigating connections between spatial representation systems and various fundamental intuitive notions such as localization, penetrability, and movement, though we shall not here go beyond some very general remarks. In this regard our purpose is mainly to illustrate what we take to be an important methodological desideratum, namely a combination of some affinity with common sense (on the one hand) with an adequate degree of formal sophistication and efficiency (on the other). Although this can only be appreciated in regard with specific case-studies, we believe that a fruitful interplay between methodological and substantial aspects will lay the groundwork for an innovative perspective on the potentialities of spatial representation theories.

### 1. Introduction

Spatial representation has in the last decade become a major focus of research in cognitive science and related areas. This includes philosophy as well as psychology, brain science, linguistics, those branches of computer science involved in the construction of machines capable of autonomous interaction with the environment. A *fil rouge* connects the efforts made by researchers from all of these areas under the rubric of spatial representation. Overall, however, there is also some ambiguity as to what 'representation' really means in this regard. Two main options are:

- (a) a theory of the way an organism (a brain, a mind) but also a cognitive sub-system (such as a language or a fragment of a language) represents its spatial environment (this representation serving the twofold purpose of organizing perceptual inputs and providing sufficiently articulated grounds for behavioral outputs);
- (b) a formal theory of the geometric representation of space capable of accounting

be presupposed by, a psychological theory in the sense of (a): one can reason about space in a totally abstract fashion, without any intuitive representation of it (and the purpose of a type (b) theory is, ideally, to provide a system of axioms and rules of inference whose execution does not depend on their intended interpretation). On the other hand, it seems highly plausible that not only lower animals, but also primates and humans possess a system of spatial representation that does not require the execution of inferential procedures in the sense of (b). Thus, one may have a type-(a) theory without having a type-(b) theory, and vice versa. Even so, this relative independence does not rule out the possibility that the study of spatial representation in either sense be useful for the study of spatial representation in the other sense. After all, the abstract concepts employed by a geometer stem from the unsophisticated (or seemingly unsophisticated) notions of common sense 'plane', 'vertex', 'point', 'concavity', 'inclusion' and these latter exhibit idiosyncratic features that can be reasonably ascribed to the functioning of the cognitive system dedicated to the representation of space. Conversely, if these intuitive notions do not constitute an undifferentiated conceptual scheme, this could be explained by the existence of a structure that can be modeled by a formal theory.

Even if spatial representation is becoming a well established field of (interdisciplinary) research, it is fair to say that many fundamental issues have not been sorted out yet. Relevant work tends to be local rather than taking place within global, over-all conceptual frameworks; and although the resulting variety of theories have a corresponding variety of merits, how to assess these and where to look for further progress is subject to the possibility of providing more extensive accounts at the foundational level. What we have in mind is a number of questions that lie in a philosophical region partly overlapping theories of type (a) as well as theories of type (b). Both kinds of theory speak of spatial entities and spatial environments. But what sort of entities are these? And how do they relate to one another how does a spatial region relate to the entities (people, material objects, events) that may inhabit it? Is the former somehow dependent on the latter? Are there things that consist merely of space? What does it mean for something to occupy space, as opposed to simply being located in space? The parts of a thing occupy parts of the region occupied by the thing, but what exactly is the relationship between truly spatial relations such as 'contained in', or 'located between' and purely mereological (part-whole) relations? These and many others are all fundamental questions that a spatial representation theory cannot eschew, regardless of whether one's concern is with (a), with (b), or with a combination of both approaches.

What follows is meant as a contribution precisely in this direction. Our long term aim is to unfold a general framework where these issues can be approached in a uniform fashion. In prospective, this is likely to require some metaphysical commitment as to the true nature of space. Here however we shall confine ourselves to some preliminary work. In fact our approach will be mostly methodological. We shall

of the questions mentioned above, as well as of other fundamental issues in spatial representation such as, for instance, the correspondence between perceptual and physical space or the representation of egocentric or indexical spatial properties.

## 2. Parts and wholes: a problem for Whitehead

Much recent work on spatial representation has focused on mereological and topological concepts, and the interaction between these two domains constitutes a first major issue.

There is, in fact, no question that a considerable portion of our reasoning about space involves mereological thinking, that is, reasoning in terms of the *part* relation. Traditionally this has been associated with a nominalistic stand, and mereology has been presented as a parsimonious alternative to set theory, dispensing with all abstract entities or, better, treating all entities as individuals. However there is no necessary internal link between mereology and nominalism, and we may simply think of the former as a theory concerned with the analysis of parthood relations among whatever entities are allowed into the domain of discourse. (This of course fits in perfectly with the spirit of type-(b) theories, but type-(a) theories may also be seen this way.) Even so, it is clear that a purely mereological outlook is too tight unless one integrates it with concepts and principles of a topological nature. Without going into much detail (see Varzi [1994]), the reason is simply that mereological reasoning by itself cannot do justice to the notion of a whole (a self-connected whole, such as a stone or a rope, as opposed to a scattered entity made up of several disconnected parts, such as a broken glass or an archipelago). Parthood is a relational concept, wholeness a global property. And in spite of a widespread tendency to present mereology as a theory of parts *and* wholes, the latter notion cannot be explained in terms of the former. For every whole there is a set of (possibly potential) parts; for every set of parts (i.e., arbitrary objects) there is in principle a complete whole, viz. its mereological sum, or fusion. But there is no way, mereologically, to draw a distinction between 'good' and 'bad' wholes; there is no way one can rule out wholes consisting of widely scattered or ill assorted entities (the sum consisting of our four eyes and Chisholm's left foot) by reasoning exclusively in terms of parthood. (Whitehead's early attempts to characterize his ontology of events provides a good exemplification of this difficulty. His mereological systems [1919, 1920] do not admit of arbitrary wholes, but only of wholes made up of entities that are 'joined' to each other. This relation is defined thus:

$$(1) \mathbf{J}(x,y) = \text{df } \exists z(\mathbf{O}(z,x) \ \& \ \mathbf{O}(z,y) \ \& \ \forall w (\mathbf{P}(w,z) \rightarrow \mathbf{O}(w,z) \vee \mathbf{O}(w,y)))$$

where 'P' indicates parthood and 'O' overlapping. But it is immediately verified that this definition falls short of this task *unless it is already assumed* that the piece  $z$

### 3. The topological option: a problem for Clarke

Mereology cannot suffice for the purpose of spatial representation even if we confine ourselves to very basic patterns. Not only is it impossible to capture the notion of one-piece wholeness; one cannot even account for such basic notions as, say, the relationship between an object and its surface, or the relation of something being inside, abutting, or surrounding something else. All of these are phenomena that run afoul of plain part-whole relations, and their systematic account requires a topological machinery of some sort. Now, in recent AI literature, this intuition has been taken to suggest that topology is truly a more basic and more general framework subsuming mereology in its entirety. In other words, if topology eludes the bounds of mereology, then one should better turn things around: start from topology right away and define mereological notions in terms of topological primitives. For just as mereology can be seen as a generalisation of the even more fundamental theory of identity (parthood, overlapping, and even fusion subsuming singular identity as a definable special case), likewise topology can be seen as a generalisation of mereology, where the relation of connection takes overlapping and parthood as special cases. (This view was actually considered by Whitehead himself [1929], but it was only with Clarke [1981, 1985] that it was fully worked out. Recently it has been widely employed in AI, e.g. for work in spatio-temporal reasoning [Randell & Cohn 1989, 1992; Randell *et al.* 1992a, 1992b] and natural language processing [Vieu 1991, Aurnague & Vieu 1993a, 1993b].)

The proposed subsumption of mereology is straightforward: given a relation of topological connection ('C'), one thing is part of another if everything connected to the first thing is also connected to the second:

$$(2) \mathbf{P}(x,y) = \text{df } \forall z (\mathbf{C}(z,x) \rightarrow \mathbf{C}(z,y))$$

Obviously, the reduction depends on the intended interpretation of 'C' (which is generally axiomatized as a reflexive and symmetric relation). If we give 'C' the same intuitive meaning as 'O', then (2) converts to a standard mereological equivalence. But things may change radically on different readings. Typically, the suggestion is to interpret the relation 'C(x,y)' as meaning that the *regions*  $x$  and  $y$  have at least one *point* in common. This means two things. First the domain of quantification consists of (spatio-temporal) regions, and not of ordinary 'things'. Second, since points are not regions, sharing a point does not imply overlapping, which therefore does not coincide with (though it is included in) connection. In other words, things may be 'externally' connected. There are of course some immediate problems with this account, for the absence of boundary elements in the domain means that things can be topologically 'open' or 'closed' without there being any corresponding mereological difference. In fact various refinings are available that avoid this unpalatable feature, so we need not go into these details. Suffice it to say that with the help of 'C' it becomes easy, on some reasonable interpretation, to capture various topological notions and to account for

So if spatio-temporal regions are the only entities of our domain, then (2) yields important conceptual achievements. However, there is a second side of the coin. For if we are to take an open-faced attitude towards real world things and events (without identifying them with their respective spatio-temporal co-ordinates), then the reduction offered by (2) seems hardly tenable, as different entities can be perfectly co-localized. A shadow does not share any parts with the wall onto which it is cast. And a stone can be wholly located inside a hole without actually being part of it. The region that it occupies is part of the region occupied by the hole but that's all. From here, intuitions diverge rapidly. And the notions of connection and parthood that we get by reasoning exclusively in terms of regions, no matter which specific interpretation we choose, just seem inadequate for dealing with the general case. (This means, among other things, that the possibility of extending this theory to neighboring domains might suffer. For instance, a theory of events which reduces mereology to topology by mapping every event onto the interval or instant of time of its occurrence (as do most AI theories of temporal reasoning developed under the impact of Allen [1981]) will not have room for co-temporal distinct events, let alone events occurring in the same spatio-temporal regions.)



Figure 2. Clarke's problem:  $x$  is inside object  $y$ ; but what is the relationship between object  $x$  and hole  $z$ ? (and what the relationship between  $z$  and  $y$ ?)

#### 4. Holes: a problem for Gotts

These objections may not be definitive. In particular, our example of the stone in the hole presupposes a friendly attitude towards the ontological status of holes, which is all but unproblematic. (We shall come back to this below.) Even so, it is apparent that the simplification introduced by (2) has critical consequences if our concern is with the foundations of general-purpose representation systems. For the basic issue of the relationship between an object and 'its' space (the space where it is located) is trivialized. Moreover, it yields a flat world in which every morphological feature is ignored, and the question of whether holes should be treated as *bona fide* entities next to ordinary objects, far from being left in the background, cannot even be raised. This, we maintain, is not only a source of conceptual poverty; it may also be misleading.

Some recent work by Nick Gotts [1994a, 1994b] is indicative of the problems we have in mind. Clarke's system and its derivatives include among their models an infinity of topological spaces. But the notion of a topological space seems to be much less specific than is required by our spatial intuitions. So Gotts asks: What additional axioms should 'C' satisfy (besides reflexivity and symmetry) in order to capture such intuitions? Gotts shows that using 'C' as a primitive we can describe toroidal structures

superficial hollows, grooves, and other discontinuities of irreducible morphological nature. Of course this is not a real problem if we treat superficial holes as uninteresting. If we confine ourselves to topology, we *must* do so, regardless of whether our primitive is 'C' or something else. This is not a real objection to Gotts; rather, it shows that topology is only one step ahead of mereology, and need be integrated by other notions and principles if we want to go beyond a world of spheres, toruses, and little else. The second trouble is more specific. For as it turns out, Gotts' system is intrinsically incapable of capturing the notion of a *knotted* hole. That is, it captures the intrinsic topology of a holed object, not the extrinsic topology. Now, of course knotted holes are just as important as straight ones, as it were. (And surely, you can hardly tell if the hole you are walking through is knotted or not.) But if the theory can't tell the *difference*, its classificatory power is seriously deficient.



Figure 3. Gotts' problem: we can't 'C' well enough to tell a straight from a knotted hole.

## 5. Components and negative parts: a problem for Biederman

We take the foregoing to imply a threefold moral. Firstly, one needs both mereology and topology, as independent (though mutually related) frameworks. Second, the limitations of topology are dramatic even at the elementary level (a long way before functional features become important for classifying shapes or providing an analysis of the containment relation; see Vandeloise [1994]). Third, one had better abandon an approach to spatial representation and spatial reasoning in terms of spatial regions, and consider from the very beginning an ontology consisting of the sort of entities that may inhabit those regions. (As we said at the beginning, this is of course tied in with the difficult metaphysical issue of whether we can *dispense* with spatial items altogether. This is the controversy between spatial absolutism the Newtonian view that space is an individual existing by itself, independently of whatever entities may inhabit it, and is in fact a *container* for the latter and spatial relativism the Leibnizian view according to which space is parasitic upon, and can be construed from, objects and relations thereof. But we believe one can remain neutral with respect to this issue at least at the beginning.)

A potential candidate in this direction is Biederman's [1987] 'Recognition By Components' (RBC) theory. This theory Biederman's concern is with shape recognition is based on the primitive notion of normalized cylinder, or 'geon', and offers a simple 'spatial syntax' whereby every object can be viewed as composed out of cylinder-like components. (The basic idea has been used by several other authors and is usually traced back to the work of Thomas Biford; Biederman should

the first at both ends. (In a more recent formulation [1990], both geons and relation among them are defined in terms of more primitive parameters, such as variation in the section size, relative size of a geon's axes with respect to its section, relative size of two geons, vertical position of a geon at the point of junction with another. The outcome is that with three geons one can theoretically describe over 1.4 billions distinct objects).

Also in this case, however, there are some intuitive problems. For instance, it is unfair to represent a table dish as a very wide thin cylinder a flat geon. This already tells us something important: the fact that a certain object *can be* represented as a normalized cylinder does not imply that it actually *is* represented that way. Furthermore and this is a more substantial problem it seems awkward to say that a doughnut (O-shaped object) consists of two joined handles (C-shaped cylinders), or of one elongated handle whose extremities are in touch. But of course the RBC theory does not want to expand its primitives by adding doughnuts. Otherwise bi-toruses, i.e., doughnuts with two holes (8-shaped objects), should also be assumed as a primitive. This seems necessary insofar as there seems to be no principled way within the putative RBC+torus theory to decompose a bitorus: as torus plus handle (C-shaped geon), or as handle plus torus? Since the same puzzle arises also for a tritorus, and for any arbitrary *n*-torus, it therefore seems that by this pattern one would have to introduce an infinite amount of primitives.

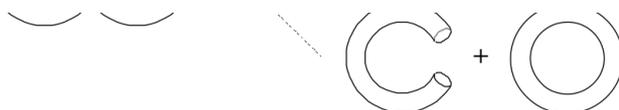


Figure 4. Biederman's problem: how do we decompose doughnuts into normalized cylinders?

One possible solution could come from Hoffman and Richards's theory of parts [1985]: analyze a torus as consisting of two (non-standard) parts a disc including the torus, and a *negative part* of such a disk (the hole). This is a solution inasmuch as both the disc and the negative part can be treated as RBC-normalized cylinders. The notion of negative part can be defined in relation to the normalization of the solid (positive) body hosting it: the closest solid for a torus is a cylinder; the negative part is the 'missing' cylinder in the middle. (The solution is obviously generalizable to arbitrary *n*-toruses.)

Note that the negative part here would correspond exactly to what we would treat as a hole. But Hoffman and Richards's theory and our theory in *Holes* [Casati & Varzi 1994] are false friends they are not notational variants. On our theory, a hole is not a *part* of its host (if you join the tips of your thumb and your index so as to form an 'O', you do not thereby create a *new part* of yourself, however negatively you look at it.) It just seems false that a torus is really two things: a disc plus a negative part. Moreover, Hoffman and Richards's theory suffers from another difficulty, which may



Figure 5. Biederman's inverse problem: should a semi-sphere be decomposed into a sphere with a negative half?

## 6. Herskovits' fly

Halfway between Gotts' problem and Biederman's inverse problem is the fly problem, widely discussed in the literature on spatial reasoning and on the semantics of spatial prepositions. Several authors (among whom Randell *et al*, 1992a, 1992b.) have suggested explaining the meaning of a preposition such as 'in' in terms of mereological inclusion in the convex hull of the containing object. As already pointed out by Annette Herskovits [1986], this approach fails to appreciate the crucial role of containing parts as opposed to other non-convex parts (a fly near the stem of a glass is not *in* the glass, though it may well fall within its convex hull; see figure 6, center). Nor could the problem be overcome by focusing exclusively on the convex hull of the object's containing parts, as suggested for instance by Vandeloise [1986] (Vandeloise [1994] defends a thoroughly functional approach). Apart from the apparent circularity, it is not difficult to find counterexamples insofar as the outer boundaries of containing parts may themselves involve concavities (figure 6, right; example from Vieu [1991: 207]).



Figure 6. The fly problem: reference to the convex hull (dashed line) is of little use for the purpose of telling the flies inside the glass from those outside.

This problem is similar to Gotts's inasmuch as the relevant role of what really counts as a container (a 'fillable' morphological discontinuity) cannot be explained in topological terms even if we extend the range of application of connection to the convex hull. Morphology runs afoul of pure topology (or mereotopology). The problem is also similar to Hoffman's, for it requires thinking about the complement of the object. By contrast, if we reason directly in terms of holes we get a radically different picture. Only the region corresponding to the hole the one on the top, not the 'groove' surrounding the stem or the top part of the glass in the right diagram can reasonably be treated as the container. And to be contained *in* the glass is to occupy (perhaps partially) that region, i.e., to fill (maybe partly) the hole. (See Varzi [1993] for a development along these lines.)

## 7. Concluding remarks

points, but even material objects have been the subject of philosophical dispute. Moreover, the methodology one is to adopt requires a careful examination of the consequences and is often a symptom of wider philosophical problems. For instance, one might choose to parsimoniously concentrate on objects and on their intrinsic properties, at the expense of the environment and of the relational ties of the objects to their environment, and thus to neglect complementary reasoning or more generally holistic components in spatial reasoning. On the other hand, one can attend to global properties of spatial situations, and fail thus to isolate relevant features of individual objects. This conflict (among others) could be seen as the reflection of a deeper philosophical conflict between spatial absolutism and spatial relationism.

Building models for knowledge representation has traditionally been seen as a top-down methodology in comparison to bottom up psychological generalizations. Nevertheless, this discrepancy between psychological, (a)-type theories, and formal, (b)-type theories does not exhaust the matter. Within both (a)-type and (b)-type theories a top-down strategy whose main constraints, such as representational economy, are formal should meet bottom-up requirements on the plausibility and well-foundedness of the theory.

## References

- Allen J. F., 1981, 'An Interval-Based Representation of Temporal Knowledge', *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, Vancouver: IJCAI [Morgan Kaufmann], Vol. 1, pp. 221-26.
- Aurnague M., Vieu L., 1993a, 'A Three-Level Approach to the Semantics of Space', in C. Z. Wibbelt (ed.), *The Semantics of Prepositions: From Mental Processing to Natural Language Processing*, Berlin: Mouton de Gruyter, pp. 393-439.
- 1993b, 'Toward a Formal Representation of Space in Language: A Commonsense Reasoning Approach', in F. Anger, H. Guesgen, and J. van Benthem (eds.), *Proceedings of the Workshop on Spatial and Temporal Reasoning. 13th International Joint Conference on Artificial Intelligence*, Chambery: IJCAI, pp. 123-58.
- Biederman I., 1987, 'Recognition-by-Components. A Theory of Human Image Understanding', *Psychological Review*, 94, 115-47.
- 1990, 'Higher-Level Vision', D. N. Osherson, S. M. Kosslyn, and J. M. Hollerbach (eds.), *An Invitation to Cognitive Science. Volume 2: Visual Cognition and Action*, Cambridge, MA, and London: MIT Press, pp. 1-36.
- Casati R. and Varzi A. C., 1994, *Holes and Other Superficialities*, Cambridge, MA, and London: MIT Press (Bradford Books).
- Clarke B. L., 1981, 'A Calculus of Individuals Based on 'Connection'', *Notre Dame Journal of Formal Logic* 22, 204-18.
- 1985, 'Individuals and Points', *Notre Dame Journal of Formal Logic* 26, 61-75.
- Gotts N. M., 1994a, 'How Far Can We 'C'? Defining a 'Doughnut' Using Connection Alone', in J. Doyle, E. Sandewall, and P. Torasso (eds.), *Principles of Knowledge Representation and Reasoning: Proceedings of the Fourth International Conference*, San Mateo, CA: Morgan Kaufmann, pp. 246-57.
- 1994b, 'Defining a 'Doughnut' Made Difficult', in C. Eschenbach, C. Habel, and B. Smith (eds.), *Topological Foundations of Cognitive Science. Papers from the Workshop at the First International Summer*

- Reasoning. Proceedings of the First International Conference*, Los Altos, CA: Morgan Kaufmann, pp. 357-68.
- 1992, 'Exploiting Lattices in a Theory of Space and Time', *Computers and Mathematics with Applications* 23, 459-76.
- Randell D. A., Cui Z., Cohn A. G., 1992a, 'An Interval Logic of Space Based on 'Connection'', in B. Neumann (ed.), *Proceedings of the 10th European Conference on Artificial Intelligence*, Chichester: John Wiley & Sons, pp. 394-98.
- 1992b, 'A Spatial Logic Based on Regions and Connection', in B. Nebel, C. Rich and W. Swartout (eds.), *Principles of Knowledge Representation and Reasoning. Proceedings of the Third International Conference*, Los Altos, CA: Morgan Kaufmann, pp. 165-76.
- Vandeloise C., 1986, *L'espace en français: sémantique des prépositions spatiales*, Paris: Seuil.
- 1994, 'Methodology and analyses of the preposition *in*', *Cognitive Linguistics*, 5-2, 157-84.
- Varzi A. C., 1993, 'Spatial Reasoning in a Holey World: A Sketch', in F. Anger, H. Guesgen, and J. van Benthem (eds.), *Proceedings of the Workshop on Spatial and Temporal Reasoning. 13th International Joint Conference on Artificial Intelligence*, Chambéry: IJCAI, pp. 47-59.
- 1994, 'On the Boundary Between Mereology and Topology', in R. Casati, B. Smith, and G. White (eds.), *Philosophy and the Cognitive Sciences. Proceedings of the 16th International Wittgenstein Symposium*, Vienna: Hölder-Pichler-Tempsky, pp. 423-442.
- Vieu L., 1991, *Sémantique des relations spatiales et inférences spatio-temporelles: Une contribution à l'étude des structures formelles de l'espace en Langage Naturel*, Université Paul Sabatier de Toulouse: PhD Thesis.
- Whitehead A. N., 1919, *An Enquiry Concerning the Principles of Human Knowledge*, Cambridge: Cambridge University Press.
- 1920, *The Concept of Nature*, Cambridge: Cambridge University Press.
- 1929, *Process and Reality. An Essay in Cosmology*, Macmillan.