(Published in N. Guarino (ed.), *Formal Ontology in Information Systems,* Amsterdam: IOS Press, 1998, pp. 29–38.)

Basic Problems of Mereotopology

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

Abstract. Mereotopology is today regarded as a major tool for ontological analysis, and for many good reasons. There are, however, a number of open questions that call for an answer. Some of them are philosophical, others have direct import for applications, but all are crucial for a proper assessment of the strengths and limits of mereotopology. This paper is an attempt to put some order into this still untamed area of research. I will not attempt any answers. But I shall try to give an idea of the problems, and of their relevance for the systematic development of formal ontological theories.

1. Introduction

Mereotopology is today regarded as a major tool for formal-ontological analysis, and for many good reasons.[1] It is highly general and highly domain independent. It is ontologically neutral, treating all entities as individuals, i.e., as entities of the lowest logical type. (Set theory, by contrast, forces a distinction in ontological status between the first and second arguments of its primitive relation.) And it is intuitively attractive, dealing with formal structures (namely: structures of part and whole) that belong to the armamentarium not only of common sense and natural language but also of every empirical science.

Of course, insofar as mereotopology is only concerned with part-whole structures, it has its limits. One needs much more than mere part-whole reasoning to account for important ontological relations, for instance relations of existential dependence, causal relevance, or spatiotemporal location. One also needs to go beyond mereotopology to account for equally basic intuitions concerning, for instance, movement of parts or interaction among wholes. The world of mereotopology is, after all, a world of spheres and toruses, and we need to step into morphology or what we might also call qualitative geometry to account for basic differences in shape. We need to step into kinematics and dynamics to account for basic differences of behavior. And so on.

Even so, the fact remains that part-whole reasoning is a crucial and arguably basic ingredient of formal ontological reasoning, hence mereotopology deserves the place that it has acquired. My concern in this paper is not with the question of what comes *next*—of how mereotopology can be strengthened so as to deal with a wider range of ontological issues. Rather, I will be concerned with the loose ends of mereotopology—with issues that are still open *within* the scope of application of mereotopology itself. Some of these may just reflect the geography of the field: there is a variety of different mereotopologies by now, and in some cases the differences among these theories bear witness to genuine disagreement concerning

fundamental questions. However, some of the loose ends have nothing to do with this. They are simply questions that have not yet been fully addressed in the literature of mereotopology, and my purpose here is to bring them up for discussion. I will not attempt any answers. But I shall try to give an idea of the problems, and of their relevance for the systematic development of formal ontological theories in the future.

2. Background principles

One may view mereotopology as consisting of two independent but mutually related components: a mereological component, concerned with the concept of *parthood* (or *overlap*), and a topological component, concerned with the concept of *wholeness* (or *connection*).[2] There is no uniform agreement on the axioms of these components, nor on what sort of bridge principles should connect one component to the other. There is, however, a modicum of basic principles which, though not uniformly accepted, may be singled out as playing a central role in many attempts to capture the structure of part-whole reasoning. A brief review of this basic apparatus will simplify much of the discussion that follows.

To begin with, there are some purely structural axioms—axioms that set out the main structural features of the intended interpretation of the mereotopological primitives. Assuming *parthood* (*P*) and *connection* (*C*) to be such primitives (other options are available), these axioms may be spelled out as follows:

P-reflexivity: everything is part of itself.
P-antisymmetry: two distinct things cannot be part of each other.
P-transitivity: any part of a part of a thing is itself part of that thing.
C-reflexivity: everything is connected to itself.
C-symmetry: if a thing is connected to a second thing, the second is connected to the first.
Monotonicity: everything is connected to anything to which its parts are connected.

In other words, parthood must be a partial ordering while connection must be symmetric, reflexive, and monotonic relative to parthood. There is actually some controversy over these principles (especially *P*-transitivity)[3], but I shall not go into that here. For all good purposes, all of them are unproblematic.

Besides these very basic structural axioms, however, there are other, more substantive mereotopological principles. The most significant are:

P-extensionality: no two distinct things have the same proper parts.
P-fusion: for any number of things there is a smallest thing of which they are all part.
External contact: everything is connected with its mereological complement.

These principles are the source of some controversy, and many theories have been considered which result from dropping or at least weakening one or more of them. For instance, a recurrent objection to the principle of *P*-extensionality has been that it rules out the possibility that some objects differ exclusively by virtue of the arrangement of their parts (as two sentences consisting of the same words). Likewise, the principle of *P*-fusion has often been rejected on the grounds of its "counterintuitive" instances, i.e., "unnatural" fusions of scattered or otherwise ill-assorted entities (such as the totality of pink things, or the fusion of this cat and that umbrella). Finally, the external-contact principle has not been openly disputed, but it too fails to hold in some theories.[4]

Let us call a mereotopological theory *full-fledged* or *modest* depending on whether or not it satisfies these three additional principles. There is no question that a full-fledged mereotopology is needed if we are interested in its applications to the foundations of mathematics,

e.g., if our aim is to recast set theory and point-set-theoretic topology in nominalistic terms: you cannot go far without extensionality, fusion, or the principle of external contact.[5] But is a modest mereotopology the best we can have when it comes to formal ontology?

3. Puzzling stories

It is hard to give up full-fledgedness. Extensionality provides us with a clear criterion of individual identity. To abandon it and start distinguishing among mereologically indistinguishable entities is like opening Pandora's box, and one may end up multiplying entities beyond tolerance. (In fact, sentences made up of the same words are best described as different *sentence tokens* made up of distinct tokens of the same *word types*.) Likewise, fusion gives us a principled way of understanding the structure of mereological composition. We may feel uneasy about treating unheard-of fusions as *bona fide* individuals; we may want to say that a cat-umbrella is a different sort of thing than this cat or that umbrella. But that is no reason for doing away with such things altogether. For on the one hand, if we do so, we are likely to end up doing away with too many things besides. (Everything is odd to a degree.) On the other hand, such things are nothing over and above the parts that constitute them: as a knife is just the blade plus the handle, so too a cat-umbrella is just the cat plus the umbrella. At least, this is how I would argue. But here is where a number of deeper puzzles begin to arise.

3.1. Actual parts and potential parts

Take the knife on the table. Unlike the cat-umbrella, this is *prima facie* a non-problematic case. The knife consists of a handle and a blade which are perfectly homogeneous and well connected to each other. But we may still ask: How many things are there on the table? One (a knife)? Two (a handle and a blade)? Or three (a handle, a blade, *and* a knife)?

The question is important, for no matter what our mereotopology looks like, we must be able to count the objects in our domain of discourse. A full-fledged mereotopology does have an answer: it rejects the second option explicitly (by the fusion axiom) and suggests the third option (by extensionality): the handle, the blade, and the knife are three distinct things. To some extent this is also confirmed by ordinary speech. We often talk about parts in the same way in which we talk about whole objects; we quantify over parts, describe them, compare them with one another. ("The blade is sharper than the handle.") However, in some cases we also want to go out of our way to make it clear that our talk about parts is talk about things that do not quite have a life of their own—things that do not count except as parts of the wholes to which they belong. In this sense we would only count the knife in the first place. The handle and the blade are not objects to be included in an "inventory of the world" *over and above* the knife. And the knife is not an object to be included in an inventory *over and above* the handle and the blade. As some like to say, the knife just *is* the handle and blade; the handle and blade just *are* the knife.[6] But how is this tension to be explained? How can one reconcile these two conflicting attitudes towards the ontological status of parts?[7]

One may, of course, consider the remaining option: there is only one thing on the table, namely, the knife. This reflects an Aristotelian standpoint which is not quite incompatible with the principles of mereology: we could maintain that proper undetached parts (i.e., parts that are connected to other parts of the same whole) are not *actual* entities to be included in the domain of quantification. They are merely *potential* entities, entities which would exist if they were detached from the rest of the wholes to which they belong.[8] This is not implausible. However, it implies a modal step: to do full justice to the *internal* mereological structure of a connected whole (in a given world at a given time), and to our discourse about it, it becomes necessary to refer to what goes on in *other worlds* at the same time. And this is a slippery slope, as the next puzzles indicate.

3.2. Eubilides' handle

Forget the knife and focus on the handle, call it H. Imagine removing one molecule from it. There are then two possibilities: either we insist that the difference between what we had before and what we have after the removal of the molecule is so negligible as to have no effect on the identity of H, or we admit that a difference—small as it may be—exists.

In the first case, we must give up the principle of extensionality, for H minus the molecule is a proper part of H, hence extensionally distinct from it. We therefore need some other criterion for individual identity. However, since the argument can be iterated at will, this approach will take us straight to a form of Eubulides' paradox:[9]

1. H minus one molecule is the same as H.
2. For all *n*, if H minus *n* molecules is the same as H, then H minus *n*+1 molecules is the same as H (one molecule makes no difference).
3. *Ergo*, H minus one billion molecules is the same as H—which is absurd.

(The argument is valid: the conclusion follows by one billion applications of *modus ponens*.)

If, on the other hand, we go for the second option and admit that the removal of one single molecule affects the identity of H, then we preserve extensionality. However, since a handle minus a molecule is still *a handle* (albeit not *the same handle*), we seem to be forced to conclude that the spatiotemporal region occupied by H is actually hosting *n*+1 handles (actual handles? potential handles?), where *n* is the number of molecules of which H is composed. Indeed, there would seem to be many more handles than that. For, of course, each of the spatiotemporal regions occupied by the *n* handles properly included in H is inhabited by *n*–1 more handles, each of which is obtained simply by removing one molecule from the handle at issue. Evidently, this is a slippery slope just as bad as Eubulides' paradox.[10]

3.3. Theseus' ship

The puzzles generalizes. Consider Theseus's vessel, which—as the legend has it—was constantly being refitted and newly fashioned by replacing old planks with new, strong ones.[11] If we accept the principle of extensionality, then the vessel keeps changing: instead of a ship that endures through time, we have a whole array of ships, each of which replaces its predecessor as time goes by and the pieces get replaced. And the same goes for lots of similar cases: just as Theseus could not sail twice in the same vessel, so you cannot step twice in the same river, look twice at the same cloud, eat twice at the same restaurant. Indeed, the same goes for all organisms, you and I included. For we too keep losing old parts (so to say) and gaining new ones.

Evidently, the problem is a general one and relates to a series of other philosophical conundrums: from the distinction between essential and accidental parts to the question of identity and individuation criteria, including criteria for identifying things through time and across possible worlds. But there is one more complication. For we can always ask what happened to those parts that get replaced as times goes by. Suppose *every* part of Theseus' ship gets replaced. And suppose someone kept all the old pieces and afterwards put them back together in the very same order as they were originally, again making a ship of them. The result would be a ship perfectly indistinguishable from Theseus' initial ship. Indeed, it would also consist exactly of the same parts. Would *it* be a different ship?[12]

3.4. Tibbles' tail

All three problems mentioned above relate to the following one. Consider Tibbles, the cat. Call her tail 'Tail' and the rest 'Tib'. Is Tibbles the same as Tib plus Tail? The postulates of a full-fledged mereotopology (at least the postulates of mereological extensionality and fu-

sion) say so: an object is just the fusion of its parts. However, if we allow for the possibility that an object undergo change without losing its individual identity, then the two creatures are different in a significant respect: Tibbles may lose her tail in an accident; Tib+Tail may not. (By definition, Tib+Tail includes Tail as a proper part.) Once again, this seems to violate extensionality: before the accident, there seem to be two distinct entities with the same parts (distinct because of their different modal properties). But is that right?[13]

To make things worse, let us now ask about the relationship between Tibbles and Tib. Clearly, at each time t of the first phase of Tibbles' life, when she is a perfectly normal cat with tail, Tibbles is distinct from Tib, for it includes Tib as a proper part. But what about their relationship at a later time t', after Tibbles has lost her tail and has become coincident with Tib? If we say that the two have become one, then we are in trouble, for we have:

1. Tibbles at $t \neq$ Tib at t (one is a proper part of the other)
2. Tibbles at t = Tibbles at t' (Tibbles survives the loss of Tail)
3. Tib at t = Tib at t' (Tib is not affected by whatever happens to Tail)
4. Tibbles at t' = Tib at t' (both have the same parts)

Yet 2–4 jointly imply the negation of 1 by transitivity of identity, so we are in plain contradiction. If, on the other hand, we deny that Tibbles and Tib have become one and the same thing, i.e., if we deny 4 (and extensionality with it), then we must abandon the traditional identity criterion according to which two distinct material bodies cannot occupy the same spatial region at the same time. And this is just as high a cost to pay.[14]

Of course we could also keep 1 and 4 and give up either 2 or 3. Rejecting 2 takes us back to the case of Theseus' ship, suggesting a form of mereological essentialism: the removal of a part (even a very tiny and seemingly inessential one) affects the identity of the whole. But rejecting 3 seems to imply an equally doubtful form of essentialism—a form of topological essentialism, to give it a name—to the effect that the removal of a part affects the identity of another, adjacent but mereologically disjoint part. And if one worries about mereological essentialism, why should one accept *that*?[15]

3.5. Tibbles' wound

There is another puzzle raised by Tibble's story. For exactly what happens when Tib and Tail get separated? An object splits into two. Two connected things become disconnected. But what is the correct story here?

One popular answer is that cutting an object is a way of revealing matter in its interior, a way of bringing to light new surfaces: at t the boundary between Tib and Tail is an *interior* boundary; at t' it is an *exterior* boundary.[16] If so, however, a question naturally arises: who inherited that boundary—Tib or Tail? And if it is Tib (say) which inherited the boundary, shall we say that at t' Tail is topologically a "semi-open" entity? What sort of entity would that be? Nor is cutting necessary for the puzzle to arise. Consider the interior boundary separating Tib and Tail at t: does it belong to Tib or to Tail? It cannot belong to neither, for otherwise something would lie *between* them, contradicting the assumption that Tib and Tail are in contact. And it cannot belong to both, for otherwise Tib and Tail would not just be in contact: they would overlap. Shall we then conclude that at t, before the cut, Tibbles consists of two parts, one of which is closed while the other is partially open? It is hard to imagine how these questions could be answered without selecting one or other term at random.[17]

Some take this to be a *reductio ad absurdum* of the very idea of a boundary and, with it, of the distinction between open and closed entities. Talk of boundaries would then be a mere *façon de parler* about other things—for example about infinite series of nested extended bodies.[18] But this is no easy way out. For boundaries can certainly be *defined* in mereotopology. This can be done in two steps:

x is a boundary part of *y* $=_{df}$ every part of *x* is connected to the complement of *y*
x is a boundary of *y* $=_{df}$ *x* is a boundary part either of *y* or of the complement of *y*

There is nothing mysterious or controversial in these definitions (though other accounts are possible[19]). So if we want to go boundary-free, we have to say that these relations are empty—in particular, there can be no boundary parts. But how can this be? By the external contact principle, every object *y* is in contact with its complement. Hence, by the fusion principle, the fusion *z* of all parts of *y* that are not in contact with *y*'s complement must be a proper part of *y* (since it is a part of *y* and *y* is not part of it). But then, by the extensionality principle, there must exist a difference between *y* and *z*, i.e., there must exist some part *x* of *y* that does not overlap *z*. And such an *x* qualifies as a boundary part of *y*. Thus, if we want to go boundary-free, one of the three principles must go. But why should the fusion principle be called into question here—why should a purely mereological notion suffer from a topological desease? And why should a boundary-free theory give up the idea that every object is connected to its complement? The only option seems once again to give up extensionality. In fact, one would have to give up an even weaker (and less controversial) principle:

Supplementation: if an object has a proper part, it has another, non-overlapping part.

And without this principle, it is hard to say that we have a mereotopology at all: how can one get a whole from a proper part without *adding* some other part?[20]

3.6. The hole in the donut

Tibbles had an accident and lost her tail. This affected her topology and her mereology. But imagine an object changing only its topology, without any consequences for its mereological composition. Imagine opening up a donut so as to transform it into a C-shaped object—the hole goes. Does the donut survive? Or imagine that a blob starts growing a finger which eventually comes round to meet the main body again, forming a sort of handle. At the instant that it does so, the topology of the blob changes: we had a sphere; now we have a torus. Has the blob survived the change?

Mereotopology is neutral with regard to these questions—even in a full-fledged version. But it is neutral only insofar as parthood and connection are not sufficient to characterize variations of topological genus. As a theory of parts and *wholes*, mereotopology says nothing about the difference between parts and *holes*.[21] Yet this is arguably a limitation, not an advantage. On the one hand, holes are not just some special kind of part—*negative* or *immaterial* parts, as some have suggested.[22] As the blob example shows, a hole is not necessarily *something missing*. And even where the hole is properly characterized as such (as with a hole in the roof or a tunnel through a mountain), it does not follow that the hole itself is *part* of its material host. It is attached to it; we may even say that it is vitally attached to it (you cannot separate a hole from its host, just as you cannot separate a boundary from the object that it bounds). But hole and host remain quite distinct things, and mereotopology should have room for the distinction.[23]

On the other hand, if holes are not negative parts, what are they? Some would argue that they are nothing at all. This attitude requires a thorough eliminative strategy, a systematic way of paraphrasing every hole-committing sentence by means of a sentence that does not refer to or quantify over holes.[24] In some cases this is easy. ("There are holes in the cheese" becomes "The cheese is perforated", where '… is perforated' is just an ordinary shape predicate like '… is flat' or '… is pyramidical'.) However, it is far from obvious that such a strategy can be carried out in full. I prefer to treat holes as *bona fide* entities. But this means adding a new primitive relation (say a binary relation '*x* is a hole in *y*') to the basic mereotopological machinery. And this may carry in its wake an array of additional philosophical conun-

drums. What exactly is the relationship between a hole and its host? What is the relationship between a hole and its guests—the things you can place into it? What happens when two holes come into contact? And can you split a hole into two parts? Can you take away a part from a hole? If you do, what remains?[25]

3.7. The edible stuff

If holes exist, they are immaterial—they are made of nothing (like ghosts or angels). A donut, by contrast, is made of edible stuff. Is the stuff *part* of the donut? Is the relation of material constitution a mereological relation?

If it is, then the stuff must *coincide* with the donut. For no part of the donut is left if you eat all the stuff it is made of (unless we want to take the mind-body problem for donuts seriously). And if no part of the donut is left upon eating the stuff, then by the principle of supplementation the stuff cannot be a proper part of the donut. On the other hand, it would seem that the stuff *cannot* coincide with the donut, for the two have distinct persistence properties: if you cut it into a thousand pieces, the donut is gone but the stuff is all there. The stuff existed before the donut and may outlast it—or so the argument goes.[26] Shall we then conclude that constitution is not a mereological relation—that the spatiotemporal region presently occupied by the donut is actually occupied by two distinct entities? If so, how exactly are they related? And how does this case differ from that of Tibbles and Tib+Tail?[27]

4. Why bother?

These are only a sample of the puzzles that surround the concepts of part and whole. Some are purely philosophical; others have direct import for applications. (The treatment of boundaries, for instance, is crucial to any application of the theory to the geographic domain; the treatment of holes is crucial to its application to the domain of mechanical engineering.[28]) But all seem to me to be critical for a proper assessment of the strengths and limits of mereotopology.

Perhaps at this point the choice between a modest or a full-fledged theory could be viewed as a function of the context: depending on the intended application, we may go for a weaker or stronger theory. For instance, if we are modeling a domain of bare spatial regions or temporal intervals, then all the axioms of a full-fledged mereotopology appear unproblematic. Indeed, in that case the topological component can even be strengthened, as it seems reasonable to assume, not only the principles of *P*-extensionality and *P*-fusion, but also the (stronger) analogues for connection:

> *C-extensionality*: no two distinct things have the same connections.
> *C-fusion*: for any number of things there is a smallest thing which is connected exactly to what they are connected to (jointly or individually).[29]

If, on the other hand, we are modeling a domain of artifacts or living organisms, then the puzzles above must be taken seriously. We may then come up with a more modest theory, or try to save the stronger principles in some form.

This may well prove to be a convenient way of proceeding, at least to avoid unnecessary complications; but I do not recommend it. I do not recommend a context-based strategy even if our purpose is applicative rather than philosophical. The whole point of recent efforts aimed at the development of large scale ontologies is that these be usable for a variety of tasks. Formal ontology is conceived of as a general tool for analysing and integrating large and often heterogeneous bodies of information—information relative to a wide variety of different domains.[30] If mereotopology is to play a role in this enterprise, it cannot proceed on an *ad hoc*

basis. The puzzles must be handled in a general way independently of their original philosophical motivations.

To put it differently, there seems to me to be a gap between recent work aimed at developing so-called top-level ontologies, on the one hand, and technical developments in mereotopology (and other branches of formal ontology), on the other. While the former is aiming at generality and domain neutrality, the latter have arisen mostly out of specific needs and for specific purposes—for instance, for applications in the area of spatial or spatiotemporal reasoning. The resulting variety of theories has a corresponding variety of merits. But the link with ontology is unclear. If we are to take an open-faced attitude towards entities of a large variety of sorts besides spatial regions and temporal intervals—if the ontology we are interested in is also an ontology of cats, ships, and donuts that move around and undergo change—then we can hardly hope to go far by confining ourselves to the mereotopology of their spatiotemporal locations. Whatever the applicative agenda here, there is serious work to be pursued. The puzzles may be bizarre. But the study of limit cases is often the best starting point for understanding the ordinary cases.[31]

Notes

[1] See [54] for an introduction and review of mereotopology and [47] for the link with formal ontology.

[2] On the demarcation between these two components (and a taxonomy of the main options) see [53].

[3] See e.g. [15, 27].

[4] The objection to *P*-extensionality goes back to [25] and [40], while the objection to the *P*-fusion principle can be traced back to [35]. (A rich literature has grown out of these objections; see [42] for a review.) As for the principle of external contact, its failure is a feature of certain boundary-free mereotopologies in the tradition of [14]. I'll get back to all of these issues below.

[5] Historically the first intended applications of mereology were indeed in the foundations of mathematics [29, 30].

[6] See [4] and [31].

[7] For more on this tension, see [57].

[8] The view that proper undetached parts are merely potential has been defended in [52]. See also the discussion in [43], Ch. 3.

[9] Eubulides' paradox is typically discussed in connection with vague predicates (see [62] for an overview). In this form, the puzzle is also known as the "problem of the many": see [50].

[10] See [18], §110, and [34] for a discussion of this complication.

[11] The story of Theseus's ship comes from Plutarch's *Lives* and the literature on the topic is very large. See [37] for a collection of texts and ch. 7 of [42] for a critical review.

[12] The duplication problem was raised by Hobbes and is now part of the standard discussion on Theseus's ship. See e.g. [13, 61].

[13] Tibbles' puzzle has been introduced into contemporary philosophical discussion in [59], though its first variants can be traced back at least to the Stoics [41]. See ch. 6 of [33] for an introduction and [39] for a selection of recent texts.

[14] See [9] for a discussion of this point.

[15] On mereological essentialism see [12, 50, 60]. The position of topological essentialism may be associated with those authors who feel uncomfortable with scattered fusions: see [11] for a discussion.

[16] See [1] for a statement of the view.

[17] For more on these problems see [46, 48, 56]; for the historical background see [63]. See also [19] for analogues in the temporal domain.

[18] This position is rooted in [58] and is well exemplified in the recent literature by [2, 21].

[19] See e.g. [45, 46].

[20] See [42] for this sort of objection.

[21] See [20] for an attempt to capture genus differences in terms or connection relations.

[22] See e.g. [26].

[23] For a more detailed criticism of the view that holes are negative parts, see [10].

[24] The eliminative strategy was first considered in [32].

[25] Some work in this direction has been done in [8]. See also [55] for formal developments.

[26] For an overview of this position, see [22]. For a recent philosophical exchange, see [28, 36]. For an AI perspective, see [5].

[27] See [17] for a discussion of the analogy. I myself believe that the cases, though conceptually different, admit of the same solution, namely that material constitution *is* identity, just as mereological composition *is* identity: see [57].

[28] See e.g. [7, 44] and [3, 6], respectively.

[29] These principles hold in the mereotopology of [14], for instance.

[30] On this see [23, 24].

[31] Thanks to Massimiliano Carrara and Barry Smith for helpful comments on a previous draft.

References

[1] Adams, E. W., 1984, 'On the Superficial', *Pacific Philosophical Quarterly* 65, 386–407.

[2] Asher N. and Vieu L., 1995, 'Toward a Geometry of Common Sense: A Semantics and a Complete Axiomatization of Mereotopology', in *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, San Mateo, CA: Morgan Kaufmann, pp. 846–852.

[3] Bäckström C., 1990, 'Logical Modelling of Simplified Geometrical Objects and Mechanical Assembly Processes', in S. Chen (ed.), *Advances in Spatial Reasoning, Volume 1*, Norwood: Ablex, pp. 35–61.

[4] Baxter D., 1988, 'Identity in the Loose and Popular Sense', *Mind* 97, 575–82.

[5] Borgo S., Guarino N., and Masolo C., 1995, 'A Naive Theory of Space and Matter', 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 E, pp. 29–32.

[6] Borgo S., Guarino N., and Masolo C., 1997, 'Qualitative Spatial Modelling Based on Parthood, Strong Connection, and Congruence', LADSEB–CNR Internal Report 03/97.

[7] Casati R., Smith B., and Varzi, A. C., 1998, 'Ontological Tools for Geographic Representation', this volume.

[8] Casati R. and Varzi A. C., 1994, *Holes and Other Superficialities*, Cambridge, MA, and London: MIT Press (Bradford Books).

[9] Casati R. and Varzi A. C., 1996, 'The Structure of Spatial Location', *Philosophical Studies* 82, 205–39.

[10] Casati R. and Varzi A. C., 1997, 'Spatial Entities', in O. Stock (ed.), *Spatial and Temporal Reasoning*, Dordrecht/Boston/London: Kluwer Academic Press, pp. 73–96.

[11] Casati R. and Varzi A. C., 199+, 'Topological Essentialism', *Philosophical Studies*, to appear.

[12] Chisholm R. M., 1973, 'Parts as Essential to Their Wholes', *Review of Metaphysics* 26, 581–603.

[13] Chisholm R. M., 1976, *Person and Object. A Metaphysical Study*, London: Allen & Unwin.

[14] Clarke B. L., 1981, 'A Calculus of Individuals Based on "Connection"', *Notre Dame Journal of Formal Logic* 22, 204–218.

[15] Cruse D. A., 1979, 'On the Transitivity of the Part–Whole Relation', *Journal of Linguistics* 15, 29–38.

[16] Eberle R. A., 1970, *Nominalistic Systems*, Dordrecht: Reidel.

[17] Fine K., 1994, 'Compounds and Aggregates', *Noûs* 28, 137–58.

[18] Geach P. T., 1980, *Reference and Generality*, 3rd Edition, Ithaca: Cornell University Press.

[19] Galton A., 1996, 'Time and Continuity in Philosophy, Mathematics, and Artificial Intelligence', *Kodikas/Code* 19, 101–19.

[20] Gotts N. M., 1994, '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.

[21] Gotts N. M., Gooday J. M., and Cohn A. G., 1996, 'A Connection Based Approach to Common–Sense Topological Description and Reasoning', *The Monist* 78, 51–75.

[22] Griffin N., 1977, *Relative Identity*, Oxford: Clarendon Press.

[23] Guarino N., 1995, 'Formal Ontology, Conceptual Analysis and Knowledge Representation', *International Journal of Human–Computer Studies* 43, 625–40.

[24] Guarino N., 1997, 'Semantic Matching: Formal Ontological Distinctions for Information Organization, Extraction, and Integration', in M. T. Pazienza (ed.) *Information Extraction: A Multidisciplinary Approach to an Emerging Information Technology*, Berlin/Heidelberg: Springer-Verlag, pp. 139-170

[25] Hempel C. G., 1953, 'Reflections on Nelson Goodman's "The Structure of Appearance"', *The Philosophical review*, 62, 108–16.

[26] Hoffman D. D. and Richards W. A., 1985, 'Parts of Recognition', *Cognition* 18, 65–96.

[27] Iris M. A., Litowitz B. E., and Evens M., 1988, 'Problems of the Part–Whole Relation', in M. Evens (ed.), *Relations Models of the Lexicon*, Cambridge: Cambridge University Press, pp. 261–288.
[28] Johnston M., 1992, 'Constitution Is Not Identity', *Mind* 101, 89–105
[29] Leonard H. S. and Goodman N., 1940, 'The Calculus of Individuals and Its Uses', *Journal of Symbolic Logic* 5, 45–55.
[30] Leśniewski S., 1927–30, 'O podstawach matematyki', *Przegld Filozoficzny*, 30, 164–206; 31, 261–91; 32, 60–101; 33, 77–105, 142–70 (English translation by D. I. Barnett, 'On the Foundations of Mathematics', in S. Leśniewski, *Collected Works*, ed. by S. J. Surma, J. T. Srzednicki, D. I. Barnett, F. V. Rickey, Dordrecht, Boston, and London: Kluwer Academic Publishers, 1992, Vol. 1, pp. 174–382).
[31] Lewis D. K., 1991, *Parts of Classes*, Oxford: Basil Blackwell.
[32] Lewis D. K. and Lewis S. R., 1970, 'Holes', *Australasian Journal of Philosophy* 48, 206–12.
[33] Loux M. J., 1998, *Metaphysics. A Contemporary Introduction*, London: Routledge.
[34] Lowe E. J., 1982, 'The Paradox of the 1,001 Cats', *Analysis* 42, 27–30.
[35] Lowe V., 1953, 'Professor Goodman's Concept of an Individual', *Philosophical Review* 62, 117–126.
[36] Noonan H. W., 1993a, 'Constitution Is Identity', *Mind* 102, 133–46.
[37] Noonan H. W. (ed.), 1993b, *Identity*, Aldershot: Dartmouth.
[38] Randell D. A., Cui Z., and Cohn A. G., 1992, 'A Spatial Logic Based on Regions and Connection', in B. Nebel, C. Rich e W. Swartout (eds.), *Principles of Knowledge Representation and Reasoning. Proceedings of the Third International Conference*, Los Altos: Morgan Kaufmann, pp. 165–76.
[39] Rea M., 1997, *Material Constitution. A Reader*, Lanham, MD: Rowman & Littlefield.
[40] Rescher N., 1955, 'Axioms for the Part Relation', *Philosophical Studies* 6, 8–11.
[41] Sedley D., 1982, 'The Stoic Criterion of Identity', *Phronesis* 27, 255–75.
[42] Simons P. M., 1987, *Parts. A Study in Ontology*, Oxford: Clarendon Press.
[43] Smith B., 1994, *Austrian Philosophy. The Legacy of Franz Brentano*, Chicago and La Salle, IL: Open Court.
[44] Smith, B., 1995, 'On Drawing Lines on a Map', in A. U. Frank and W. Kuhn (eds.), *Spatial Information Theory. A Theoretical Basis for GIS*, Berlin/Heidelberg: Springer-Verlag, pp. 475–84.
[45] Smith B., 1996, 'Mereotopology: A Theory of Parts and Boundaries', *Data and Knowledge Engineering* 20, 287–304.
[46] Smith, B., 1997, 'Boundaries: An Essay in Mereotopology', in L. H. Hahn (ed.), *The Philosophy of Roderick Chisholm*, Chicago and La Salle, IL: Open Court, pp. 534–61.
[47] Smith B., 1998, 'The Basic Tools of Formal Ontology', this volume.
[48] Smith B. and Varzi A. C., 199+, 'Fiat and Bona Fide Boundaries', submitted.
[49] Tarski A., 1935, 'Zur Grundlegung der Booleschen Algebra. I', *Fundamenta Mathematicae* 24, 177–198 (English translation by J. H. Woodger, 'On the Foundations of the Boolean Algebra', in A. Tarski, *Logics, Semantics, Metamathematics, Papers from 1923 to 1938*, Oxford: Clarendon Press, 1956, pp. 320–41).
[50] Unger P., 1980, 'The Problem of the Many', *Midwest Studies in Philosophy* 6, 411–467.
[51] Van Cleve J., 1986, 'Mereological Essentialism, Mereological Conjunctivism, and Identity Through Time', *Midwest Studies in Philosophy* 11, 141–56.
[52] Van Inwagen P., 1981, 'The Doctrine of Arbitrary Undetached Parts', *Pacific Philosophical Quarterly* 62, 123–37.
[53] Varzi A. C., 1994, 'On the Boundary Between Mereology and Topology', in R. Casati, B. Smith, and G. White (eds.), *Philosophy and the Cognitive Sciences*, Vienna: Hölder-Pichler-Tempsky, pp. 423–42.
[54] Varzi A. C., 1996a, 'Parts, Wholes, and Part–Whole Relations: The Prospects of Mereotopology,' *Data and Knowledge Engineering* 20, 259–86.
[55] Varzi A. C., 1996b, 'Reasoning about Space: The Hole Story', *Logic and Logical Philosophy* 4, 3–39.
[56] Varzi A. C., 1997, 'Boundaries, Continuity, and Contact', *Noûs* 31, 26–58.
[57] Varzi A. C., 199+, 'Mereological Minimalism', submitted.
[58] Whitehead, A. N., 1929, *Process and Reality. An Essay in Cosmology*, New York: Macmillan.
[59] Wiggins D., 1968, 'On Being in the Same Place at the Same Time', *Philosophical Review* 77, 90–95.
[60] Wiggins D., 1979, 'Mereological Essentialism: Asymmetrical Essential Dependence and the Nature of Continuants', *Grazer philosophische Studien* 7, 297–315.
[61] Wiggins D., 1980, *Sameness and Substance*, Oxford: Basil Blackwell.
[62] Williamson T., 1994, *Vagueness*, London/New York: Routledge.
[63] Zimmerman, D. W., 1996, 'Indivisible Parts and Extended Objects: Some Philosophical Episodes from Topology's Prehistory', *The Monist* 79, 148–80.