(In Oliviero Stock (ed.), *Spatial and Temporal Reasoning*, Dordrecht: Kluwer, 1997, pp. 73-96.)

Spatial Entities

Roberto Casati
CNRS, Seminaire d'Epistémologie Comparative, Aix-en-Provence (France)

Achille C. Varzi
Department of Philosophy, Columbia University, New York (USA)

1. INTRODUCTION

Common-sense reasoning about space is, first and foremost, reasoning about things *located in* space. The fly is inside the glass; hence the glass is not inside the fly. The book is on the table; hence the table is under the book. Sometimes we may be talking about things *going on* in certain places: the concert took place in the garden; then dinner was served in the solarium. Even when we talk about "naked" (empty) regions of space—regions that are not occupied by any macroscopic object and where nothing noticeable seems to be going on—we tipically do so because we are planning to move things around, or because we are thinking that certain actions or events did or should take place in certain sites as opposed to others. The sofa should go right here; the aircraft crashed right there. Spatial reasoning, whether actual or hypothetical, is typically reasoning about spatial entities of some sort.

One might—and some people do—take this as a fundamental claim, meaning that spatial entities such as objects or events are fundamentally (cognitively, or perhaps even metaphysically) prior to space: there is no way to identify a region of space except by reference to what is or could be located or take place at that region. (This was, for instance, the gist of Leibniz' contention against the Newtonian view that space is an individual entity in its own right, independently of whatever entities may inhabit it.) It is, however, even more interesting to see how far we can go in our understanding of spatial reasoning without taking issue on such matters. Let us acknowledge the *fact* that space has little use *per se* in the ordinary representation of our environment (that is, the representation implicit in our everyday interaction with the spatial environment). What is the meaning of that for the theory of spatial reasoning? How does that affect our construction of a general model of our spatial competence?

These questions have both methodological and substantial sides. On the substantial side, they call for a clarification of the relevant ontological presuppositions. A good theory of spatial reasoning must be combined with (if not grounded on) an account of the sort of entities that may enter into the scope of the theory, an account of the sort of entities that can be located or take place in space—in short, an account of what may be collected under the rubric of *spatial entities* (as opposed to *purely spatial items* such as points, lines, regions). What are they? What exactly is their relation to space, and how are they related to one another? In short, what *special* features make them *spatial* entities?

On the methodological side, the issue is the definition of the basic tools required by a good theory of spatial representation and reasoning, understood as a theory of the representation of and reasoning about these entities. In fact there may be some ambiguity here, as there is some ambiguity as to how "reasoning" and "representation" should really be understood in this regard. We may think of (i) a theory of *the way* a cognitive system represents its spatial environment (this representation serving the twofold purpose of organizing perceptual inputs and synthesizing behavioral outputs); or (ii) a theory of the spatial *layout* of the environment (this layout being presupposed, if not explicitly referred to, by such typical inferences as those mentioned above). The two notions are clearly distinct; and although a comprehensive theory of space should eventually provide a framework for dealing with their mutual interconnections, one can presumably go a long way in the development of a theory of type (i) without developing a theory of type (ii), and vice versa. On the other hand, both notions share a common concern; both theories require an account of the geometric representation of our spatial competence before we can even start looking at the mechanisms underlying our actual performances. (This is obvious for option (ii). For option (i) this is particularly true if we work within a symbolic paradigm, i.e., if we favor some sort of mental logic over mental models of reasoning. For then the specificity of a spatial theory of type (i) is fundamentally constrained—if not determined—by the structure of the domain.) It is this common concern that we have in mind here. What are the basic tools required by a theory of this competence? How, for instance, should we deal with the interplay between truly spatial concepts—such as "contained in", or "located between"—and purely mereological (part-whole) notions? What are the underlying principles? And how do they relate to other important tools of spatial representation, such as topology, morphology, kynematics, or dynamics? Moreover, do the answers to these questions hold for all sorts of spatial entities? Or is there a difference between, say, material objects and events? Why for instance do spatial boundaries seem to play a crucial role for the former but not for the latter? Are there spatial entities whose spatial location is more than a contingent fact?

These and many other questions arise forcefully as soon as we acknowledge the legitimacy of the more substantial issues mentioned above. Our contention is that the shape of the theory of space depends dramatically on the answers one gives. Over the last few years there has been considerable progress in the direction of sophisticated theories both of type (i) and of type (ii), particularly under the impact of AI projects involved in the construction of machines capable of autonomous interaction with the environment. We think at this point there is some need for a philosophical pause, so to speak. Our purpose in what follows is to offer some thoughts which may help to fulfil (albeit very partially and asystematically) this need.

PARTS AND WHOLES

Much recent work on spatial representation has focused on mereological and topological concepts, and the question of the interaction between these two domains will be our main concern here.

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. How is this relation to be characterized? How are the spatial parts of an object spatially related to one another? Traditionally mereology has been associated with a nominalistic stand, and 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 the philosophical position of nominalism. 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 (including sets, if one will, as in Lewis [1991]). This certainly fits in well with the spirit of type-(ii) theories, in the terminology of the previous section, but type-(i) theories may also be seen this way. So mereology is ontologically neutral. The question is, rather, how far we can go with it—how much of the universe can be grasped and described by means of purely mereological notions.

We are going to argue that one cannot go very far. In our view (and this is a view we share with others, though we may disagree on how to implement it), a purely mereological outlook is too restrictive unless one integrates it at least with concepts and principles of a topological nature. There are several reason for this, in fact, and some of them will keep us occupied for quite a while in the second part of this paper. However one basic motivation seems easily available. Without going into much detail (see Varzi [1994]), the point 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 (in its ordinary understanding) 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 Caesar's left foot) by reasoning exclusively in terms of parthood. If we allow for the possibility of scattered entities, then we lose the possibility of discriminating them from integral, connected wholes. On the other hand, we cannot just keep the latter without some means of discriminating them from the former.

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 parts that are "joined" to each other. This relation is defined thus:

(1) $\mathsf{J}(x,y) =_{\mathrm{df}} \;\; z(\mathsf{O}(z,x) \;\; \mathsf{O}(z,y) \;\;\; w(\mathsf{P}(w,z) \rightarrow \mathsf{O}(w,x) \;\; \mathsf{O}(w,y)))$

where ‘P’ indicates parthood and ‘O’ overlapping. This should rule out scattered wholes. But it is easily verified that this definition falls short of its task *unless it is already assumed* that the piece z overlaying two “joined” events x and y be itself connected. In other words, the account works if the general assumption is made that only self-connected entities can inhabit the domain of discourse. But this is no account, for it just is not possible to make the assumption explicit.

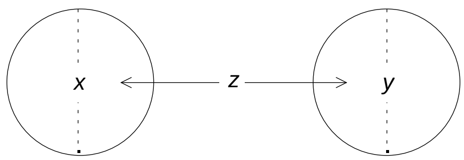

Figure 1. Whitehead’s problem: x and y are not connected unless the overlaying piece z is itself assumed to be (self-)connected.

THE TOPOLOGICAL OPTION

Mereology can hardly serve 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; mereologically 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. These and similar notions are arguably fundamental for spatial reasoning (for type (i) theories as well as for type (ii) theories). Yet they run afoul of plain part-whole relations, and their systematic account seems to require an explicit topological machinery of some sort.

Now, in recent philosophical and 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—so goes the argument—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 over 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, as reported in Cohn's contribution to this volume.)

The subsumption of mereology to topology proposed on this approach 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) $\mathsf{P}(x, y) =_{\mathrm{df}} \quad z(\mathsf{C}(z, x) \rightarrow \mathsf{C}(z, y))$.

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: whatever overlaps a part overlaps the whole. But things may change radically on different readings.

Typically, the suggestion is to interpret the relation '$\mathsf{C}(x,y)$' as meaning that the *regions* x and y have at least one *point* in common. This means two things. First, the domain of quantification is viewed as consisting of spatial (or 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 complications 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 (see Varzi [1996a]). 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 various patterns of topological reasoning. For instance, self-connectedness is immediately defined:

(3) $\mathsf{SC}(x) =_{\mathrm{df}} \quad y \quad z(\quad w(\mathsf{O}(w, x) \leftrightarrow \mathsf{O}(w, y) \quad \mathsf{O}(w, z)) \rightarrow \mathsf{C}(y, z))$.

So if spatio-temporal regions are the only entities of our domain, then (2) yields important conceptual achievements: the basic limits of mereology are overcome. However, there is a second side of the coin. For if we really are to take an open-faced attitude towards real world things and events, as we urged in the Introduction (without identifying sych entities 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: holes are immaterial, and can therefore share their location with other entities. For another example (from Davidson [1969]), the rotation and the getting warm of a metal ball that is simultaneously rotating and getting warm are two distinct event. Yet they occupy exactly the same spatio-temporal region; because events, unlike material objects, do not *occupy* the space at which they are located. 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 the 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. Compare Casati [1995] and Pianesi & Varzi [1994, 1996a, 1996b].)

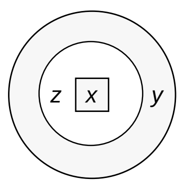

Figure 2. Clarke's problem: *x* is inside object *y*; but what is the relationship between object *x* and hole *z*? (and what the relationship between *z* and *y*?)

THE HOLE TROUBLE

These concerns may not be definitive. In particular, our examples of the shadow and of the stone in the hole presuppose a friendly attidude towards the ontological status of shadows and holes, which is far from unproblematic. A non-realist would simply say that such “things” do not exist (for shadows and holes are—after all—paradigm examples of *nothings*)—hence the above question would not even arise. However this would require some radical eliminative strategy. It would, for instance, require some systematic way of paraphrasing every shadow- or hole-committing sentence by means of sentences that do not refer to or quantify over shadows or holes (the cheese is holed, but there *is* no hole in it; the wall is darkened, but there *is* no shadow). On the other hand, if we want to take common sense seriously, we should resist these ways out in favor of a realist, common-sense attitude. Holes and shadows are enigmatic. Yet, if there is an ontology inherent in our everyday reasoning about the world, then this ontology comprises shadows and holes (and cognate entities such as waves, knots, cuts, grooves, cracks, fissures, smiles, grims) along with stones and chunks of cheese.

We have defended this view at large in previous work (especially in *Holes* [1994]), and we refer to it for further discussion of the underlying philosophical issues. In fact, we take this to be a good example of the sort of general ontological concern that we mentioned at the beginning: a general theory of spatial representation calls for a clarification of the relevant ontological presuppositions. It must be combined with (if not grounded on) an explicit account of the sort of entities that may enter into the scope of the theory. Be it as it may, it is apparent that the simplification introduced by (2) has critical consequences if our concern is with the foundations of general-purpose representation systems, *even if* we take a non-realist attitude toward holes and shadows. For the basic issue of the relationship between an entity and “its” space (the space where it is located) is then trivialized: every entity is *reduced* to its space. 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 mantain, is not only a source of conceptual poverty; it may also be misleading.

Let us focus on holes. Some recent work by Nick Gotts [1994a, 1994b] is indicative of the difficulties 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 go as far as to describe toroidal structures (doughnuts). Hence we can in principle describe perforated objects without directly resorting to holes. (Just focus on the doughnut, and ignore the hole, as it were.) This is indeed remarkable, but closer inspection show that the results are necessarily partial.

There are two troubles. The first is that the notion of a torus is only capable of capturing one type of hole, viz. perforations ("tunnels", as we call them). It remains thoroughly blind in front of 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 an objection to Gotts; rather, it merely points out 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 and doughnuts (and little else) without reflecting on the ontology. The second trouble is more specific. For as it turns out, an account in terms of 'C' is intrinsecally 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 a theory can't tell the *difference*, its classificatory power is deficient in an important sense, at least from our present perspective.

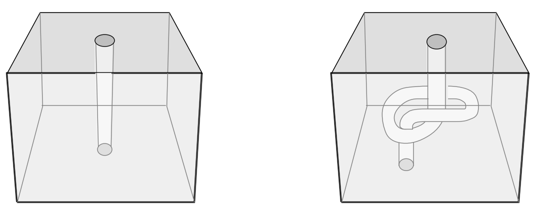

Figure 3. We can't 'C' well enough to tell a straight hole from a knotted one.

THE COMPOSITIONAL APPROACH

We take the foregoing to imply a threefold moral. Firstly, it appears that one needs both mereology and topology as *independent* (though mutually related) frameworks. Mereology alone is too weak; topology alone is too strong (in a sense) but at the same time very limited (in another sense). Second, the limitations of topology are significant even at a fairly elementary level (a long way before functional features become important for classifying shapes or providing an analysis of such relations as containment; as urged e.g. by Vandeloise [1994] and Aurnague & Vieu [1993]). Third, and perhaps more importantly, 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 the direction dictated by the last *desideratum* 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 a normailzed 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 Binford; Biederman should nevertheless be given credit for formulating it in purely qualitative terms, without resorting to sophisticated abstract hyerarchies). The related cognitive thesis is that the human shape recognition system is based on the capacity to decompose an object into cylinders. Thus, for instance, a coffee cup would consist of a main semi-concave cylinder (the containing part) with a small bended cylinder (the handle) attached to the first at both ends. (In a more recent formulation [1990], both geons and relations 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, some problems arise immediately. For one thing, the theory is based on a general assumption pertaining to the cognitive dimension of part-whole reasoning which seems false. RBC is meant to do justice to the intuition that the mereological module is crucial to object recognition. However, recent data by Cave and Kosslyn [1993] show very clearly that a module for decomposition into parts does *not* act prior to, and is not a necessary condition of, object recognition. Their results indicate first of all that the recognition of an object depends crucially on the proper spatial relations among the parts: when the parts are scrambled or otherwise scattered, naming times and error rates increase. Secondly, Cave and Kosslyn's results show that the mereological parsing of an object affects the object's identificatiton "only under the most impoverished viewing conditions". This is not a disproof of the existence of a merelogical module *per se* (for instance, the way objects are partitioned tends to be rather robust across individuals). However, Cave and Kosslyn contend that the module need not be activated for the purpose of object recognition, and their results leave little room for a rebuttal. (We tend to rely on data of this sort, because they dispose of the issue of object recognition in our discussion. In particular, the structure of the as yet putative mereological module should be considered independent of the pressures of object recognition.)

A second problem is more technical and, in a sense, farther reaching. Take a flat object, say a disc. In spite of the "generative" power of the notion of a normalized cylinder, it would seem that in cases like this its representational adequacy is at the limit: it seems unfair to represent a disc as a wide, short *cylinder*—a flat geon. It might be replied that this is an objection only if our concern is with type-(i) theories, with *the way* a cognitive system represents its spatial environment. (Surely the fact that a certain object *can be* represented as a normalized cylinder does not imply that it actualy *is* represented that way by a cognitive system.) But if we are looking for a purely geometric theory of type (ii), one could argue that this sort of artificiality is inessential. After all, for the purpose of spatial *reasoning*, it does not matter what we take a disc to be: the important thing is to keep the number of primitives to a minimum. If so, however, consider then a disc with a hole, or a doughnut for that matter. How is such an object to be represented? Here the problem is twofold. On the one hand, we would again say that it is awkward to regard a disc or doughnut (an **O**-shaped object) as consisting of two joined handles (**C**-shaped cylinders), or perhaps of a single elongated handle whose extremities are in touch. This is the type (i) misgiving. But there is also a type (ii) misgiving. For how do we choose between the possible decompositions? More generally, how do we go about decomposing an object with holes in terms of its non-holed parts—is there any principled way of doing that? There isn't. And of course we wouldn't want to expand our primitives by adding doughnuts. Otherwise bitoruses, i.e., doughnuts with two holes (**8**-shaped objects), should also be assumed as primitives. That would be 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 more generally for arbitrary *n*-toruses, it therefore seems that by this pattern one would have to introduce an infinite amount of primitives—and *that* is unacceptable also from the perspective of a type-(ii) theory.

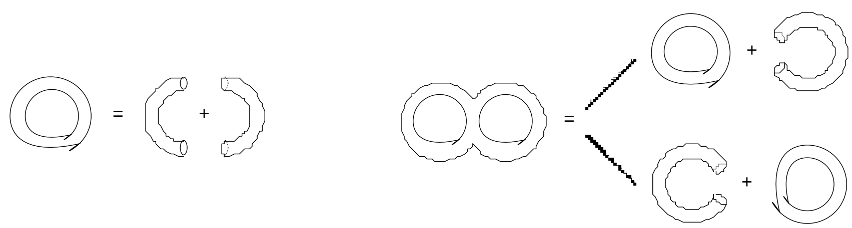

Figure 4. How do we decompose a doughnut into normalized cylinders? How do we decompose a double doughnut?

NEGATIVE PARTS

Once again, the problem is that the theory under consideration aims to account for one *desideratum*, but neglects the others. We welcome the suggestion of investigating a spatial compositional structure which is not simply a mereology of *space*, but of spatial entities. This is as it should be. But we already saw that a pure mereological module is not going to do all the work; *a fortiori*, one can't go very far by reasoning exclusively in terms of such well-behaved parts as geons.

However, here one might be tempted to reconsider our earlier conclusion: perhaps *that* is precisely the problem; perhaps the problem is precisely that the relevant notion of part (or component) is not broad enough to do all the work. With a broader notion—not only broader than geons, but at this point even broader than the notion of part underlying standard parthood theories—mereology might be enough after all.

One way of implementing this intuition could come from Hoffman and Richards's original theory of parts [1985]. Analyze a doughnut as consisting of two parts—not just two ordinary parts, but two "complementary" parts (as it were): a *positive* part (in the shape of a disc) and a *negative* part (the hole). This would be a solution to the above problem inasmuch as both the disc and the negative part can be treated as RBC-normalized cylinders. And 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 doughnut is a cylinder; the negative part is the "missing" cylinder in the middle. (The solution is obviously generalizable to arbitrary *n*-toruses.)

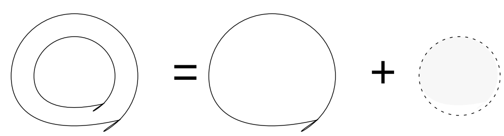

Figure 5. Holes as negative parts.

This proposal has some independent attractive features, which should not be overlooked. One is that it deals neatly with complementary or dual structures, such as those constituted by grooves (or notches, dents, indentations) and ridges. A groove is a negative, intruding part just like a ridge is an ordinary, protruding part. There is a rather obvious reason for the desire to treat such dual structures on a par. A natural way to produce an indent in a body is to act on it with another body's protrusion; conversely, a natural way to produce a protrusion is to fuse some material in anothers body's indent. We can immediately predict, by observing the processes of fusion and of indentation, that the shapes of the notch and of the protrusion will fit perfectly. (Are they actually one single shape? This is an interesting question, pertaining to the more general issue of the status of complementary shapes.) Another advantage of negative parts is that when it comes to holes the notions of *completion* of an object, or of a hole's being a lack in the object (a missing something), are immediately and rather nicely implemented. For holes are exactly there where some part of the object could conceivably have been; and as they *always* are where some part of the object could have been, it does not make a big difference if we have them coincide with some *actual* negative part of the object.

Let us stress indeed that these negative parts would exactly correspond to what we would treat as a hole. Wherever we have a hole, this theory would have a negative part, and viceversa. But the negative-part theory and our theory in *Holes* (the "hole-theory", hereafter) are false friends—they are not just notational variants. On the hole-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. On the negative-part theory, by contrast, a hole is precisely that: it is just a part, albeit of a somewhat special and hiterto neglected sort.

Now, this may well be a disadvantage of the hole-theory. It requires a special primitive 'H' ('… is a hole in …') logically distinct from the parthood primitive 'P'. The two primitives are not only distinct; they stand for two relations that are totally disjoint: as we said, holes are never parts of their hosts. They don't even overlap, as reflected in the axiom:

(4) $\mathsf{H}(x, y) \rightarrow \neg \mathsf{P}(x, y)$

By contrast, if holes are treated as parts (albeit parts of a special kind) the possibility is left open that the 'P' primitive (perhaps combined with a suitable inversion functor) be sufficient for most purposes. The difference is important, of course. Conceptual economy may be very advantageous, especially from the perspective of a type-(ii) theory. (Ironically, this is not the perspective of Hoffman and Richards. But think of an expert system whose task is, say, to classify shapes. One may imagine using, in addition to several shape primitives, the 'P' primitive and an inversion functor which maps (suitable) positive parts of an object's complement onto corresponding negative parts of the object itself, and viceversa—an idea that can be traced as far back as to the theory of Franz Reuleaux [1875].)

At the same time, the price of this conceptual economy may be too high. On the hole-theory a doughnut is just a doughnut—an object *with a hole.* On the negative-part theory a doughnut is really the *sum* of two things: a disc *plus* a negative part. Is *that* what a doughnut is? More importantly, what does it mean to represent a doughnut that way? What kind of mereology is required? And when are we allowed to speak of negative parts any-way? Are holes (and grooves, notches, etc.) the only sort of negative part? Take a sphere and cut it in half. According to one intuition, the closest approximation for each piece is the sphere itself; each piece has a missing part. Yet surely it would be absurd to treat a semi-sphere as a *whole sphere plus a negative part* (the missing half). How can the negative-part theory rule that out? Perhaps in this specific case the difficulty might be dealt with simply by stipulating that objects must be approximated to their convex hull. But that has the force of an *ad hoc* solution. A champagne glass would by that pattern involve two large negative parts, one surrounding the stem, and one in the wine cup. It is hard to find satisfaction in that picture.

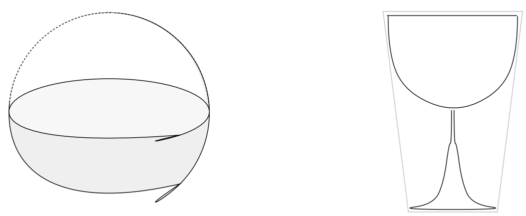

Figure 5. Problems with negative parts: is a semi-sphere composed of a whole sphere plus a negative half? What negative parts does a champagne glass consist of?

HYBRID SUMS

Let us look at these questions more closely. If we are right, the answer will eventually be that no mereological module could function *reasonably* if it had to operate on negative parts. And this will be relevant to both a type-(i) and a type-(ii) perspective on spatial reasoning and representation.

Consider how the mereological module can operate on holes if these are construed not as negative parts but as immaterial individuals which are *not* part of their material hosts. According to the hole-theory, anytime there is a hole in an object there is some mereological composition around. Not only because the theory implies that atoms are holeless. (A hole is always in some proper parts of an object, therefore, if an object has a hole, it must have parts). There is also the trivial fact that a hole is a part of the mereological sum of the host and of the hole itself.

Now define a (cognitively) natural object as an object which is taken by the cognitive system as unitary (typically, a cognitively natural object is a unit for counting). As we said, not all mereological sums are natural objects. Think again of the sum consisting of our four eyes and Caesar's left foot. But even so, a large number of mereological sums are unitary—they have that cosy, peculiar naturalness and wholeness. Our question now is whether the mereological sum of an object and its hole(s) is a cognitively natural object, and, if it is, how it is related to a normal and holeless natural object. For topological connection plays a hand in the game, but exactly which hand it is unclear.

Mind the fact that such a sum—call it *s*—is not a mereological sum of the most obvious kind. It is a sum of two objects, one of which (the hole) *depends* existentially upon the other (you cannot remove the hole from the doughnut), which in turn depends geometrically or conceptually upon the former (you cannot have a doughnut without a hole). And these dependences are more than mereological: they involve a form of topological dependence too. The sum *s* is not decomposable into hole *h* and host *o* in the same sense in which the sum of two solid objects, *a* plus *b*, is decomposable into *a* and *b*. For the hole exists only insofar as it is topologically connected with its host. And if you *eliminate* the hole (e.g., by elastically deforming the host or, if the host is a doughnut, by cutting it open), the host is no longer holed. Thus, even if the sum *s* of a hole *h* together with its host *o* is indeeed a sum of a hole and a *holed* object, metaphysically it has rather peculiar features. For instance, it does not behave as the ordinary sum of a holed object and of its perfect filler (imagining the hole to be filled).

Observe now that the hole-theory allows us to express these facts by making a distinction between a hole's being *in* something and a hole's being *part of* something. The hole, *h*, is part of the hybrid sum *s* (hole + host) but it is not a hole in *s*. For *h* overlaps *s*, and by axiom (4) no hole overlaps its host. Moreover, the following principle of left-monotonicity hold:

$$(5) \qquad \mathsf{H}(x, y) \quad \mathsf{P}(y, z) \rightarrow \mathsf{O}(x, z) \quad \mathsf{H}(x, z)$$

That is, any object that includes the host of a hole is a host of that hole, unless its parts also include parts of that very hole. You can produce a holed object by taking just another holed object and by attaching a part thereto (in an appropriate way).

Connectedness (between a hole and its host) is thus mandatory for binding the salient parts of *s*. But one must add that the bind between a hole and its host is much stronger—topologically, not only ontologically—than the bind between two ordinary solid objects in touch with each other. (The bind resembles the one between the various non-salient, only-potential parts of an homogeneous, self-connected chunk of an ideal stuff. Any part of the chunk cannot but ideally be detached from the object—real detachment produces *two* new objects.) We do not see, in any case, how these facts should be related to the property of being a natural whole. We do not see any reason why the sum of an object and its holes should be a natural object. On the other hand, it seems that the negative-part theory requires that such a sum *be* a natural object, for this is where the cognitive system should start from when it comes to holed objects. The sum is assumed to be *cognitively prior* to the analysis into object + holes. But this is far from obvious. We are prepared to accept that in some cases a holed object is considered a sort of incomplete object (a statue with a perforation, say), but this is not the rule. And the proof is, quite simply, that in so many cases we would not be able to tell what parts are missing from what object.

NEGATIVE PARTS OF *WHAT*?

We thus come to what seems to be the major problem of an ontology of negative parts. We have a number of characters here. To begin with, there is an ordinary holed object; call it the *solid object, o*. Add to this its *hole* (or holes), *h*. The hole is not part of the solid object (which is impenetrable, and thus cannot have penetrable parts). It is nevertheless part of the sum of the solid object with the hole itself; call this sum the *holed sum*, *s*. Thus, $s = h + o$ and $o = s - h$ where '+' and '–' can be defined in the usual way:

(6) $x+y =_{df} \iota z \quad w(\mathrm{O}(w,z) \leftrightarrow (\mathrm{O}(w,x) \quad \mathrm{O}(w,y)))$

(7) $x-y =_{df} \iota z \quad w(\mathrm{P}(w,z) \leftrightarrow (\mathrm{P}(w,x) \quad \neg\mathrm{O}(w,y)))$

Then there are two relevant mereological complements: the complement of the solid object, o, and the complement of the holed sum, s, where in general we define

(8) $x =_{df} \iota z \quad w(\mathrm{O}(w,z) \leftrightarrow \neg\,\mathrm{O}(w,x))$.

The hole is part of the complement of the solid object (which for the sake of simmetry may be called the *complement sum*), but it is not part of the complement of the holed sum. For the complement sum and the holed sum mereologically overlap: they share a negative part, viz. the hole. In short, $s = (\ \ o) - h$ and $o = (\ \ s) + h$,

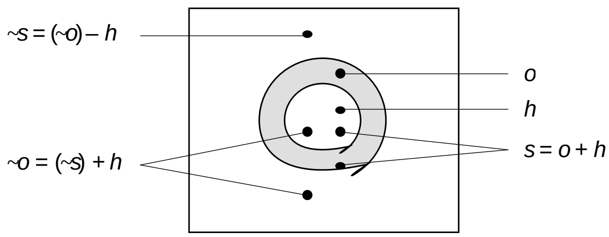

Figure 6. Doughnut, hole, sums, and complements.

It is enough to formulate these distinctions to see a problem emerge. A pure mereological module founders because a negative part is at the same time part of the holed sum and of the complement sum. Now the theory of negative parts does not founder because of that. On that theory, the hole is not a negative part of the holed sum: it is a positive part thereof, and a negative part of something else. But of what, exactly?

More characters must be added to the picture. First, in the negative-part theory we have this entity (partly solid and partly immaterial) which has h as a negative part. Call that entity o'. This is not to be confused with the holed sum, s, because h is a standard part of s, not a negative part. Nor is o' to be confused with o, because h is not a part of o, whether positive or negative. Rather, o' is a third object, distinct both from s and from o. It is what, on the negative-part theory, the doughnut really is. Furthermore, the result of subtracting h from o' gives you an other object still, distinct from all of s, o, and o': it give you a disc, the disc we *would* have in case our doughnut were holeless. Call this last character s'; we then have the followig equations:

(9) $\quad s = o + h \qquad o = s - h \qquad s' = o' - h \qquad o' = s' + h.$

The first two of these (on the left) are standard mereological equations, corresponding to (6) and (7). The last two (on the right) are not. The way h is added to o to yield s is not the same way h is added to s' to obtain o'; for the former operation yields a bigger object (in an intuitive, compositional, non-metric sense) than the one we start with, whereas the latter—a form of negative sum—yields a smaller object. And the way *h* is subtracted from *s* to yield *o* is not the same way *h* is subtracted from *o'* to obtain *s'*; for the former operation yields a smaller object than the one we start with, whereas the latter—a form of negative difference—yields a bigger object. But is there any way to characterize the latter operations in terms of the former? Is there any way to characterize $\dotplus$ and $\dot{-}$ in terms of + and –? It seems not, *unless* negative parthood is assumed as a primitive next to parthood simpliciter. But if we do so, then we have two mereologies, not one; that is, we have two mereological primitives. And one seeming advantage of the negative-part theory over the hole-theory (conceptual simplicity) is lost.

Nor is this the whole story. Consider the complement operation, as defined in (8). How does it behave with respect to negative parts? We are not asking for the negative counterpart of the complement operation, which could presumably be defined as in (8) but using the negative counterpart of 'O'. We are asking how ' ' behaves when its arguments are among the additional characters envisioned by the negative-part theory. For instance, how is the following table to be completed?

(10) $s = (\ o) - h$ $\qquad$ $o = (\ s) + h$ $\qquad$ $s' = ?$ $\qquad$ $o' = ?$

Is *h* part of *s'*? It would seem so, for surely *h*, a hole, does not overlap *s'*, a solid disc; so it must be part of the complement. But then, what is the difference between *s'* and *o*? It can only be $s' - o$, i.e., the small *solid* disc in the middle of *s'* which is conceiled, so to speak, by the negative part of *o'*. However, that means that the list of entities at stake is still growing, giving rise to further questions. Call this new "invisible" part *d*. Is *d* also part of the complement of *o*? Of *s*? Of what entities? And what sort of entity is $d + h$? Finally, what sort of entity is $d \dotplus h$—the entity obtained from *d* by "adding" a perfectly congruent negative part? How does it differ from nothing at all?

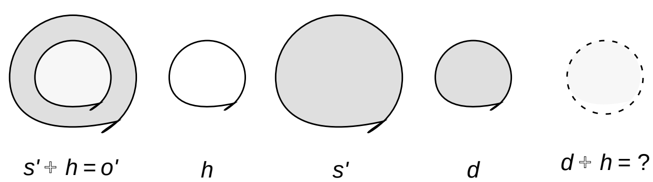

Figure 8. Spatial entities in negative-part mereology.

THE NEED FOR EXPLICIT THEORIES

The upshot of all this seems clear. If we don't take holes seriously, we end up with a theory which is formally just as rich, due to the need of two distinct mereological primitives, and ontologically much more dubious, due to the presence of such mysterious entities as s', d, $d + h$, and so on. Seen from another perspective, it is the notion of *complement* that founders conceptually. If the doughnut is really somewhat bigger than its edible part, if it also consists of a negative part, then *its* complement does not comprise the negative part. But if the mysterious negative parts are not parts of the complement—that is, if the negative parts of the doughnut are not parts of the doughnut's complement—then why are they negative? This intuition is not negotiable. And if the negativity does not lie in the complement, then why not allow for "negative" entities to begin with—why not allow for holes?

We have thus reached again a general conclusion concerning the interplay between ontology, mereology, and topology. And the conclusion is that we need all of them. We need mereology because topology is mereologically unsophisticated. We need topology because mereology is topologically blind. And we need ontology because both topology and mereology—even if we try to relax or supplement the relevant primitives—are intrinsically incapable of making sense of important ontological distinctions.

It now goes beyond the aims of the present work to give specific indications of how these three domains can actually be combined into a systematic theory. Some developments can be found in *Holes* as well as in Casati [1995a, 1996], Casati and Varzi [1996], Pianesi & Varzi [1994], and Varzi [1993, 1994, 1996a, 1996b]. By way of illustration, however, consider briefly how a mereotopological theory developed in this spirit can provide the foundations for some basic patterns of spatial reasoning of the sort mentioned at the beginning. The fly is inside the glass; hence the glass is not inside the fly. But under what conditions does a fly qualify as being *in* a glass? (Annette Herskovits gives a thorough analysis of this issue in her contribution to this volume, examining all the intricacies allowed by the use of prepositions in natural language. Here we only interested in the geometry of the problem, as it were—a much more modest task.)

Some authors have suggested that the answer could be given in terms of mereological inclusion in the convex hull of the containing object (Figure 9, left). But as already Herskovits [1986] pointed out, such an account-would fail 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 9, middle). Nor can the problem be overcome by focusing exclusively on the convex hull of the object's containing parts, as initially suggested 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 9, right; example from Vieu [1991: 207]).

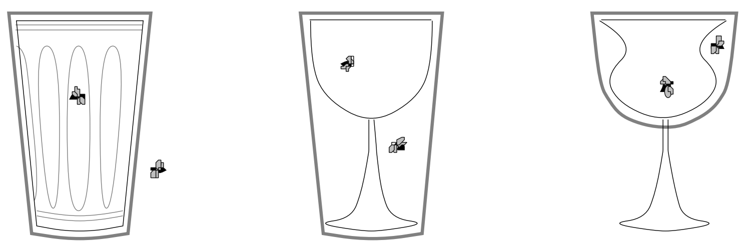

Figure 9. The fly problem: reference to the convex hull (dashed line) is of little use for the purpose of telling the flies inside the glass from those outside.

Now, this problem is halfway between what we called the hole problem (a problem for a purely topological theory, which we discussed earlier in connection with Gotts's work) and Hoffman and Richards' problem (one of the problems for a purely mereological theory). It is similar to the former 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. And the problem is similar to the latter (and more generally to the crucial dilemma of the negative-part theory) insofar as 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. Mereology and topology give us the structure of the entities at issue (the glass, the hole, the fly, the corresponding regions); and containment is explained in terms of simple inclusion relations between the region of the fly and the region of the hole. We need the fly and the hole to begin with. *Then* we look at their regions. (See Casati & Varzi [1996] for a closer examination of the structure of spatial location.)

Here the point is of course that the containing part of a glass determines a true hole—a hollow, in effect. No doubt there are other senses in which a glass can be said to be holed. However, *what* exactly counts as a hole or a containing part is not at issue: the account will be effective precisely insofar as the existence of independent criteria for holehood is presupposed—e.g., insofar as the space around the stem of a glass is not taken to be a hole. So if we have holes in the ontology (along with corresponding identity and individuation criteria), the problem dissolves; whereas the lack of holes gives rise to the difficulties illustrated in Figure 9.

This is not to suggest that mereology + topology + explicit hole commitment will give us a complete account of the notion of spatial containment, or even a full solution of the fly-in-the glass problem. In fact it is not difficult to find instances where ordinary intuitions are not adequately captured by the above suggestion. For example, the two patterns in Figure 10 provide apparent counterexamples whose solution seems to call for a decisive step into other territories than purely geometrical reasoning. Most likely these include at least some pragmatics (as suggested, e.g., by Aurnague & Vieu [1993] and Vandeloise [1994, 1996]), or causal factors at large. Even so, several useful refinements can be introduced already at the geometrical level, including some applications to naive-physical reasoning about containment. (See Varzi [1993] for details.)

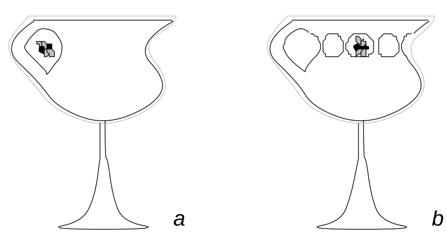

Figure 10. Further difficulties with the relation of containment: in both cases the fly is in a hole hosted by the glass, but not in the glass itself.

Moreover, the basic theory can be improved in various ways, by fully investigating the class of entities compatible with it. For instance, from a classificatory perspective the gist of the theory fragment presented in *Holes* is that holes come in three kinds: superficial hollows, perforating tunnels, and internal cavities (plus some significant mixed cases). The basic patterns are illustrated in Figure 11. There are in fact important distinctions among these three kinds of hole, i.e., more precisely, among the entities affected by such holes. And there is, correspondingly, a simple decision tree. This is illustrated in Figure 12 (left). But this basic taxonomy can then be extended. For example, Figure 12 (right) shows the result of including grooves. The idea is that grooves are a kind of hole, though geometrically rather bizarre. (We could say that a groove is a sort of “external” tunnel.)

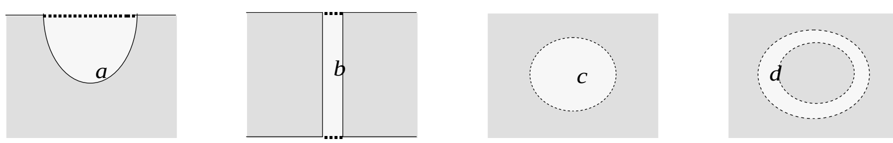

Figure 11. Holes come in superficial hollows (*a*), perforating tunnels (*b*), internal cavities (*c*), and some mixed cases, e.g. internal tunnel-cavities (*d*).

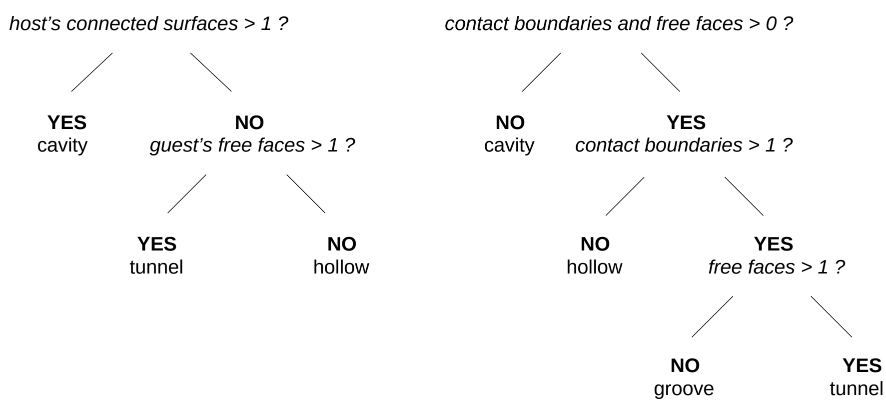

Figure 12. Basic and extended classification trees for holes. Here a free face is any (maximally connected) part of the surface of the hole’s filler (the ideal object that can be used to perfectly fill the hole) that is not connected with the host’s surface, and a contact boundary is any (maximally connected) boundary of such a free face.

In all of these cases, the same intuition is at work. To classify holes and hole-like entities, we need look at more than just the topology of their material host: we also need to look at their potential “guests”, so to speak. For holes are *fillable* entities, and much of our reasoning about holes and hole-like entities involves reasoning about how one can fill them. (In our view this is the only way to avoid two serious classificatory deficiencies of pure topology. One is the fact, already mentioned in connection with Gotts’s work, that topology is blind against non-perforationg holes. The second is that it is also too strong, for it treats as equivalent holed objects that are clearly different from the standpoint of common sense, and which should be kept distinct for most purposes: see Figure 13. By focusing on the patterns of interaction between host and guests (fillers), the hole-theory aims at finer distinctions, as illustrated in Figure 14: the morphological complexity of a hole is reflected in the topological complexity of the host-guest contact surface.)

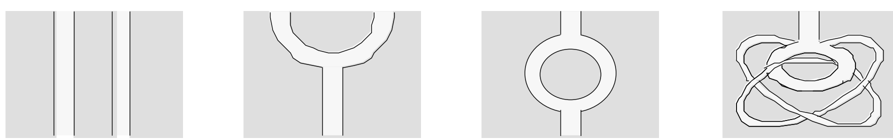

Figure 13. Some configurations that elementary topology cannot distinguish.

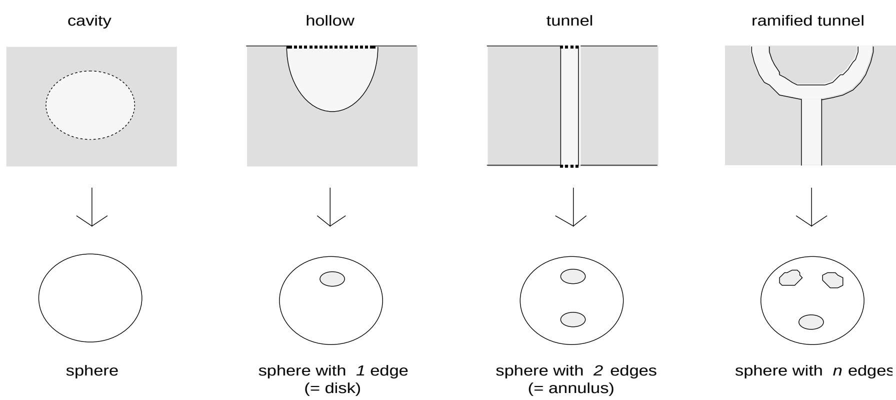

Figure 14. The morphological complexity of a hole is reflected in the topologial complexity of the contact surface of its perfect filler.

In a similar fashion, we can also extend our taxonomy by adding suitable *branching* and *knot* theories to account for further morphological complexities. We may, for instance, wish to distinguish between an **X**-shaped and an **H**-shaped tunnel/groove by counting the relevant number of nodes, or junction points. And we may want to distinguish between a straight **I**-shaped tunnel and a knotted one, between a simple **O**-shaped tunnel-cavity and a trefoil knotted one. (In all of these cases, the base theory is insensitive to what goes on inside the hole, and only considers its relations with the external surface of the object. Thus, for instance, the basic principles underlying the decision tree illustrated in figure 12 does not extend to the rightmost pattern of figure 11.) Indeed, the possibility of relying on an explicit knot theory is an immediate and advantagious consequence of the main ontological choice of the hole-theory: it is because holes are full-fledged entities—*that* type of entities, viz. immaterial spatial bodies bodies rather than negative parts—that one can investigate the ways they can be knotted together. And with the help of this extended machinery, more complex patterns of spatial interaction between holed objects and their environment can be fruitfully studied.

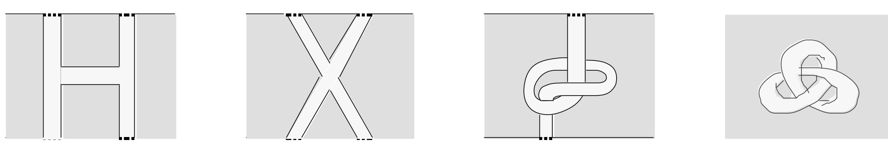

Figure 14. Taking ramifications and knots into account.

CONCLUDING REMARKS

We have focused so much on holes and the problems they pose because we take them to be indicative of the issues involved in any spatial theory aiming to combine some affinity with common sense and a suitable degree of formal sophistication. This, we believe, is the main aim for a good theory of spatial reasoning and representation that will overcome the apparent discrepancy between “psychological”, type-(i) theories, and “formal”, type-(ii) theories. As we proceed, we discover layers of problems that are recalcitrant to simple solutions and that are a sign of the presence of unresolved conceptual issues. And holes show how radical the need can be for a revision—or at least a re-examination—of the conceptual categories required for this task.

Holes are not an exception, though. Similar problems arise virtually for *every* spatial entity: not only holes or regions of space, but even material objects have been the subject of philosophical dispute. And the particular strategy one is to adopt is often a symptom of wider philosophical concerns. Thus, one might choose to concentrate exclusively on the spatial entities and on their intrinsic properties (at the expense of the environment and of the relational ties linking the objects to their environment) and thereby neglect complementary reasoning or more generally holistic components in spatial reasoning. Or one can attend to global properties of spatial situations, and fail to isolate relevant features of individual objects. To some extent this conflict (among others) can be seen as a sign of a deeper conflict between spatial absolutism and spatial relationism. But even so, we suggest the conflict may partly be resolved by integrating both perspectives: *common-sense* reasoning about space is, by one and the same token, reasoning about spatial entities.

ACKNOWLEDGMENTS. The first part of this paper includes revised material originally appeared in Casati and Varzi [1995]. Some sections also draw from Varzi [1993, 1994] and Casati [1995b]. We owe much to conversations we have enjoyed with people attending the Bolzano summer school on "Spatial Reasoning".

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., 1993, '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.

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., 1995a, 'Temporal Entities in Space', in P. Amsili, M. Borillo, and L. Vieu, eds. *Time, Space and Movement: Meaning and Knowledge in the Sensible World*. Proceedings of the 5th International Workshop, Toulouse: COREP, Part D, pp. 66-78.

— 1995b, 'The Shape Without', Workshop *Topology and Dynamics in Cognition and Perception*, International Center for Semiotic and Cognitive Studies, San Marino, December 1995.

— 1996, 'Rest and Motion: A Newtonian Framework', forthcoming in the *Proceedings of the Third Italian Conference of Logic and the Philosophy of Science* (Roma, January 1996).

Casati R., and Varzi A. C., 1994, *Holes and Other Superficialities*, Cambridge, MA: MIT Press.

— 1995, 'Basic Issues in Spatial Representation', in M. De Glas and Z. Pawlak (eds.), *Proceedings of the 2nd World Conference on the Fundamentals of Artificial Intelligence*, Paris: Angkor, pp. 63–72.

— 1996, 'The Structure of Spatial Location', *Philosophical Studies*, 82, 205–39.

Cave, C. B., and Kosslyn, S. M., 1993, 'The role of parts and spatial relations in object identification', *Perception*, 22, 229-248.

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.

Davidson, D., 1969, 'The Individuation of Events', in N. Rescher (ed.), *Essays in Honor of Carl G. Hempel*, Dordrecht: Reidel, pp. 216–34; reprinted in Davidson's *Essays on Actions and Events*. Oxford: Oxford University Press, 1980, pp. 163–80.

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 (KR94)*, San Mateo, CA: Morgan Kaufmann, 1994, 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 Institute in Cognitive Science*, Hamburg: Universität Hamburg, 1994, pp. 105-29.

Herskovits A., 1986, *Language and Spatial Cognition. An Interdisciplinary Study of the Prepositions in English*, Cambridge: Cambridge University Press.

Heydrich W., 1988, 'Things in Space and Time', in J. S. Petöfi (ed.), *Text and Discourse Constitution. Empirical Aspects, Theoretical Approaches*, Berlin and New York: de Gruyter, pp. 377-418.

Hoffman D. D. and Richards W. A., 1985, 'Parts of Recognition', *Cognition* 18, 65-96.

Lewis D. K., 1991, *Parts of Classes*, Oxford: Basil Blackwell.

Pianesi F., Varzi A. C., 1994, 'The Mereo-Topology of Event Structures', in P. Dekker and M. Stokhof (eds.), *Proceedings of the 9th Amsterdam Colloquium*, Amsterdam: ILLC, pp. 527–46.
— 1996a, 'Events, Topology, and Temporal Relations', *The Monist*, 78, 89–116.
— 1996b, 'Refining Temporal Reference in Event Structures', *Notre Dame Journal of Formal Logic*, 37, 71–83.
Reuleaux F., 1875, *Theoretische Kinematik: Grundzüge einer Theorie des Maschinenwesens*, Braunschweig: Vieweg (Eng. trans. by A. B. W. Kennedy, *The Kinematics of Machinery: Outline of a Theory of Machines*, London: Macmillan, 1876; reprinted by Dover, 1963).
Smith B., 1993, 'Ontology and the Logistic Analysis of Reality', in N. Guarino and R. Poli (eds.), *International Workshop on Formal Ontology in Conceptual Analysis and Knowledge Representation*, Padova: CNR, pp. 51–68 (revised and expanded version as 'Mereotopology: A Theory of Parts and Boundaries', *Data & Knowledge Engineering*, forthcoming).
Vandeloise C., 1986, *L'espace en français: sémantique des prépositions spatiales*, Paris: Seuil (Eng. trans. by A. R. K. Bosch, *Spatial Prepositions. A Case Study from French*, Chicago/London: University of Chicago Press, 1991).
— 1994, 'Methodology and analyses of the preposition *in*', *Cognitive Linguistics*, 5-2, 157-84.
— 1996, *Langue et Physique*, forthcoming.
Varzi A. C., 1993, 'Spatial Reasoning in a Holey World', in P. Torasso (ed.), *Advances in Artificial Intelligence. Proceedings of the 3rd Congress of the Italian Association for Artificial Intelligence*, Berlin/Heidelberg: Springer-Verlag, pp. 326–36.
— 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.
— 1996a, 'Boundaries, Continuity, and Contact,' *Noûs*, forthcoming.
— 1996b, 'Parts, Wholes, and Part-Whole Relations: The Prospects of Mereotopology,' *Data & Knowledge Engineering*, forthcoming.
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.