

Chapter 1

SPATIAL REASONING AND ONTOLOGY: PARTS, WHOLES, AND LOCATIONS

Achille C. Varzi
Columbia University

[Published in M. Aiello, I. Pratt-Hartmann, and J. van Benthem (eds.), *Handbook of Spatial Logics*, Berlin: Springer-Verlag, 2007, pp. 945-1038]

Spatial reasoning is no abstract business. It is, to a great extent, reasoning about entities *located in space*, and such entities have spatial *structure*. If the table is in the kitchen, then it follows that the table top is in the kitchen, and it follows because the top is *part of* the table. If the concert took place at the stadium, then it didn't take place in the theater, for concerts are spatially *continuous*. Even when we reason about empty places, we typically do so with an eye to the anatomy of their potential tenants: space as such is perceptually remote and we can hardly understand its structure without imagining what *could* fill the void.

This general feature of our spatial competence might suggest a deep metaphysical truth, to the effect that concrete entities such as objects and events are fundamentally prior to, and independent of, their spatial receptacles. It might even suggest that space itself is just a fiction, a picture of some kind: really there are only objects and events spatially related to one another in various ways. Such was, for instance, the gist of Leibniz's stern *relationism* against Newton's *substantivalism*, in spite of the major role the idea of space plays in the sciences. At the same time, one might argue that our understanding of the spatial structure of objects and events, including their spatial relationships, depends significantly on our understanding of the structure of space *per se*: that the spatial features we attribute to objects and events are somehow inherited from those of the spatial regions they occupy. Thus, for example, we are

inclined to say that ordinary objects have parts insofar as their spatial regions have parts. We may be inclined to say that the top is part of the table because of its salience and functional role, but we may just as easily talk about the top half of a sphere, or its inner parts, in spite of their lacking any cognitive or functional salience: we identify (and reify) such parts in terms of the parts of the region the sphere occupies.

This tension (not to say this ambiguity) between concrete object-oriented thinking and abstract space-oriented thinking is responsible for many of the philosophical issues that lie behind any formal theory of spatial reasoning. On the one hand, it is natural (if not necessary) to supplement the theory with an explicit account of what kinds of thing may enter into its scope, an account of the sorts of entity that can be located in space—in short, an account of what may be collected under the rubric of “spatial entities”. On the other hand, we also want the theory to be independent of any specific ontological biases we might have. Whatever spatial entities we are inclined to build into the basic furniture of the world—subatomic particles, middle-size objects of the garden variety, large scattered entities such as crowds, forests, archipelagos, galaxies—our reasoning about their spatial properties and relations appears to be governed by the same general principles, and it is natural to think that such principles must reflect our understanding of the structural features of the spatial environment in which such entities are located. In short, although a theory of spatial reasoning may be viewed as an example of applied logic, it may also be regarded as an example of a formal theory whose principles do not necessarily depend on the intended domain of application except to the extent that the domain must include entities that properly qualify as *spatial* entities of some sort.

The purpose of this chapter is to take a closer look at these delicate matters. Rather than doing this in general, however, we shall look at how the subtle interplay between purely spatial intuitions and intuitions about concrete spatial entities shows up in the construction of a formal theory. More specifically, we shall consider three sorts of theory, each of which occupies a prominent position in recent literature: (1) mereology, or the theory of parthood relations; (2) topology, broadly understood as a theory of qualitative spatial relations such as continuity and contiguity; and (3) the theory of location proper, which deals explicitly with the relationship between an entity and the spatial region it occupies. Arguably, such theories may be viewed as jointly contributing to an overall appraisal of our spatial competence, and over the last few years there has been considerable progress in each direction. At this point there is some need for a philosophical pause, and our purpose in this chapter is to go some way in the direction of a systematic assessment.

1. Philosophical issues in mereology

Let us begin with mereology. This is often defined as the theory of the part-whole relation, but such a definition is misleading. It suggests that mereology has something to say about both parts and wholes, which is not true. As we shall see in Section 2, the notion of a whole goes beyond the conceptual resources of mereology and calls for topological concepts and principles of various sorts. By itself, mereology is best understood as the theory of the parthood relation, regardless of whether the second term of the relation may be said to qualify as a *whole* entity. Thus, for instance, it is a mereological fact that my hand is part of my arm just as it is a mereological fact that it is part of my body, although it may plausibly be argued that only my body qualifies as a whole (maximally connected) object; my arm doesn't.

It is also worth pointing out that mereology has a long pedigree, which makes it a central chapter, not only of formal theories concerned with spatial reasoning, but of any theory in the realm of formal logic and ontology broadly understood. Its roots can be traced back to the early days of philosophy, beginning with the Pre-Socratic atomists and continuing throughout the writings of Plato (especially the *Parmenides* and the *Thaetetus*), Aristotle (the *Metaphysics*, but also the *Physics*, the *Topics*, and *De partibus animalium*), and Boethius (*In Ciceronis Topica*). Mereology occupies a prominent role also in the writings of medieval ontologists and scholastic philosophers such as Peter Abelard, Thomas Aquinas, Raymond Lull, and Albert of Saxony, as well as in Jungius's *Logica Hamburgensis* (1638), Leibniz's *Dissertatio de arte combinatoria* (1666) and *Monadology* (1714), and Kant's early writings (especially the *Monadologia physica* of 1756). As a formal theory of the parthood relation, however, mereology made its way into our times mainly through the work of Franz Brentano and of his pupils, especially Husserl's third *Logical Investigation* (1901). The latter may rightly be considered the first attempt at a rigorous formulation of the theory, though in a format that makes it difficult to disentangle the analysis of mereological concepts from that of other formal notions (such as the relation of ontological dependence). It is not until Leśniewski's *Foundations of a General Theory of Manifolds* (1916) that a pure theory of parthood was given an exact formulation. And because Leśniewski's work was largely inaccessible to non-speakers of Polish, it is only with the publication of Leonard and Goodman's *The Calculus of Individuals* (1940) that mereology has become a chapter of central interest for modern ontologists and logicians. Indeed, although Leśniewski's and Leonard and Goodman's theories came in different logical guises, they are sufficiently similar to

be recognized as a common basis for most subsequent developments. The question that interests us here is how such developments—and the variety of motivations that lie behind them—reflect and affect our understanding of mereology as a formal theory of spatial reasoning.¹

1.1 ‘Part’ and parthood

To this end, the first thing to observe is that the word ‘part’ has many different meanings in ordinary language, not all of which correspond to the same relation. In a way, it can be used to indicate any portion of a given entity, regardless of whether the portion itself is attached to the remainder, as in (1), or undetached, as in (2); cognitively salient, as in (1)–(2), or arbitrarily demarcated, as in (3); self-connected, as in (1)–(3), or disconnected, as in (4); homogeneous, as in (1)–(4), or gerrymandered, as in (5); material, as in (1)–(5), or immaterial, as in (6); extended, as in (1)–(6), or unextended, as in (7); spatial, as in (1)–(7), or temporal, as in (8); and so on.

- (1) The handle is part of the mug.
- (2) This cap is part of your pen.
- (3) The left half is your part of the cake.
- (4) The cutlery is part of the tableware.
- (5) This stuff is only part of what I bought.
- (6) That area is part of the living room.
- (7) The outermost points are part of the perimeter.
- (8) The first act was the best part of the play.

All of these cases illustrate the notion of parthood that forms the focus of mereology. Often, however, the English word ‘part’ is used in a restricted sense. For instance, it may be used to designate only the cognitively salient relation illustrated in (1) and (2). In this sense, the parts of an object x are just its “components”, i.e., those parts that are available as individual units regardless of their interaction with the other parts of x . (A component is *a part* of an object, rather than just *part* of it; see e.g. Tversky 1989). Clearly, the properties of such restricted relations may not coincide with those of parthood broadly understood, so the principles of mereology should not be expected to carry over automatically.

Also, the word ‘part’ is sometimes used in a broader sense, for instance to designate the relation of material constitution, as in (9), or the relation of mixture composition, as in (10), or even a relation of conceptual inclusion, as in (11):

¹For a historical survey of mereology, see Henry (1991), Burkhardt and Dufour (1991), and Simons (1991c). For systematic comparisons, see Simons (1987) and Ridder (2002).

- (9) The clay is part of the statue.
- (10) Gin is part of martini.
- (11) Writing detailed comments is part of being a good referee.

The mereological status of these relations, however, is controversial. The constitution relation exemplified in (9) was included by Aristotle in his threefold taxonomy (*Metaphysics*, Δ , 1023b), but many contemporary authors would rather construe it as a *sui generis*, non-mereological relation (see Rea 1997 and references therein).² Similarly, the ingredient-mixture relationship exemplified in (10) is subject to controversy, as the ingredients may involve significant structural connections besides spatial proximity and may therefore fail to retain important characteristics they have in isolation (see Sharvy 1983). As for statements such as (11), it may simply be contended that the term ‘part’ appears only in the surface grammar and disappears at the level of logical form, for instance if (11) is paraphrased as “A good referee is one who writes detailed comments”. For more examples and tentative taxonomies, see Winston *et al.* (1987), Gerstl and Pribbenow (1995), and Iris *et al.* (1988).

Finally, it is worth stating explicitly that mereology is typically construed as a piece of formal ontology, i.e., a theory of certain formal properties and relations that are exemplified across a wide range of domains, *whatever the nature of the entities in question*. Thus, although both Leśniewski’s and Leonard and Goodman’s original theories betray a nominalistic stand, reflecting a conception of mereology as a parsimonious alternative to set theory,³ most contemporary formulations assume no ontological restriction on the field of ‘part’. The relata can be individual entities as in (1)–(8), but also abstract entities such as propositions, sets, types, or properties, as in:

- (12) That premise is part of my argument.
- (13) The domain of quantification is part of the model.
- (14) The colon is part of the title.
- (15) Humanity is part of personhood.

(The example in (11) may perhaps be read as expressing a mereological relation between properties, too.) This “ontological innocence” of mereology plays of course an important role in the appraisal of what principles

²Actually, if the statue is identified with the lump of clay, as some would argue (e.g. Noonan 1993 vs. Johnston 1993), and if identity is treated as a limit case of (improper) parthood, as we shall indeed suppose, then the relation of material constitution is a mereological relation. This, however, is the subject of controversy and we shall come back to it in due time.

³To be sure, the original calculus of individuals had variables for classes; the class-free version is due to Goodman (1951). On the link between mereology and nominalism, see Eberle (1970).

should hold unrestrictedly: greater generality means fewer axioms, and here the tension between the tasks of an applied logic and those of a purely formal theory shows up most vividly. In the following we focus primarily on the spatially salient uses of ‘part’, but it is important to keep this tension in mind when it comes to assessing the philosophical underpinnings of the most controversial tenets of mereology.

1.2 Basic principles

With these provisos, let us proceed to unpack the theory. We may ideally distinguish two sorts of mereological principles. On the one hand, there are principles that may be thought of as purely “lexical” axioms fixing the intended meaning of the relational predicate ‘part’. On the other, there are principles that go beyond the obvious and aim at greater sophistication and descriptive power. Exactly where the boundary should be drawn, however, is by itself a matter of controversy.

1.2.1 Parthood as a partial ordering. The obvious is this: regardless of how one feels about matters of ontology, if ‘part’ stands for the general relation exemplified by all of (1)–(8) above, then it stands for a partial ordering—a reflexive, transitive, antisymmetric relation:

- (16) Everything is part of itself.
- (17) Any part of any part of a thing is itself part of that thing.
- (18) Two distinct things cannot be part of each other.

As it turns out, virtually every theory put forward in the literature accepts (16)–(18), though it is worth mentioning some misgivings that have occasionally been raised.

Concerning reflexivity, one might observe that many legitimate senses of ‘part’ do not countenance saying that a whole is a part of itself. For instance, Rescher (1955: 10) cited the biologists’ use of ‘part’ for the functional subunits of an organism as a case in point. This is of little import, though. Taking reflexivity as constitutive of the meaning of ‘part’ amounts to regarding identity as a limit (improper) case of parthood. A stronger relation, whereby nothing counts as part of itself, can obviously be defined in terms of the weaker one, hence there is no loss of generality (see Section 1.2.2). Vice versa, one could frame a mereological theory by taking proper parthood as a primitive instead. This is merely a question of choosing a suitable primitive.

The transitivity principle, (17), is more controversial. Several authors have observed that many legitimate senses of ‘part’ are non-transitive. Examples would include: (i) a biological subunit of a cell is not a part of the organ of which that cell is a part; (ii) a handle can be part of a door

and the door of a house, though a handle is never part of a house; (iii) my finger is part of me and I am part of the team, yet my finger is not part of the team. (See again Rescher 1955, Cruse 1979, and Winston *et al.* 1987, respectively; for other examples see Iris *et al.* 1988, Moltman 1997, and Johansson 2004 *inter alia*). Arguably, however, such misgivings stem again from the ambiguity of ‘part’. What counts as a biological subunit of a cell may not count as a subunit, i.e., a *distinguished* part of the organ, but that is not to say that it is not part of the organ at all. Similarly, if there is a sense of ‘part’ in which a handle is not part of the house to which it belongs, or my finger not part of my team, it is a restricted sense: the handle is not a *functional* part of the house (a “component”), though it is a functional part of the door and the door a functional part of the house; my finger is not *directly* part of the team, though it is directly part of me and I am directly part of the team. It is obvious that if the interpretation of ‘part’ is narrowed by additional conditions (e.g., by requiring that parts make a functional or direct contribution to the whole), then transitivity may fail. In general, if x is a ϕ -part of y and y is a ϕ -part of z , x need not be a ϕ -part of z : the predicate modifier ‘ ϕ ’ may not distribute over parthood. But that shows the non-transitivity of ‘ ϕ -part’, not of ‘part’. And within a sufficiently general framework this can easily be expressed with the help of explicit predicate modifiers (Varzi 2005). In any event, it seems clear that *spatial* parthood is transitive: whether we construe this as a restricted notion or identify it with the general notion of parthood, (17) holds.

Finally, concerning the antisymmetry postulate (18), two sorts of worry are worth mentioning. On the one hand, some authors maintain that the relationship between an object and the stuff it is made of provides a perfectly ordinary example of symmetric parthood: according to Thomson (1998), for example, a statue and the clay it is made of are part of each other, yet distinct. This is highly controversial and there is a large philosophical literature devoted on this topic (see e.g. the papers in Rea 1997). For the moment, let us simply observe that the example trades once again on the ambiguity of ‘part’. We have already mentioned that material constitution is best regarded as a *sui generis*, non-mereological relation. Whether this relation may obtain between two spatially coincident objects is an interesting question, but we should postpone its discussion to where it belongs: the theory of spatial location (Section 3.3). On the other hand, one may wonder about the possibility of *unordinary* cases of symmetric parthood relationships. Sanford (1993: 222) refers to Borges’s *Aleph* as a case in point: “I saw the earth in the Aleph and in the earth the Aleph once more and the earth in the Aleph ...”. In this case, a plausible reply is simply that

fiction delivers no guidance to conceptual investigations: conceivability may well be a guide to possibility, but literary fantasy is by itself no evidence of genuine conceivability (van Inwagen 1993: 229). Still, one may observe that the possibility of mereological loops is not pure fantasy. In view of certain developments in non-well-founded set theory (Aczel 1988), one might indeed suggest building mereology on the basis of a notion of parthood that may violate (18). This is particularly significant insofar as set theory itself may be reformulated in mereological terms—a possibility that is explored in the works of Bunt (1985) and especially Lewis (1991). At present, however, no systematic study of non-well-founded mereology has been put forward in the literature. Moreover, we are interested here in mereology as a tool for spatial reasoning, and in this regard the possibility of symmetric loops does indeed appear to be pure fantasy. In the following we shall therefore confine ourselves to theories that accept the antisymmetry postulate along with reflexivity and transitivity: parthood is a partial ordering.

1.2.2 Other mereological concepts. It is convenient at this point to introduce some degree of formalization. Let us use ‘P’ for the binary predicate constant ‘... is part of ...’. Taking the underlying logic to be a standard predicate calculus with identity, the above minimal requisites on parthood may then be regarded as forming a first-order theory characterized by the following proper axioms for ‘P’:

- | | | |
|-------|------------------------------------|--------------|
| (P.1) | Pxx | Reflexivity |
| (P.2) | $Pxy \wedge Pyz \rightarrow Pxz$ | Transitivity |
| (P.3) | $Pxy \wedge Pyx \rightarrow x = y$ | Antisymmetry |

(Here and in the following we simplify notation by dropping all initial universal quantifiers. Unless otherwise specified, all formulas are to be understood as universally closed.) A number of additional mereological predicates can then be introduced by definition. For example:

- | | | |
|------|--|------------------|
| (19) | $EQxy =_{df} Pxy \wedge Pyx$ | Equality |
| (20) | $PPxy =_{df} Pxy \wedge \neg Pyx$ | Proper Parthood |
| (21) | $PExy =_{df} \neg Pxy \wedge Pyx$ | Proper Extension |
| (22) | $Oxy =_{df} \exists z(Pzx \wedge Pzy)$ | Overlap |
| (23) | $Uxy =_{df} \exists z(Pxz \wedge Pyz)$ | Underlap |

An intuitive model for these relations, with ‘P’ interpreted as spatial inclusion, is given in the diagram of Figure 1.1.

Note that ‘Uxy’ is bound to hold if we assume the existence of a “universal entity” of which everything is part. Conversely, ‘Oxy’ would always hold if we assumed the existence of a “null entity” that is part




Figure 1.1. Basic mereological relations. (Shaded cells indicate parthood.)

of everything. In the domain of spatial entities, the latter assumption is of course implausible (Geach 1949).⁴ The former assumption may be challenged, too (Simons 2003, Varzi 2006), but it seems reasonable, if not obvious, in case the only spatial entities countenanced by the theory are regions of space. We shall come back to these issues in Section 1.4.

Note also that the definitions imply (by pure logic) that EQ, O, and U are reflexive and symmetric; in addition, EQ is also transitive—an equivalence relation. By contrast, PP and PE are irreflexive and asymmetric, and it follows from (P.2) that both are transitive. Since the following biconditional is also a straightforward consequence of the axioms:

$$(24) \quad Pxy \leftrightarrow (PPxy \vee x = y)$$

it should now be obvious that one could in fact use proper parthood as an alternative starting point for the development of mereology, using the right-hand side of (24) as a definiens for ‘P’. This is, for instance, the option followed in Simons (1987), where the partial ordering axioms for ‘P’ are replaced by the strict ordering axioms for ‘PP’:

$$(25) \quad \neg PPxx$$

$$(26) \quad PPxy \wedge PPyz \rightarrow PPxz$$

$$(27) \quad PPxy \rightarrow \neg PPyx$$

Ditto for ‘EP’, which was in fact the primitive relation in Whitehead’s (1919) semi-formal treatment of the mereology of events. Other options may be considered, too. For example, Goodman (1951) used ‘O’ as a primitive and Leonard and Goodman (1940) used its opposite:

$$(28) \quad Dxy =_{df} \neg Oxy \qquad \text{Disjointness}$$

⁴In other contexts one may feel differently: see Martin (1965) and Bunt (1985) for theories with a null individual, and Bunge (1966) for a theory with *several* null individuals.

However, the relations corresponding to such predicates are weaker than PP and PE and no biconditional is provable from (P.1)–(P.3) that would yield a corresponding definiens of ‘P’ (though one could of course define ‘P’ in terms of ‘O’ or ‘D’ in the presence of further axioms; see below *ad* (45)). Thus, other things being equal, ‘P’, ‘PP’, and ‘PE’ appear to be the only reasonable options. Here we shall stick to ‘P’.

Finally, note that Identity could itself be introduced by definition, due to the following corollary of the antisymmetry postulate (P.3):

$$(29) \quad x = y \leftrightarrow EQxy$$

Accordingly, the theory could be formulated in a pure first-order language by assuming (P.1) and (P.2) and replacing (P.3) with the following variant of the standard axiom schema for ‘=’ (where ϕ is any formula):

$$(P.3') \quad EQxy \rightarrow (\phi x \leftrightarrow \phi y) \qquad \text{Indiscernibility}$$

One may in fact argue on these grounds that parthood is in some sense conceptually prior to identity (as in Sharvy 1983: 234), and since ‘EQ’ is not definable in terms of ‘PP’ or ‘PE’ without resorting to ‘=’, the argument would also provide evidence in favor of ‘P’ as the most basic primitive. As we shall see in Section 1.3.2, however, the link between parthood and identity is philosophically problematic. In order not to compromise our discussion, in the following we shall therefore continue to work with a language with both ‘P’ and ‘=’ as primitives.

1.3 Decomposition principles

Let M be the theory defined by the three basic principles (P.1)–(P.3). M may be viewed as embodying the common core of any mereological theory. Not just any partial ordering qualifies as a part-whole relation, though, and deciding what further principles should be added to (P.1)–(P.3) is precisely the question a good mereological theory is meant to answer. It is here that philosophical issues begin to arise.

Generally speaking, such refinements may be divided into two main groups. On the one hand, one may extend M by means of *decomposition principles* that take us from a whole to its proper parts. For example, one may consider the idea that whenever something has a proper part, it has more than one—i.e., that there is always some *mereological difference* (a remainder) between a whole and its proper parts. This need not be true in every model for M : a world with only two items, only one of which is part of the other, would be a counterexample, though not one that could be illustrated with the sort of geometric diagram used in Figure 1.1. On the other hand, one may extend M by means of *composition principles*




Figure 1.2. Three unsupplemented models. (Parthood relationships are represented by connecting lines going uphill.)

that go in the opposite direction—from the parts to the whole. For example, one may consider the idea that whenever there are some things there exists a whole that consists exactly of those things—i.e., that there is always a *mereological sum* (or fusion) of two or more parts. Again, this need not be true in a model for M , and it is a matter of controversy whether the idea should hold unrestrictedly.

1.3.1 Parts and remainders. Let us begin with the first sort of extension. And let us start by taking a closer look at the intuition according to which a whole cannot be decomposed into a single proper part. There are various ways in which one can try to capture this basic intuition. Consider the following (from Simons 1987: 26–28):

(P.4 _a)	$PPxy \rightarrow \exists z(PPzy \wedge \neg z = x)$	Weak Company
(P.4 _b)	$PPxy \rightarrow \exists z(PPzy \wedge \neg Pzx)$	Strong Company
(P.4)	$PPxy \rightarrow \exists z(Pzy \wedge \neg Ozx)$	Supplementation

The first principle, (P.4_a), is a literal rendering of the idea in question: every proper part must be accompanied by another. However, there is an obvious sense in which (P.4_a) only captures the letter of the idea, not the spirit: it rules out the unintended model mentioned above (see Figure 1.2, left) but not, for example, an implausible model with an infinitely descending chain in which the additional proper parts do not leave any remainder (Figure 1.2, center).

The second principle, (P.4_b), is stronger: it rules out both models as unacceptable. However, (P.4_b) is still too weak to capture the intended idea. For example, it is satisfied by a model in which a whole can be decomposed into several proper parts all of which overlap one another (Figure 1.2, right), and it is may be argued that such models do not do justice to the meaning of ‘proper part’: after all, the idea is that the removal of a proper part should leave a remainder, but it is by no means clear what would be left of z once x (along with its parts) is removed.

It is only the third principle, (P.4), that appears to provide a full formulation of the idea that nothing can have a single proper part. According to this principle, every proper part must be “supplemented” by another, *disjoint* part, and it is this last qualification that captures the notion of a remainder.

Should (P.4) be incorporated into M as a further fundamental principle on the meaning of ‘part’? Most authors (beginning with Simons himself) would say so. Yet here there is room for disagreement. In fact, it is not difficult to conceive of mereological scenarios that violate not only (P.4), but also (P.4_b) and even (P.4_a). For example, in Brentano’s (1933) theory of accidents, a soul is a proper part of a thinking soul even though there is nothing to make up for the difference (see Chisholm 1978; Baumgartner and Simons 1994). Similarly, in Fine’s (1982) theory of *qua* objects, every basic object (John) qualifies as the only proper part of its incarnations (John *qua* philosopher, John *qua* husband, etc.). Now, such putative counterexamples are controversial and, more importantly for our present concerns, they appear to be of little significance if mereology is to be thought of as a theory of space. The spatial relations illustrated by our initial examples (1)–(7) all seem to satisfy (P.4) and, *a fortiori*, (P.4_a) and (P.4_b). Nonetheless there are counterexamples also in the realm of truly spatial mereologies. The best illustration comes from Whitehead’s (1929) theory of extensive connection: on this theory, a topologically closed region includes its open interior as a proper part in spite of there being no boundary elements to distinguish them—the domain only consists of extended regions. Whether the omission of boundary elements such as points, lines, and surfaces is a reasonable thing to do when it comes to the task of modeling our understanding of space, and whether in the absence of such elements the distinction between open and closed regions is still legitimate, are questions that every theory of space must of course address. In Section 2.4 we shall see that answering in the affirmative involves serious philosophical and technical complications. But we shall also see that several theories are available to do the job, including theories that occupy a prominent role in the current literature on qualitative spatial reasoning. One may rely on the intuitive appeal of (P.4) to discard such theories as implausible, but one may as well turn things around and regard the adequacy of such theories as a good reason not to accept (P.4) unrestrictedly. As things stand, it therefore seems appropriate to regard such a principle as providing a minimal but substantive addition to (P.1)–(P.3), one that goes beyond the mere lexical characterization of ‘part’ provided by M . For future reference, let us label the resulting mereological theory MM (for *Minimal Mereology*).



Figure 1.3. A weakly supplemented model.

1.3.2 Supplementation, extensionality, identity. There is another way of expressing the supplementation intuition that is worth considering. It corresponds to the following axiom, which differs from (P.4) in the antecedent:

$$(P.5) \quad \neg P y x \rightarrow \exists z (P z y \wedge \neg O z x) \quad \text{Strong Supplementation}$$

Intuitively, this says that if an object *fails* to include another among its parts, then there must be a remainder. It is easily seen that (P.5) implies (P.4), so any theory rejecting at least (P.4) will *a fortiori* reject (P.5). (For instance, on Whitehead's boundary-free theory of extensive connection, a closed region is not part of its interior even though each part or the former overlaps the latter.) However, the converse does not hold. The diagram in Figure 1.3 illustrates a model in which (P.4) is true, since each proper part counts as a supplement of the other; yet (P.5) is false.

The theory obtained by adding (P.5) to (P.1)–(P.3) is thus a proper extension of *MM*. Let us label this stronger theory *EM*, for *Extensional Mereology*, the attribute ‘extensional’ being justified precisely by the exclusion of countermodels that, like the one just mentioned, contain distinct objects with the same proper parts. In fact, the following is a theorem of *EM*:

$$(30) \quad \exists z P P z x \rightarrow (\forall z (P P z x \rightarrow P P z y) \rightarrow P x y)$$

from which it follows that no composite objects with the same proper parts can be distinct:

$$(31) \quad (\exists z P P z x \vee \exists z P P z y) \rightarrow (x = y \leftrightarrow \forall z (P P z x \leftrightarrow P P z y))$$

(The analogue for ‘P’ is, of course, already provable in *M*, since part-hood is reflexive and antisymmetric.) Thus, *EM* is truly an extensional theory incorporating the view that an object is exhaustively defined by its constituent parts. This goes far beyond the intuition that lies behind the weak supplementation principle (P.4). Does it go too far?

On the face of it, it is not difficult to envisage scenarios that would correspond to the diagram in Figure 1.3. For example, we can obtain a counterexample to (P.5) by identifying *x* and *y* with the sets $\{\{z_1\}, \{z_1, z_2\}\}$

and $\{\{z_2\}, \{z_1, z_2\}\}$ (i.e., with the ordered pairs $\langle z_1, z_2 \rangle$ and $\langle z_2, z_1 \rangle$, respectively), interpreting ‘P’ as the ancestral of the improper membership relation (i.e., of the union of \in and $=$). But sets are abstract entities; can we also envisage similar scenarios in the spatial domain?

Here is a case where the answer may differ crucially depending on whether we are interested in modeling a domain of concrete spatial entities or just the domain of the regions of space that they occupy. In the latter case there is little room for controversy: spatial regions are extensional, if anything is, unless of course we favor a Whiteheadian conception of space. In the former case, however, the answer is controversial. There are two sorts of objection worth considering. On the one hand, it is sometimes argued that sameness of parts is not *sufficient* for identity, as some entities may differ exclusively with respect to the arrangement of their parts. For example, it is sometimes argued that: (i) two words can be made up of the same letters, as with ‘fallout’ and ‘outfall’; (ii) the same flowers can compose a nice bunch or a scattered bundle, depending on the arrangements of the individual flowers; (iii) a cat can survive the annihilation of its tail, but the amount of feline tissue consisting of the cat’s tail and the rest of the cat’s body cannot survive the annihilation of the tail, hence they have different properties and must be distinct by Leibniz’s law in spite of their sharing exactly the same ultimate mereological constituents. (See Hempel 1953: 110, Eberle 1970: §2.10, and Wiggins 1968, respectively; variants of (iii) may also be found in Doepke 1982, Lowe 1989, Johnston 1992, and Baker 1999, *inter alia*.) On the other hand, it is sometimes argued that sameness of parts is not *necessary* for identity, as some entities may survive mereological change. If a cat survives the annihilation of its tail, then the tailed cat (before the accident) and the tailless cat (after the accident) are one and the same in spite of their having different proper parts (Wiggins 1980). If any of these arguments is accepted, then clearly (31) is too strong a principle to be imposed on the parthood relation. And since (31) follows from (P.5), it might be concluded that *EM* is on the wrong track.

Let us look at these objections separately. Concerning the necessity of mereological extensionality, i.e., the left-to-right conditional in the consequent of (31):

$$(32) \quad x = y \rightarrow \forall z(PPzx \leftrightarrow PPzy)$$

it is perhaps enough to remark that the difficulty is not peculiar to extensional mereology. The objection proceeds from the consideration that ordinary entities such as cats and other living organisms (and possibly other entities as well, such as cars and houses) survive all sorts of gradual mereological changes. Yet the same can be said of other types of change

as well: bananas ripen, houses deteriorate, people sleep at night and eat at lunch. How can we say that they are the same things, if they are not quite the same? Indeed, (32) is just an instance of the identity axiom

$$(ID) \quad x = y \rightarrow (\phi x \leftrightarrow \phi y)$$

and it is well known that this axiom calls for revisions when '=' is given a diachronic reading. Arguably, any such revisions will affect the case at issue as well, and in this sense the above-mentioned objection to (32) can be disregarded. For example, if the basic parthood predicate were reinterpreted as a time-indexed relation (Thomson 1983), then the problem would disappear as the tensed version of (P.5) would only warrant the following variant of (32):

$$(32') \quad x = y \rightarrow \forall t \forall z (PP_t z x \leftrightarrow PP_t z y)$$

Similarly, the problem would disappear if the variables in (32) were taken to range over four-dimensional entities whose parts may extend in time as well as in space (Heller 1984, Sider 2001), or if identity itself were construed as a contingent relation that may hold at some times but not others (Gallois 1998). Such revisions may be regarded as an indicator of the limited ontological neutrality of extensional mereology. But their independent motivation also bears witness to the fact that controversies about the necessity of extensionality stem from larger and more fundamental philosophical conundrums and cannot be assessed by appealing to our intuitions about the meaning of 'part'.

The worry about the sufficiency of mereological extensionality, i.e., the right-to-left conditional in the consequent of (31):

$$(33) \quad \forall z (PP z x \leftrightarrow PP z y) \rightarrow x = y$$

is more to the point. However, there are various ways of resisting such counterexamples as (i)–(iii) on behalf of *EM*. Consider (i)—two words made up of the same letters. Insofar as we are dealing with truly spatial entities, this is best described as a case of different word *tokens* made up of distinct tokens of the same letter *types*. There is, accordingly, no genuine violation of (33) in the opposition between 'fallout' and 'outfall' (for instance), hence no reason to reject (P.5) on these grounds. (Besides, even with respect to abstract types, it could be pointed out that 'fallout' and 'outfall' do not share *all* their proper parts: the string 'lo', for instance, is only included in the first word.) What if one of the two word tokens is obtained from the other by rearranging the same letter *tokens*? In that case the reply misfires, but so does the objection: the issue becomes once again one of diachronic non-identity, with all that it entails (Lewis 1991: 78f).

Case (ii)—the flowers—is not significantly different. The same, concrete flowers cannot compose a nice bunch and a scattered bundle *at the same time*. Case (iii), however, is more delicate. There is a strong intuition that a cat really is something over and above the amount of feline tissue consisting of its tail and the rest of its body—that they have different survival conditions and, hence, different properties—so it may be thought that here we have a genuine counterexample to mereological extensionality. On behalf of *EM*, it should nonetheless be noted that the appeal to Leibniz’s law in this context is debatable. Let ‘Tibbles’ name our cat and ‘Tail’ its tail, and let us grant the truth of

(34) Tibbles can survive the annihilation of Tail.

There is, indeed, an intuitive sense in which the following is also true:

(35) The amount of feline tissue consisting of Tail and the rest of Tibbles’s body cannot survive the annihilation of Tail.

However, this intuitive sense corresponds to a *de dicto* reading of the modality, where the description in (35) has narrow scope:

(35a) Necessarily, the amount of feline tissue consisting of Tail and the rest of Tibbles’s body has Tail as a proper part.

On this reading (35) is hardly negotiable (in fact, logically true). Yet this is irrelevant in the present context, for (35a) does not amount to an ascription of a modal property and cannot be used in connection with Leibniz’s law. (Compare the following fallacious argument: The number of planets might have been even; 9 is necessarily odd; hence the number of planets is not 9.) On the other hand, consider a *de re* reading of (35):

(35b) The amount of feline tissue consisting of Tail and the rest of Tibbles’s body necessarily has Tail as a proper part.

On this reading the appeal to Leibniz’s law would be legitimate (modulo any concerns about the status of modal properties) and one could rely on the truth of (34) and (35), i.e., (35b), to conclude that Tibbles is distinct from the relevant amount of feline tissue. However, there is no obvious reason why (35) should be regarded as true on this reading. That is, there is no obvious reason to suppose that the amount of feline tissue that in the actual world consists of Tail and the rest of Tibbles’s body—that amount of feline tissue that is now resting on the carpet—cannot survive the annihilation of Tail. Indeed, it would appear that any reason in favor of this claim *vis-à-vis* the truth of (34) would have to *presuppose*




Figure 1.4. A strongly supplemented model violating complementation

the distinctness of the entities in question, so no appeal to Leibniz’s law would be legitimate to *establish* the distinctness on pain of circularity (Varzi 2000). This is not to say that the putative counterexample to (34) is wrong-headed. But it requires genuine metaphysical work to defend it and it makes the rejection of the strong supplementation principle (P.5) a much harder task.

1.3.3 Complementation. There is a way of expressing the supplementation intuition that is even stronger than (P.5). It corresponds to the following thesis, which differs from (P.5) in the consequent:

$$(P.6) \quad \neg P y x \rightarrow \exists z \forall w (P w z \leftrightarrow (P w y \wedge \neg O w x)) \quad \text{Complementation}$$

This says that if y is not part of x , there exists something that comprises exactly those parts of y that are disjoint from x —something that we may call the *difference* or *relative complement* between y and x . It is easily checked that this principle implies (P.5). On the other hand, the diagram in Figure 1.4 shows that the converse does not hold: there are two parts of y disjoint from x , namely z_1 and z_2 , but there is nothing that consists exactly of such parts, so we have a model of (P.5) in which (P.6) fails.

Any misgivings about (P.5) may of course be raised against (P.6). But what if we agree with the above arguments in support of (P.5)? Do they also give us reasons to accept the stronger principle (P.6)? The answer is in the negative. Plausible as it may sound, (P.6) has consequences that even an extensionalist may not be willing to accept. For example, Figure 1.5 depicts a scenario that—it may be argued—corresponds exactly to the model of Figure 1.4. It may be argued that although x and z_1 jointly constitute a larger part of y (the difference between y and z_2), and similarly for x and z_2 (the difference between y and z_1), there is nothing consisting of z_1 and z_2 (the difference between y and x), since these two pieces are disconnected. More generally, it appears that (P.6) would force us to accept the existence of scattered entities, such as the “sum” of your left and right arms, or the “sum” of Canada and Mex-




Figure 1.5. Does the difference between y and x exist?

ico, and since Lowe (1953) many authors have objected to this thought *regardless* of how one feels about extensionality. (One philosopher who explicitly agrees to extensionality while distrusting scattered entities is Chisholm 1987.) As it turns out, the extra strength of (P.6) is therefore best appreciated in terms of the sort of mereological aggregates that this principle would entail, aggregates that are composed of two or more parts of a given whole. This suggests that any additional misgivings about (P.6), besides its extensional implications, are truly misgivings about matters of composition. We shall accordingly postpone their discussion to Section 1.4, where we shall attend to these matters more fully. For the moment, let us simply say that (P.6) is, on the face of it, not a principle that can be added to M without further argument.

1.3.4 Atomism and other options. One last important family of decomposition principles concerns the question of atomism. Mereologically, an atom (or “simple”) is an entity with no proper parts:

$$(36) \quad Ax =_{df} \neg \exists y PPyx \quad \text{Atom}$$

Are there any such entities? And if there are, is everything entirely made up of atoms? Does everything comprise at least some atoms? Or is everything made up of atomless gunk? These are deep and difficult questions, which have been the focus of philosophical investigation since the early days of philosophy and have been center stage also in many recent disputes in mereology (see, for instance, van Inwagen 1990, Sider 1993, Zimmerman 1996a, and the papers collected in Hudson 2004). Here we shall confine ourselves to a brief examination.

The two main options, to the effect that there are no atoms at all, or that everything is ultimately made up of atoms, correspond to the following postulates, respectively:

$$(P.7) \quad \exists y PPyx \quad \text{Atomlessness}$$

$$(P.8) \quad \exists y (Ay \wedge P yx) \quad \text{Atomcity}$$

These postulates are mutually incompatible, but taken in isolation they can consistently be added to any mereological theory X considered here.

Adding (P.8) yields a corresponding *Atomistic* version, AX ; adding (P.7) yields an *Atomless* version, $\bar{A}X$. Since finitude together with the anti-symmetry of parthood (P.3) jointly imply that mereological decomposition must eventually come to an end, it is clear that any finite model of M (and *a fortiori* of any extension of M) must be atomistic. Accordingly, an atomless mereology $\bar{A}X$ admits only models of infinite cardinality. (A world containing such wonders as Borges’s Aleph, where parthood is not antisymmetric, might by contrast be finite and yet atomless.) An example of such a model, establishing the consistency of the atomless version of most mereological theories considered in this chapter, is provided by the regular open sets of a Euclidean space, with ‘P’ interpreted as set-inclusion (Tarski 1935).

Now, one thing to notice is that, independently of their motivations, atomistic mereologies admit of significant simplifications in the axioms. For instance, AEM can be simplified by replacing (P.5) and (P.8) with

$$(P.5') \quad \neg P y x \rightarrow \exists z (A z \wedge P z y \wedge \neg P z x)$$

which in turn implies the following atomistic variant of the extensionality thesis (31):

$$(37) \quad x = y \leftrightarrow \forall z (A z \rightarrow (P z x \leftrightarrow P z y))$$

Thus, any atomistic extensional mereology is truly “hyperextensional” in Goodman’s (1958) sense: things built up from exactly the same atoms are identical. An interesting question, discussed at some length in the late 1960’s (Yoes 1967, Eberle 1968, Schuldenfrei 1969) and taken up more recently by Simons (1987: 44f) and Engel and Yoes (1996), is whether there are atomless analogues of (37). Is there any predicate that can play the role of ‘A’ in an atomless mereology? Such a predicate would identify the “base” of the system and would therefore enable mereology to cash out Goodman’s hyperextensional intuitions even in the absence of atoms. This question is particularly significant from a nominalistic perspective, but it also bears on our present concerns. For example, it is a relevant question to ask in connection with the Whiteheadian conception mentioned in Section 1.3.1, according to which space contains no parts of lower dimensions such as points or lines (see Forrest 1996, Roeper 1997). In special cases there is no difficulty in providing a positive answer. For example, in the $\bar{A}EM$ model consisting of the open regular subsets of the real line, the open intervals with rational end points form a base in the relevant sense. It is unclear, however, whether a general answer can be given that applies to any sort of domain. If not, then the only option would appear to be an account where the notion of

a “base” is relativized to entities of a given sort. In Simons’s terminology, we could say that the ψ -ers form a base for the ϕ -ers if and only if the following variants of (P.5’) and (P.8) are satisfied:

$$\begin{aligned} \text{(P.5}^*) \quad & \phi x \wedge \phi y \rightarrow (\neg P y x \rightarrow \exists z(\psi z \wedge P z y \wedge \neg P z x)) \\ \text{(P.8}^*) \quad & \phi x \rightarrow \exists y(\psi y \wedge P y x) \end{aligned}$$

An atomistic mereology would then correspond to the limit case where ‘ ψ ’ is identified with ‘A’ for every choice of ‘ ϕ ’. In an atomless mereology, by contrast, the choice of the base would depend each time on the level of “granularity” set by the relevant specification of ‘ ϕ ’.

A second important consideration concerns the possibility of theories that lie between the two extreme options afforded by Atomicity and Atomlessness. For instance, it can be held that there are atoms, though not everything need have a complete atomic decomposition, or it can be held that there is atomless gunk, though not everything need be gunky (Zimmerman 1996a). Again, formally this amounts to endorsing a restricted version of either (P.7) or (P.8) in which the variables are suitably restricted so as to range over entities of a certain sort:

$$\begin{aligned} \text{(P.7}_\phi) \quad & \phi x \rightarrow \exists y P P y x \\ \text{(P.8}_\phi) \quad & \phi x \rightarrow \exists y (A y \wedge P y x) \end{aligned}$$

At present, no thorough formal investigation of such options has been entertained (but see Masolo and Vieu 1999). Yet the issue is particularly significant from the perspective of a mereological theory aimed at modeling the spatial world, especially if the theory is to countenance concrete spatial entities along with the regions of space that such entities may occupy. It is, after all, a plausible thought that while the question of atomism may be left open with regard to the mereological structure of material objects (pending empirical findings from physics, for example), it must receive a definite answer with regard to the structure of space itself. This would amount to endorsing a version of either (P.7_φ) or (P.8_φ) in which ‘ ϕ ’ is understood as a condition that is satisfied exclusively by regions of space. Such a condition, of course, cannot be formulated in the language of a purely mereological theory, but we shall see in Section 3 that a suitably enriched theory, in which the relation of location is explicitly articulated, can do the job properly. (Actually, it is hard to conceive of a world in which an atomistic space is inhabited by entities that can be decomposed indefinitely, so in this case it is reasonable to suppose that any theory accepting (P.8_φ) for regions would also endorse the stronger principle (P.8). However, (P.7_φ) would be genuinely independent of (P.7), unless it is assumed that every mereologically atomic entity should also be spatially atomic, i.e., unextended.)

Similar considerations apply to other decomposition principles that may come to mind at this point. For example, one may consider a requirement to the effect that ‘PP’ forms a dense ordering, as already Whitehead (1919) had it:

$$(P.9) \quad PPxy \rightarrow \exists z(PPxz \wedge PPzy) \quad \text{Density}$$

As a general decomposition principle, (P.9) might be deemed too strong, especially in an atomistic setting. However, it is plausible to suppose that (P.9) should hold at least in the domain of spatial regions, regardless of whether these are construed as atomless gunk or as aggregates of spatial atoms. Evidently much depends on the link one establishes between the mereology of an object and that of its spatial location and this, again, is a question to which we attend more fully in Section 3. For the moment, let us simply observe that the sort of philosophical issues that lie behind these options is significantly different from those considered in the previous sections. Whether something can have a single proper part, whether parthood is extensional, or even whether it satisfies the complementation principle (P.6) are issues that depend greatly on our understanding of the parthood relation. They are, in an important sense, conceptual questions. Whether there are mereological atoms, by contrast, or whether mereological decomposition should obey a density principle, are substantive questions that have nothing to do with our understanding of parthood as such. (For more on these questions, and on their general historical background, see Pyle 1995 and Holden 2004.)

1.4 Composition principles

Let us now consider the second way of extending M mentioned at the beginning of Section 1.3. Just as we may want to fix the logic of P by means of decomposition principles that take us from a whole to its proper parts, we may look at composition principles that go in the opposite direction—from the parts to the whole. More generally, we may consider the idea that the domain of the theory ought to be closed under mereological operations of various sorts: not only fusions, but also products, differences, and more. Here, again, there is room for several philosophical considerations, some of which are particularly indicative of the tension between space-oriented and object-oriented intuitions.

1.4.1 Bounds and fusions. Conditions on composition are many. Beginning with the weakest, consider the claim that any two suitably related entities have an upper bound, i.e., underlap:

$$(P.10_\psi) \quad \psi xy \rightarrow \exists z(Pxz \wedge Pzy) \quad \text{Boundedness}$$




Figure 1.6. A fusion_a that is not a fusion, and a fusion that is not a fusion_b.

Exactly how ‘ ψ ’ should be construed is an important question by itself—a version of what van Inwagen (1990) calls the “special composition question”. Perhaps the most natural choice is to identify ψ with mereological overlap, the rationale being that such a relation establishes an important tie between what may count as two distinct parts of a larger whole. As we shall see momentarily, with ψ so construed, (P.10 _{ψ}) is indeed uncontroversial. However, regardless of any specific choice, it is apparent that (P.10 _{ψ}) is pretty weak, as it holds trivially in any domain with a universal entity of which everything is part.

A somewhat stronger condition would be to require any pair of suitably related entities to have a *smallest* underlapper—something composed exactly of them and nothing else. This requirement is sometimes stated by saying that such entities must have a mereological “sum”, or “fusion”, though it is not immediately obvious how that should be formulated in the formal language. Consider:

$$\begin{array}{ll}
 \text{(P.11}_{\psi a}\text{)} & \psi xy \rightarrow \exists z(\mathbf{P}xz \wedge \mathbf{P}yz \wedge \forall w(\mathbf{P}xw \wedge \mathbf{P}yw \rightarrow \mathbf{P}zw)) & \text{Fusion}_a \\
 \text{(P.11}_{\psi b}\text{)} & \psi xy \rightarrow \exists z(\mathbf{P}xz \wedge \mathbf{P}yz \wedge \forall w(\mathbf{P}wz \rightarrow \mathbf{O}wx \vee \mathbf{O}wy)) & \text{Fusion}_b \\
 \text{(P.11}_{\psi}\text{)} & \psi xy \rightarrow \exists z\forall w(\mathbf{O}wz \leftrightarrow \mathbf{O}wx \vee \mathbf{O}wy) & \text{Fusion}
 \end{array}$$

In a way, (P.11 _{ψa}) would seem the obvious choice, corresponding to the idea that the fusion of two objects is just their least upper bound relative to \mathbf{P} . (See e.g. Bostock 1979, van Benthem 1983.) However, this condition is too weak to capture the intended notion of a mereological fusion. For example, with ψ construed as overlap, (P.11 _{ψa}) is satisfied by the model of Figure 1.6, left: here the least upper bound of x and y is z , yet z hardly qualifies as something “made up” of x and y since its parts also include a third, disjoint item w . In fact, it is a simple fact about partial orderings that among finite models (P.11 _{ψa}) is equivalent to (P.10 _{ψ}), hence just as weak. By contrast, (P.11 _{ψb}) corresponds to a notion of fusion (to be found e.g. in Tarski 1929) that may seem too strong: it rules out the model on the left of Figure 1.6; but it also rules out the model on the right, which depicts a situation in which z may be viewed as an entity truly made up of x and y insofar as it is ultimately

composed of atoms to be found either in x or in y . Of course, such a situation violates the strong supplementation principle (P.5), but that's precisely the sense in which (P.11 $_{\psi b}$) is too strong: an anti-extensionalist might want to have a notion of fusion that does not presuppose strong supplementation. The formulation in (P.11 $_{\psi}$) is the natural compromise: it is strong enough to rule out the model on the left, but weak enough to be compatible with the model on the right. This is, in fact, the formulation that best reflects the notion of fusion to be found in standard treatments of mereology, and in the sequel we shall mostly stick to it. Note, however, that *if* (P.5) holds, then (P.11 $_{\psi}$) is equivalent to (P.11 $_{\psi b}$). Moreover, it turns out that if the stronger complementation axiom (P.6) holds, then all of these principles are trivially satisfied in any domain in which there is a universal entity: in that case, regardless of ψ , the fusion of any two entities is just the complement of the difference between the complement of one minus the other. (Such is the strength of (P.6), a genuine cross between decomposition and composition principles.)

We can further strengthen these principles by considering infinitary bounds and fusions. For example, (P.10 $_{\psi}$) can be generalized to a principle to the effect that any non-empty set of entities satisfying a suitable condition ξ has an upper bound. Strictly speaking there is a difficulty in expressing such a principle in a language without set variables. We can, however, achieve a sufficient degree of generality by relying on an axiom schema where classes are identified by open formulas. Since an ordinary first-order language has a denumerable supply of formulas, at most denumerably many sets (in any given domain) can be specified in this way. But for most purposes this limitation is negligible, as normally we are only interested in those sets of objects or regions that we *are* able to specify. Thus, the following axiom schema will do, where ' ϕ ' is any formula in the language and ' ξ ' expresses the condition in question:

$$(P.12_{\xi}) \quad \exists w \phi w \wedge \forall w (\phi w \rightarrow \xi w) \rightarrow \exists z \forall w (\phi w \rightarrow P w z)$$

Likewise, the fusion axiom (P.11 $_{\psi}$) can be strengthened as follows:

$$(P.13_{\xi}) \quad \exists w \phi w \wedge \forall w (\phi w \rightarrow \xi w) \rightarrow \exists z \forall w (\mathbf{O} w z \leftrightarrow \exists v (\phi v \wedge \mathbf{O} w v))$$

and similarly for (P.11 $_{\psi a}$) and (P.11 $_{\psi b}$). (The condition ' $\exists w \phi w$ ' guarantees that ' ϕ ' picks out a non-empty set, so there is no danger of asserting the unconditional existence of "null entities"—a mereological fiction that we have already mentioned as implausible in the context of spatial ontology.) It can be checked that these generalized formulations include the corresponding finitary principles as special cases, taking ' ϕw ' to be the formula ' $(w = x \vee w = y)$ ' and ' ξw ' the condition ' $(w = x \rightarrow \psi w y) \wedge (w = y \rightarrow \psi x w)$ '.

Finally, we get the strongest version of these composition principles by asserting them as axiom schemas that hold for *every* condition ξ , i.e., effectively, by foregoing any reference to ξ altogether. Formally this amounts in each case to dropping the second conjunct of the antecedent. For example, the following schema is the unrestricted version of (P.13 $_{\xi}$), to the effect that every specifiable non-empty set of entities has a fusion:

$$(P.13) \quad \exists w\phi w \rightarrow \exists z\forall w(\mathbf{O}wz \leftrightarrow \exists v(\phi v \wedge \mathbf{O}wv)) \quad \text{Unrestricted Fusion}$$

The extension of *EM* obtained by adding every instance of this schema has a distinguished pedigree and is known as *General Extensional Mereology*, or *GEM*. It corresponds to the classic systems of Leśniewski and of Leonard and Goodman. In fact, it turns out that adding (P.13) to *MM* yields the same theory *GEM*, since (P.13) implies that every pair of overlapping things has a maximal common part (a product):

$$(38) \quad \mathbf{O}xy \rightarrow \exists z\forall w(\mathbf{P}wz \leftrightarrow (\mathbf{P}wx \wedge \mathbf{P}wy))$$

which, in turn, implies the equivalence between the weak supplementation principle (P.4) and the stronger (P.5) (Simons 1987: 31). This is by itself remarkable, for it might be thought that a composition principle such as (P.13) should be compatible with the rejection of a decomposition principle that is committed to extensionality. On the other hand, mereological extensionality is really a double-barreled thesis: it says that two wholes cannot be decomposed *into* the same proper parts but also, by the same token, that two wholes cannot be composed *out of* the same proper parts. So it is not entirely surprising that as long as PP is well behaved, as per (P.4), extensionality might pop up like this in the presence of substantive composition statements. (It is, however, noteworthy that it pops up as soon as (P.4) is combined with a seemingly innocent thesis such as (38), so the anti-extensionalist should keep that in mind.)

The intuitive idea behind all these principles is in fact best appreciated in the presence of extensionality, for in that case the relevant fusions must be unique. Just to confine ourselves to *GEM*, it is natural to consider the following fusion operator (where ‘ v ’ is the definite descriptor⁵):

$$(39) \quad \Sigma x\phi x =_{df} \iota z\forall w(\mathbf{O}wz \leftrightarrow \exists v(\phi v \wedge \mathbf{O}wv)) \quad \text{fusion}$$

Then (P.13) and (P.5) can be simplified to a single axiom schema:

$$(P.14) \quad \exists x\phi x \rightarrow \exists z(z = \Sigma x\phi x) \quad \text{Unique Unrestricted Fusion}$$

⁵We assume a classical logical background, with ‘ ι ’ defined as usual. Much of what follows, however, would also apply in case a free logic were used instead, with ‘ ι ’ assumed as part of the logical vocabulary proper. (See Simons 1991b for a free formulation of mereology.)

and the full strength of the theory can be seen by considering that its models are all closed under the following functors, modulo the absence of a null entity:

$$\begin{array}{ll}
 (40) & x + y =_{df} \Sigma z(Pzx \vee Pzy) & \text{sum} \\
 (41) & x \times y =_{df} \Sigma z(Pzx \wedge Pzy) & \text{product} \\
 (42) & x - y =_{df} \Sigma z(Pzx \wedge Dzy) & \text{difference} \\
 (43) & \sim x =_{df} \Sigma zDzx & \text{complement} \\
 (44) & \mathcal{U} =_{df} \Sigma zPzz & \text{universe}
 \end{array}$$

(Absent the null entity, \mathcal{U} has no complement while products are defined only for overlapping pairs and differences for pairs that leave a remainder). Since these functors are the natural mereological analogue of the familiar Boolean operators, with fusion in place of set abstraction, it follows that the parthood relation axiomatized by *GEM* has the same properties as the set-inclusion relation. More precisely, it is isomorphic to the inclusion relation restricted to the set of all non-empty subsets of a given set, which is to say a complete Boolean algebra with the zero element removed—a fact that has been known since Tarski (1935).

There are other equivalent formulations of *GEM* that are noteworthy. For instance, it is a theorem of every extensional mereology that parthood amounts to inclusion of overlappers:

$$(45) \quad Pxy \leftrightarrow \forall z(Ozx \rightarrow Ozy)$$

This means that in an extensional mereology ‘O’ could be used as a primitive and ‘P’ defined accordingly, and it can be checked that the theory defined by postulating (45) together with the unrestricted fusion principle (P.13) and the antisymmetry axiom (P.3) is equivalent to *GEM*. Another elegant axiomatization of *GEM*, due to an earlier work of Tarski (1929), is obtained by taking just the transitivity axiom (P.2) and the unique unrestricted fusion axiom (P.14).

1.4.2 Composition, existence, and identity. Arguably, the algebraic strength of *GEM* speaks in favor of this theory as an account of the structure of space, since it is rather intuitive (and common practice) to understand spatial regions in terms of sets of points mereologically related by set-inclusion. As a general theory of the mereology of all spatial entities, however, *GEM* reflects substantive postulates whose philosophical underpinnings are controversial. Indeed, all composition principles turn out to be controversial, just as the decomposition principles examined in Section 1.3. For, on the one hand, it appears that the weaker, conditional formulations, from (P.10 _{ψ}) to (P.13 _{ξ}), are just not doing enough work: not only do they depend on the specification of

the limiting conditions expressed by the predicates ‘ ψ ’ and ‘ ξ ’; they also treat such conditions as merely sufficient for the existence of bounds and fusions, whereas ideally we are interested in conditions that are both sufficient and necessary. On the other hand, the stronger, unconditional formulations—most notably (P.13)—appear to go too far, not only because they tend to obliterate any difference between weak and strong supplementation, i.e., extensionality, but because they commit the theory to the existence of a large variety of *prima facie* implausible mereological composites. So what is the right way to go, at least insofar as we are interested in the compositional structure of the spatial realm?

Concerning the first sort of worry, one could of course strengthen every conditional formulation to a biconditional expressing both a necessary and sufficient condition for the existence of an upper bound, or a fusion. But then the question of how such conditions should be construed becomes crucial. For example, in connection with (P.10 _{ψ}) we have mentioned the idea of construing ‘ ψ ’ as ‘O’, the rationale being that mereological overlap establishes an important connection between what may count as two distinct parts of a larger whole. However, as a necessary condition overlap is arguably too stringent. We may have misgivings about the existence of scattered entities consisting of spatially unrelated parts, such as the top of my body and the bottom of yours, or the collection of my umbrellas and your left shoes. But in some cases no such misgivings arise. In some cases it appears perfectly natural to countenance wholes that are composed of two or more disjoint entities—a bikini, the solar system, a printed inscription consisting of separate letter tokens (Cartwright 1975). More generally, intuition and common sense suggest that some and only some mereological composites exist, not all; yet it is doubtful whether the question of *which* composites exist—van Inwagen’s “special composition question”—can be answered successfully. Consider a series of almost identical mereological aggregates that begins with a case where composition appears to obtain (e.g., the body cells that currently make up my body) and ends in a case where composition would seem not to obtain (e.g., the same body cells after their relative distance has been gradually increased to a huge extent). Where should we draw the line? It may well be that whenever some entities compose a bigger one, it is just a brute fact that they do so (Markosian 1998b). But if we are unhappy with brute facts, if we are looking for a principled way of drawing the line so as to specify the circumstances under which the facts obtain, then the question is truly challenging. As Lewis (1986: 213) put it, no restriction on composition can be vague, since existence cannot be a matter of degree; but unless it is vague, it cannot fit the intuitive *desiderata*.

For these reasons, although the axiom of unrestricted fusion has been a major source of complaint since the early days of mereology (see again Lowe 1953 and Rescher 1955, with replies in Goodman 1956, 1958), it is a fact that most formally accomplished theories accept unrestricted composition principles of some sort. Apart from whatever algebraic considerations might motivate them, such principles suggest themselves as the only *non-arbitrary* ways of answering the composition question. Besides, it might be observed that any complaints against such principles rest on psychological biases that have no bearing on how the world is actually structured. In the words of van Cleve (1986: 145), little would follow even if we did manage to come up “with a formula that jibed with all ordinary judgments” about what counts as a unit and what does not, for such judgments need not have any ontological transparency.

All of this speaks in favor of (P.13) and the like against their weaker, conditional formulations, providing also an answer to the second worry mentioned above: the *prima facie* ontological extravagance of a theory such as GEM is not by itself a sign that the theory has gone too far. There is, however, another worry that is worth mentioning in this connection, and this further worry concerns the ontological exuberance—if not the extravagance—of the theory. For even granting the impossibility of drawing a principled line between natural fusions and unnatural ones, one could still object that positing every conceivable fusion is utterly unjustified. Why should mereology be committed to the existence of all such things *over and above* their constituent proper parts?

There are two lines of response to this question. First, it could be observed that the ontological exuberance associated with the relevant composition principles is not substantive. This is obvious in the case of a modest principle in the spirit of (P.10 _{ψ}), to the effect that entities of the right sort have an upper bound. After all, there are small things (my fingers) and large things (my body), and it is just a fact that the latter encompass the former. But the same could be said with respect to those stronger principles that require the large thing to be composed *exactly* of the small things—to be their mereological fusion. For one could argue that even a fusion is, in an important sense, nothing over and above its constituent parts. The fusion is just the parts “taken together” (Lewis 1991: 81); it is the parts “counted loosely” (Baxter 1988: 580); it is, effectively, the same portion of reality, which is strictly a multitude and loosely a single thing. This thesis, known in the literature as “composition as identity”, is by no means undisputed (see e.g. van Inwagen 1994, Yi 1999, Merricks 1999). Nonetheless it should be carefully evaluated in connection with any worry about the ontological exuberance of fusion principles. And if the thesis is accepted, then the charge

of ontological extravagance loses its force, too. If a fusion is nothing over and above its constituent parts, and if the latter are all right, there can be nothing particularly extravagant in countenancing the former: it just is them.

Secondly, one may observe that the worry in question bites at the wrong level. If, given two entities, positing their sum were to count as further ontological commitment, then, given any mereologically composite entity, positing its proper parts should also count as further commitment. After all, every entity is distinct from its proper parts. But then the worry has nothing to do with the composition axioms; it is, rather, a question of whether there is any point in countenancing a whole along with its parts, or vice versa. And if the answer is in the negative, then there seems to be no use for mereology *tout court*. From the point of view of the present worry, it would seem that the only truly parsimonious account would be one that rejects, not only *some*, but *all* logically admissible fusions—in fact, all mereological composites whatsoever. Philosophically such an account is defensible (see Rosen and Dorr 2002) and the corresponding axiom is compatible with *M*:

$$(P.15) \quad Ax \qquad \text{Strong Atomicity}$$

The following immediate corollary, however, says it all: nothing would be part of anything else and parthood would collapse to identity.

$$(46) \quad Pxy \leftrightarrow x = y$$

(This account is known as mereological *nihilism*, in contrast to the mereological *universalism* expressed by (P.13); see van Inwagen 1990: 72ff.)

In recent years, further worries have been raised concerning mereological theories with non-trivial composition principles—especially concerning the full strength of *GEM*. It has been argued that unrestricted composition does not sit well with certain intuitions about persistence through time (van Inwagen 1990: 75ff), that it requires every entity to necessarily have the parts it has (Merricks 1999), or that it leads to paradoxes similar to the ones afflicting naive set theory (Bigelow 1996). Such arguments are still the subject of on-going controversy and a detailed examination is beyond the scope of this chapter. Some discussion of the first point, however, is already available in the literature: see especially Rea (1998), McGrath (1998), and Hudson (2001: 93ff). Hudson (2001: 95ff) also contains a discussion of the last point.

1.5 The problem of vagueness

Let us conclude this discussion of mereology by considering a question that is not directly related to specific mereological principles but, rather,

to the underlying notion of parthood that mereology seeks to systematize. All the theories examined so far, from M to GEM , assume that parthood is a perfectly determinate relation: given any two entities x and y , there is always a definite fact of the matter as to whether or not x is part of y . However, it may be argued that this is a simplification. Perhaps there is no room for vagueness in the idealized mereology of pure space, but what about the real world? Think of a cloud, a forest, a pile of trash. What parts do they have, exactly? What are the mereological boundaries of a desert, a river, a mountain? Some stuff is positively part of Mount Everest and some stuff is positively not part, yet there is borderline stuff whose mereological relationship to Everest seems indeterminate. Even living organisms may, on closer look, give rise to vagueness issues. Surely John's body comprises his heart and does not comprise mine. But what about the candy he is presently chewing: Is it part of John? Will it be part of John only after he swallowed it? After he started digesting it? After he digested it completely?

In the face of such examples, it might be thought that the conceptual apparatus on which M and its extensions are based is too rigid. It might be thought that the world includes various sorts of vague entities, and that relative to such entities the parthood relation need not be fully determined (van Inwagen 1990: ch. 13, Parsons and Woodruff 1995). There are, in fact, various ways one could seek greater flexibility. One could leave everything as is but change the underlying logic (and semantics), for instance by allowing statements of the form ' Pxy ' to receive no determinate truth-value (as in Tye 1990), or to receive truth-values that are intermediate between classical truth and falsity (as in Copeland 1995). Or one could change the very basic apparatus of mereology, replacing the 'part of' predicate with a new primitive 'part of *to a degree*': this is, for example, the approach that led to the development of Polkowsky and Skowron's (1994) "rough mereology", where parthood undergoes a fuzzification parallel to the fuzzification of membership in Zadeh's (1965) fuzzy set theory. No matter how exactly one proceeds, obviously many among the principles discussed above would have to be reconsidered, not because of what they say but because of their classical, bivalent presuppositions. For example, the extensionality theorem of EM , (31), says that composite things with the same proper parts are identical, and this would call for qualifications: the model in Figure 1.7, left, depicts x and y as non-identical by virtue of their having distinct determinate parts; yet one might prefer to describe a situation of this sort as one in which the identity between x and y is itself indeterminate, since it is indeterminate whether they really have distinct parts. Conversely, the model on the right depicts x and y as non-identical in spite of their having the



Figure 1.7. Objects with indeterminate parts (in grey)

same determinate proper parts; yet again one might prefer to suspend judgment owing to the indeterminacy the middle element.

That there are vague objects in this sense, however, i.e., objects whose mereological composition may to some extent be objectively indeterminate, is all but obvious. Surely a statement such as

(47) x is part of Everest

may lack a definite truth-value, if x lies somewhere in the borderline area. But—it could be argued—this need not be due to the way the world is. The indeterminacy of (47) may be due exclusively to semantic factors—not to the vagueness of Everest but to the vagueness of ‘Everest’. When the members of the Geodetic Office of India baptized the mountain after the name of their British founder, they simply did not specify exactly which parcel of land they were referring to (or which parcel of land constituted the mountain they meant to name).⁶ The referent of their term was vaguely fixed and, as a consequence, the truth conditions of a statement such as (47) are not fully determined; yet this is not to say that the stuff out there is mereologically vague. Each one of a large variety of slightly distinct parcels of land has an equal claim to the vaguely introduced name ‘Everest’, and each such thing has a perfectly precise mereological structure. To put it differently, a statement such as

(48) It is indeterminate whether x is part of Everest

admits of a *de re* reading, as in (48a), but also of a *de dicto* reading, as in (48b):

(48a) Everest is a y such that: it is indeterminate whether x is part of y .

(48b) It is indeterminate whether: Everest is a y such that x is part of y .

The first reading corresponds to the initial thought, to the effect that Everest’s parts are indeed indeterminate, with the consequence that

⁶That mountains are just parcels of land is, of course, a substantive assumption: an anti-extensionalist may want to deny it, as with Tibbles and the relevant amount of feline tissue (Section 1.3.2). On the ontology of topographic entities, see e.g. Smith and Mark (2003).

mereology ought to be revised as seen above. The second reading, by contrast, corresponds to the idea that it is the semantics of ‘Everest’ that is indeterminate, and there is no reason to suppose that this is due to some objective deficiency in the parthood relation—hence no reason to require revisions in the apparatus of mereology itself. (Ditto for the other cases mentioned above. The reason why it’s indeterminate whether a certain molecule is part of a cloud, a tree part of a forest, or the candy part of John, is not that such things are mereologically indeterminate; rather, on a *de dicto* understanding the indeterminacy lies entirely in our words, in the terms we use to pick out such things from a multitude of slightly distinct but perfectly determinate potential referents.)

If the semantic conception is accepted, then, the problem of vagueness dissolves. Or rather: it ceases to be an issue for mereology and it becomes a problem for semantics broadly understood—a problem that manifests itself in many contexts besides those under consideration. (How much money do you need to be *rich*? How slowly can you *run*? How *late* can I call you?) Again, there are many things one could do at this point. A favored option is afforded by so-called supervaluational semantics, whose first application to vagueness can be traced back to Fine (1975). According to such semantics, the truth-value of a sentence involving vague terms is a function of its truth-values under the admissible precisifications of those terms: the sentence is true if it is true under every precisification, false if false under every precisification, and indeterminate otherwise. Thus, if x is in the borderline area, then the indeterminacy of (47) is explained by the fact that among the many admissible ways of precisifying the term ‘Everest’, some would pick out a referent that extends far enough to include x among its parts whereas others would not, which is to say that (47) would be true on some but not all precisifications. By contrast, if x were clearly part of Everest given the way the name is used in ordinary circumstances, or if it were clearly not part of Everest, then (47) would have a definite truth-value, for every precisification would yield the same response (always true and always false, respectively). We need not go into the details here. But three things are worth noting.

First, none of this will have any impact on the mereological axioms considered so far. For those axioms are expressed as (implicitly) universally quantified formulas involving no singular terms except for variables, and variables are not the sort of expression that can suffer from the phenomenon of vagueness. Variables range over all entities included in the domain of quantification and pick out their values independently of any vagueness that may affect the non-logical vocabulary.

Second, any model that satisfies a given axiom or theorem satisfies also any substitution instance that can be obtained by replacing one

or more variables with corresponding names or descriptive terms. For example, the following sentence is a substitution instance of (P.1):

(49) Everest is part of Everest.

and it is easily verified that a supervaluational semantics will make (49) true in every model of M . For insofar as reflexivity is meant to hold for every entity in the domain, the truth of (49) is guaranteed no matter which entity we elect as the referent of ‘Everest’. Likewise, the following is a substitution instance of (31), the extensionality theorem of EM :

(50) As long as they are non-atomic, Everest and Sagarmatha are the same if and only if they have the same proper parts.

(‘Sagarmatha’ is the Nepalese translation of ‘Everest’, though there is no guarantee that both names admit of the same precisifications.) Again, it is easily verified that a supervaluational semantics will make (50) true in every EM model. For no matter how we precisify the terms ‘Everest’ and ‘Sagarmatha’ by tracing a precise boundary around their referents, the extensionality of parthood will guarantee that the referents coincide just in case their proper parts coincide too.

Finally, it is worth emphasizing that a supervaluational semantics is perfectly adequate to classical logic (Fine 1975, McGee 1997, Varzi 2001). For example, although it does not obey to the semantic principle of *bivalence*, as with various instances of

(51) Either ‘ x is part of Everest’ is true or ‘ x is part of Everest’ is false,

it certainly satisfies the logical law of *excluded middle*: any instance of

(52) Either x is part of Everest or x is not part of Everest

is bound to be true, for it is true on any precisification of ‘Everest’. Things would change, however, if the language were enriched by adding an explicit sentential operator to express indeterminacy. In that case, the following principle would give expression to the assumption that parthood admits of no objective borderline cases:

(P.16) It is determinate whether Pxy , Determinacy

though it is obvious that this principle may have invalid instances as soon as ‘ x ’ or ‘ y ’ is replaced by a vague singular term such as ‘Everest’, as in (48). At the moment, the logic of determinacy operators is an open area of research, especially in the context of a supervaluational

semantics. (See Keefe 2000, §7.4 and references therein.) It is, however, an important tool that any good theory of vagueness should countenance. And it is bound to play a significant role in any application of the theory to mereology and spatial reasoning broadly understood.

2. Philosophical issues in topology

Let us now move on to the second major ingredient of a comprehensive theory of spatial reasoning—topology. There are many reasons for this move, but the main one is simply this: one need go beyond the bounds of a pure theory of parthood to come out with a true theory of parts and *wholes*. For as we have already mentioned, mereology by itself cannot do justice to the notion of a whole (a one-piece, self-connected whole such as a stone or a whistle, as opposed to a scattered entity made up of several disconnected parts, such as a bikini or a broken glass). Parthood is a relational concept, wholeness a global property, and the latter just runs afoul of the former.

Whitehead's early attempts to characterize his ontology of events, as presented at length in his *Enquiry* (1919) and in *The Concept of Nature* (1920), exemplify this difficulty most clearly. The mereological system underlying Whitehead's ontology was not meant to admit of arbitrary wholes, but only of wholes made up of parts that are “joined” or connected to one another—specifically, finitary sums of such parts. Thus, Whitehead was working with a composition principle patterned after (P.11_ψ), in fact with the corresponding biconditional, with ‘ψ’ understood as a predicate expressing the relevant relation of connection. And Whitehead's characterization of this relation was purely mereological:⁷

$$(53) \quad \psi xy =_{df} \exists z(Ozx \wedge Ozy \wedge \forall w(Pwz \rightarrow Owx \vee Owy))$$

Looking at the spatial patterns in Figure 1.8, we can see how this definition is intended to work. What distinguishes the connected sum $x + y$ on the left from the disconnected sum in the middle? Well, in the former case it is easy to find regions, such as z , that overlap both x and y without outgrowing the sum—regions that lie entirely *within* $x + y$. By contrast, in the middle pattern it would seem that every z overlapping both x and y will also overlap their complement—the entity that *surrounds* $x + y$. Thus, only the left pattern satisfies the condition expressed by ‘ψ’; the pattern in the middle violates it. However, Figure 1.8 also shows why this account is defective. For nothing guarantees that the

⁷The definition below corresponds to the formulation given in Whitehead (1920: 76). Whitehead's earlier definition (1919: 102) is slightly different but essentially equivalent.




Figure 1.8. A connected sum (left) and two disconnected sums (middle, right)

item z overlapping two “joined” items x and y be itself in one piece, so the pattern on the right depicts two entities that satisfy the condition expressed by ‘ ψ ’, too. Yet this is a case where we should like to say that x and y are *not* connected. Of course, Whitehead would disqualify the counterexample because his ontology does not contain any disconnected z s—but this is plain circularity. The account works on the assumption that only self-connected entities can inhabit the domain of discourse, yet that is precisely the assumption that (53) is meant to characterize.

These considerations apply *mutatis mutandis* to other attempts to subsume topological connectedness within a bare mereological framework (see e.g. Bostock 1979, Needham 1981, Ridder 2002). Nor is this exclusively an ontological concern. These limits show up in any attempt to account for a number of important spatial concepts besides connectedness, such as the distinction between a completely interior part and a tangential part that is connected with the exterior, of the difference between an open entity and a closed one. All of these—and many others indeed—are relations that any theory concerned with the spatial structure of the world should supply and which cannot, however, be defined directly in terms of plain mereological primitives.

2.1 ‘Contact’ and connection

It is here that topology comes into the picture. The connection relation that Whitehead was seeking to characterize is a topological relation. And if it cannot be defined in mereological terms, it must be formally treated on independent grounds.

Before looking at how this can be done, it is important again to begin with a couple of terminological caveats. As with ‘part’, the term ‘connection’, and cognate terms such as ‘contact’ and ‘touching’, have different meanings in ordinary language, only some of which correspond to the intended relation. Most notably, in ordinary language we do not draw a clear distinction between a truly topological notion of connection and a merely metric notion of contact. Consider:

- (54) The handle is attached to the mug.
- (55) The table is touching the wall.

The relation exemplified by (54) is topological: the handle and the rest of the mug form a unitary whole. For practical purposes there may be room for free rein, depending on whether the handle is glued to the rest of the mug (what Galton 2000, §4.2, calls “adhesion”) or truly continuous with it (“cohesion”), but either way there is an obvious sense in which we are dealing with a single, one-piece object. By contrast, the relation exemplified by (55) is not topological but metric: the table is so close to the wall that we are inclined to say they are connected. If space were discrete, this might be the right thing to say. But if space is dense, as we may plausibly assume, then the surfaces of two bodies can never truly be connected, short of overlapping: there will always be a narrow gap separating them. The narrower this gap, the easier it is to disregard it for practical purposes, but genuine topological connection can only obtain when the gap is reduced to zero. We shall see in Section 2.5.1 that some interesting topological relations may be introduced to capture at least some uses of the metric relation of contact. Overall, however, metric relations cannot be squeezed into the conceptual apparatus of topology without the help of strongly simplifying assumptions on the structure of space. (In this sense, a diagram such as Figure 1.8, left, is ambiguous, since one might think of x and y as being merely close to each other, as when we draw a picture of a table against a wall.)

A related issue concerns the distinction between connection patterns that involve a single point of contact, as in (56), or an extended boundary portion, as in (57), if not of an entire boundary, as in (58):

(56) Colorado is connected to Arizona.

(57) France is connected to Germany.

(58) The Vatican is connected to Italy.

To some extent this is a matter of convention. We may disregard (56) as irrelevant, or we may treat it as an acceptable case.⁸ Since most theories go for the second option, we shall follow them on this score. As it turns out, under suitable conditions one can on such basis draw all the relevant distinctions (Section 2.5), so the choice proves to be a convenient one.

Finally, let us just mention the fact that topological connection is, in a way, an idealized relation. Physically, as we know, the world consists of objects that are not continuous (or dense) in the relevant sense, and speaking of their boundaries is like speaking of the “flat top” of a fakir’s bed of nails (Simons 1991a: 91). Physically, a mug is just a

⁸This may also depend on context. In chess, for example, the choice depends on whether we are a rook or a bishop.

swarm of subatomic particles whose exact shape and extension involves the same degree of arbitrariness as a mathematical graph smoothed out of scattered data. In this sense, the intuition behind the claim that (54) provides a good example of topological connection betrays a naive conception of the mid-size world. However, this is not to say that topology is inadequate as a tool for practical spatial reasoning. For, on the one hand, we are generally interested in describing the spatial structure of the world precisely insofar as objects are conceptualized as finite chunks of dense matter with closed, continuous boundaries. Even if talk of boundaries and contact were deemed unsuited to the ontology of the physical sciences, one would therefore need it when it comes to the dense entities carved out by ordinary discourse and to the spatial regions that these occupy. On the other hand, the geographic examples in (56)–(58) illustrate that at least *some* entities countenanced by common sense measure up to the strict standards of topological connection. We do want to say that a geopolitical unit occupies a region of space that is strictly dense in the relevant sense, even though the underlying territory may consist of material stuff that on closer inspection is best described as a gerrymandered aggregate of zillions of disconnected subatomic particles.

2.2 Basic principles and definitions

We are now in a position to take a closer look at the idea of a topological extension of mereology—an extension that would take us beyond the prospects of a pure theory of parthood to deliver a genuine theory of parts and *wholes*. To this end, let us expand our formal language by adding a second distinguished predicate constant, ‘C’, to be understood intuitively as the relation of topological connection. The question of how mereology can actually be expanded to a richer *part-whole* theory may then be addressed by investigating how a P-based mereological system of the sort outlined in Section 1 can be made to interact with a C-based topological system.

Again, we may distinguish for this purpose “lexical” from substantive postulates for ‘C’, regarding the former as embodying a set of minimal prerequisites that any system purporting to explicate the meaning of the concept of ‘connection’ must satisfy. And a natural starting point is to assume that such lexical principles include at least the twofold requirement that ‘C’ be reflexive and symmetric:

- | | | |
|-------|-----------------------|-------------|
| (C.1) | Cxx | Reflexivity |
| (C.2) | $Cxy \rightarrow Cyx$ | Symmetry |

There is little room for controversy concerning the intuitive adequacy of (C.1)–(C.2), provided that we take ‘C’ to express, not just the relation of

external connection that may obtain between two disjoint entities that share a common boundary (in some intuitive sense to be made precise), but the relation of connection that may obtain between any two entities that share *at least* a boundary. In this sense, mereological overlap qualifies as connection, too. We shall come back to this idea shortly. First, let us note that ‘C’ need not, on the intended interpretation, express a transitive relation. France is connected to Germany and Germany to Poland, but France and Poland are not connected—they do not share any common boundary. We can, however, consider a notion of connection that captures the fact that France is connected to Poland *by* Germany. De Laguna (1922), a forerunner in the area of qualitative topological reasoning, actually based his account on a three-place primitive corresponding to this relation.⁹ In terms of ‘C’, De Laguna’s primitive is easily defined:

$$(59) \quad \text{BC}xyz =_{df} \text{C}xz \wedge \text{C}zy \quad \text{By-Connection}$$

and we can accordingly introduce the desired notion of (possibly) indirect or mediate connection as follows:

$$(60) \quad \text{MC}xy =_{df} \exists z \text{BC}xyz \quad \text{Mediate Connection}$$

By an obvious generalization, we can also define:

$$(61) \quad \text{MC}^n xy =_{df} \exists z_1 \dots z_n (\text{C}xz_1 \wedge \dots \wedge \text{C}z_n y) \quad n\text{-Connection}$$

It follows immediately from (C.1) and (C.2) that each MC^n is reflexive and symmetric, and the union of all such relations is transitive. In the absence of further principles, however, e.g., principles guaranteeing the existence of an entity connected to all the intermediate links, such a transitive union cannot be defined in the object language—unless we allow for quantification over positive integers:

$$(62) \quad \text{TC}xy =_{df} \exists n \text{MC}^n xy \quad \text{Transitive Connection}$$

Now, let T be the first-order theory defined by the two basic axioms (C.1) and (C.2), in analogy with the theory M defined by the basic mereological axioms (P.1)–(P.3). T is, of course, extremely weak and a lot will have to be added before we can say we have an interesting topology. In particular, a model of T can be obtained simply by interpreting ‘C’ as mereological overlap, and what further principles should be added to T so as to distinguish C from O is precisely one of the questions a good

⁹Strictly speaking, De Laguna’s primitive is interpreted as “ x can be connected to y by z ”, so it involves a modal ingredient. For a formal treatment, see Giritli (2003).

topological theory is meant to answer. For instance, should one assume that connection is extensional, i.e., that things that are connected exactly to the same entities are identical? Should one assume that any two connected entities satisfy at least some form of Whitehead’s account in (53)? Or consider the binary relation defined by

$$(63) \quad Exy =_{df} \forall z(Czx \rightarrow Czy) \qquad \text{Enclosure}$$

It follows from (C.1)–(C.2) that this relation is reflexive and transitive, and if C is extensional, then E is also antisymmetric—a partial ordering. Should one assume this relation to satisfy any analogues of the axioms for parthood? For each mereological predicate defined in Section 1 using ‘ P ’ one could now introduce a corresponding topological predicate using ‘ E ’ instead. Should one assume any corresponding axioms?

As it turns out, it is difficult to answer these questions in an abstract setting (see Cohn and Varzi 2003). Obviously, much depends on how exactly ‘ C ’ is interpreted, and that in turn may depend on how one thinks ‘ C ’ and ‘ P ’ should interact. Rather than pursuing these questions in isolation, then, let us proceed immediately to examining the main options for combining mereology and topology.

2.3 Bridging principles

The simplest option is just to append the T -axioms to our preferred mereological theory, X , to obtain a corresponding “mereotopology” $X + T$, which can then be strengthened by supplying further axioms for ‘ C ’. However, this would be of little interest unless one also adds some mixed principles to establish an explicit “bridge” between X and T .

2.3.1 Parts and wholes. There is one sort of bridging principle that most theories, if not all, accept: it centers around the intuition that no matter how P and C are fully characterized, they must be related in such a way that a whole and its parts are tightly connected. Here are three ways one can try to capture this intuition:

$$\begin{array}{ll} (C.3_a) & Pxy \rightarrow Cxy \qquad \text{Integrity} \\ (C.3_b) & Oxy \rightarrow Cxy \qquad \text{Unity} \\ (C.3) & Pxy \rightarrow Exy \qquad \text{Monotonicity} \end{array}$$

The first principle, (C.3_a), is perhaps the most immediate: just as everything is connected to itself by (C.1), everything must be connected to its constitutive parts. This is not to say that the parts must all be connected to one another: the two main parts of a bikini are not. But they are, in an obvious sense, connected to their sum; they are detached from each other but not from the whole bikini.

As it stands, however, (C.3_a) is extremely weak. It doesn't even capture the idea that if something is part of *two* things, then those things are thereby connected. This is not to say that they are connected *by* that common part, in the sense defined in (59); they are connected *because* of that common part. In other words, if sharing a common boundary is to count as sufficient for connection, then *a fortiori* sharing a common part ought to be sufficient, too. It is in this sense that overlap is to be regarded a special (and somewhat trivial) case of connection. And the second principle, (C.3_b), makes this explicit.

(C.3_b) is stronger than (C.3_a), since parthood implies overlap. Moreover, since the converse need not hold (on pain of trivializing the notion of connection), (C.3_b) provides the intuitive grounds for defining a non-trivial notion of external connection, or touching, that can only hold between mereologically disjoint entities:


$$(64) \quad \text{EC}xy =_{df} Cxy \wedge Dxy \qquad \text{External Connection}$$

(Note that this relation is still symmetric, but not reflexive; it is actually irreflexive, due to the irreflexivity of D.) This is an important notion, which makes all the difference between mereology and mereotopology. Yet (C.3_b) is still too weak to capture the fundamental intuition that we are after. For while this principle guarantees that overlapping a part is sufficient for being connected to the whole, it doesn't secure that *touching* a part is also sufficient. Surely something can touch a mug (say) just by touching its handle. So it is only with the third principle, (C.3), that we get a plausible formulation of the basic idea. Connection, if it is to behave properly, must be monotonic with respect to parthood.

It is easily checked that (C.3) implies (C.3_b), hence (C.3_a), so let us just focus on (C.3), and let us call *MT* (for *Minimal (mereo)Topology*) the corresponding extension of *T*.¹⁰ Whether (C.3) is to be classified as a "lexical" principle may be controversial and will depend, in an obvious sense, on the underlying axioms for 'P'. Nonetheless, the principle itself is part of virtually every mereotopological theory in the literature, either as an axiom (Varzi 1996a, Donnelly 2004) or as a theorem. And although *MT* is still far from providing an adequate characterization of the relation of topological connection, it provides the basis for the definition of a number of important spatial relations which, like EC, cannot be distinguished within a purely mereological setting. In particular, we can now express the difference between a proper part that lies entirely

¹⁰In fact, the result of adding (C.3) to *T* yields a slightly redundant theory; a more elegant formulation can be obtained by dropping (C.2) and replacing (C.3) with the following variant:

$$(C.3') \quad Pxy \wedge Cxz \rightarrow Czy$$



	Ux,y	ECx,y	Ox,y	$TPPx,y$	$IPPx,y$	EQx,y	$TPEx,y$	$EPEx,y$
$n = 1$	+	-	-	-	-	-	-	-
$n = 2$	+	+	-	-	-	-	-	-
$n = 3$	+	-	+	-	-	-	-	-
$n = 4$	+	-	+	+	-	-	-	-
$n = 5$	+	-	+	-	+	-	-	-
$n = 6$	+	-	+	-	-	+	-	-
$n = 7$	+	-	+	-	-	-	+	-
$n = 8$	+	-	+	-	-	-	-	+

Figure 1.9. Basic mereotopological relations. (Shaded cells indicate connection; darker shading stands for parthood.)

within the interior of the whole and a proper part that is connected with the exterior:

$$\begin{aligned}
 (65) \quad & IPPxy =_{df} PPxy \wedge \forall z(Czx \rightarrow Ozy) && \text{Interior PP} \\
 (66) \quad & TPPxy =_{df} PPxy \wedge \neg IPPxy && \text{Tangential PP} \\
 (67) \quad & EPExy =_{df} PExy \wedge \forall z(Czy \rightarrow Ozx) && \text{Exterior PE} \\
 (68) \quad & TPExy =_{df} PExy \wedge \neg IPExy && \text{Tangential PE}
 \end{aligned}$$

Note that, given (64), interior parts satisfy the following:

$$(69) \quad IPPxy \leftrightarrow (PPxy \wedge \neg \exists z(ECzx \wedge ECzy))$$

Thus, tangential parts are those parts that reach far enough to touch something with which the whole itself is just in touch. Similarly for proper extensions. The diagram in Figure 1.9 indicates how these predicates may represent a genuine addition to the mereological vocabulary introduced in (19)–(23) and illustrated in Figure 1.1. We shall see in Section 2.4 that this diagram may actually be misleading, owing to the delicate role played by boundaries in the proper understanding of the connection relation; but for the moment we take the intuitive, geometric interpretation of the diagram to be adequate enough to serve its purpose.

2.3.2 Parthood vs. enclosure. Things begin to be controversial as soon as we consider the possibility of stronger bridging principles. Consider again the three principles above. Clearly the converse of Integrity, (C.3_a), is unacceptable. And unacceptable is also the converse of Unity, (C.3_b), for then connection would collapse on the relation of mereological overlap and the definitions in (64)–(68) would lose their intuitive appeal. On the other hand, the converse of the Monotonicity principle (C.3) is not obviously unreasonable:

$$(C.4) \quad Exy \rightarrow Pxy \quad \text{Converse Monotonicity}$$

This says that a sufficient condition for one thing to be part of another is that whatever is connected to the former is also connected to latter. This sounds intuitive, and several authors would actually include (C.4) as a further bridging principle on top of *MT*. Indeed, a principle along these lines may already be found in Whitehead’s latest work, *Process and Reality* (1929)—a remarkable fact, since the conjunction of (C.3) and (C.4) yields a biconditional that would allow one to define parthood in terms of connection:

$$(70) \quad Pxy \leftrightarrow Exy$$

Thus, if (53) turned out to be a defective attempt to reduce topological concepts to purely mereological ones, (C.4) (together with (C.3)) reflects a reductionist attempt in the opposite direction, to the effect that mereological concepts can be defined in terms of purely topological ones. And although few have followed Whitehead on the first route, it is a fact that many authors have taken the second strategy into serious consideration. Clarke (1981) provides the most influential example, and the so-called “Region Connection Calculus” originated with Randell *et al.* (1992) is the best case in point when it comes to theories designed specifically for applications to spatial reasoning (see also Gotts *et al.* 1996 and Cohn *et al.* 1997). So the question deserves close scrutiny: Is (C.4) a reasonable addition to the basic postulates of *MT*?

Never mind the fact that working with just one primitive may be mathematically attractive. As it turns out, it is equally possible to rely on a single primitive even in the absence of (C.4). For instance, one can rely on the ternary relation ‘*x* and *y* are connected parts of *z*’ (Varzi 1994). Writing this as ‘*CPxyz*’, one could define ‘*P*’ and ‘*C*’ as follows:

$$(71) \quad Pxy =_{df} \exists z Pxyz$$

$$(72) \quad Cxy =_{df} \exists z Pxyz$$

and then go on to develop a theory based on the relative irreducibility of these two predicates. (Note that (71) and (72) only carry minimal presuppositions: that parthood be reflexive and that connected entities have an upper bound.) So the issue is not formal economy—the use of a single primitive. It is, rather, conceptual economy. In the extension of *MT* obtained by adding (C.4)—henceforth *RMT*, for *Reductive Mereotopology*—the *notion* of parthood is fully subsumed under that of connection, and the limits of mereology are overcome by turning the original problem upside down: parthood cannot deliver the full story, but connection can. However, there are at least two worries here.

The first worry concerns the material adequacy of the reduction. As Masolo and Vieu (1999) have observed (but the point goes back to Ran-



Figure 1.10. Two counterexamples to converse monotonicity

dell *et al.* 1992: §5.1), (C.4) appears to have implausible consequences if the domain contains entities with atomic proper parts. Consider an extended region (or object) x , and let y be x minus an atomic part z (Figure 1.10, left). On any reasonable understanding of ‘C’, everything connected to x is connected to y , since z is connected to both. So (C.4) would force x to be part of y . Yet, intuitively, x should count as an extension of y : it is bigger, it contains z , it contains y as a further *proper* part. Things get worse if we consider that (C.4) forces z itself to count as part of y , since z is connected to y and anything else is connected to z only if it overlaps y . Yet y was defined as x *minus* z . Of course, such models would be ruled out if *RMT* were strengthened by adding an atomlessness postulate such as (P.7). But this is precisely the point: (C.4) does not merely reinforce the bridge between P and C; it actually embodies more substantive views about the mereological structure of space. Besides, even the atomless variant *ARM*T would be open to counterexamples. For we have the same sort of problem if we suppose that z is a non-atomic proper part of x with no interior proper parts of its own (Figure 1.10, right). To rule out this model, a stronger assumption than (P.7) would be needed, corresponding to the thesis that everything has *interior* proper parts:

$$(C.5) \quad \exists x \text{IPP}xy \qquad \text{Boundarylessness}$$

In fact, it is precisely with the help of an axiom like (C.5) that we can give expression to a Whiteheadian, boundary-free conception of space (see below, Section 2.4.3). However, this is a controversial conception. Why should our analysis of parthood force upon us a rebuttal of the boundary concept? Why should we assume (C.5) in order to ensure a coherent implementation of a basic bridging principle like (C.4)?

The second worry is more general. Consider an object and the stuff it is made of—for instance, a statue and the corresponding amount of clay. As we have seen, few would regard the relationship of material constitution that holds between such entities as a case of proper parthood. And there are many philosophers for whom constitution is not identity

(improper parthood) either: see again the papers in Rea 1997. This is by no means a pacific thesis, but never mind. The point is simply that the relationship between the clay and the statue is not obviously an instance of parthood. Yet, on any plausible understanding of ‘C’, whatever is connected to the clay is bound to be connected to the statue, too, so (C.4) would immediately settle the issue: the clay *is* part of the statue. Indeed, since it is equally plausible to suppose that the same applies in the opposite direction—whatever is connected to the statue is connected to the clay—(C.4) implies that the statue and the clay are one and the same thing. And this is a substantive tenet which, as such, ought not to be built into the basic apparatus of mereotopology *at the outset*.

This worry is perhaps best appreciated by noting that the following theorem is an immediate consequence of (C.1)–(C.4), provided P satisfies the basic mereological axioms (P.1)–(P.3):

$$(73) \quad \forall z(Czx \leftrightarrow Czy) \rightarrow x = y$$

In fact, (P.1) and (P.2) (reflexivity and transitivity) are derivable from (70), and *RMT* turns out to be equivalent to the theory defined by taking (73) as an axiom along with (C.1) and (C.2). (That is actually the customary axiomatization since Clarke 1981; see also Biacino and Gerla 1991.) Now, with parthood construed as enclosure, (73) is nothing but the antisymmetry principle (P.3). Yet (73) does not merely assert the antisymmetry of parthood; it says that connection is extensional—that different things cannot connect to the same things. And this is just as controversial as the thesis that parthood is extensional.

It could be replied that the analogy with mereological extensionality is in fact helpful, since the original arguments in support of (P.5) (Section 1.3.2) could now be offered on behalf of (C.4). Indeed, if a statue and the clay are construed as things that, at some level of decomposition, share the same constituents—e.g., the same molecules—the question of whether constitution is identity just *is* the question of whether parthood is extensional. However, the worry does not only apply to cases of material constitution. Consider a shadow cast onto a wall. The shadow is not part of the wall, yet anything connected to the shadow is—arguably—connected to the wall. Or consider a stone inside a hole. The stone is not part of the hole, yet one could argue that anything connected to the stone is connected to the hole. Broadly speaking, the problem arises as soon as we allow for the possibility that distinct entities occupy the same space (Casati and Varzi 1999). That this is a real possibility is by itself contentious and is one of the questions to be addressed by an explicit theory of location (Section 3). But precisely for this reason, ruling it out on mereotopological grounds would be utterly inappropriate.

This last point is particularly worth stressing, for it shows once again how the choice of a suitable set of principles may depend crucially on whether we are interested in a theory aimed at modeling the domain of all spatial entities or just a domain of pure spatial regions. The worry mentioned above arises forcefully in the context of theories of the first sort. It does, however, lose its force in relation to the second sort of theory, since two regions cannot overlap spatially without overlapping mereologically. Now, it is a fact that most authors committed to (C.4) have been working on such a narrower understanding of their theory. Whitehead’s own account was explicit in this regard: in the theory put forward in *Process and Reality* (as in Clarke’s 1981 reformulation), the field of C was meant to consist exclusively of spatial regions. On the other hand, it is also a fact that a major motivation for developing a theory of this sort has been the assumption that connection thus understood is all that matters for practical purposes. For one can always treat the relation of connection as the “shadow” (in De Laguna’s 1922: 450 apt terminology) of the relation of physical contact or overlap that may obtain between actual, concrete entities. In other words, such theories have typically been developed on the assumption that the following principle provides the necessary and sufficient link between the mereotopology of pure space and the mereotopology of spatial entities broadly understood:

$$(74) \quad x \text{ is connected to } y \text{ if and only if the region occupied by } x \\ \text{is connected to the region occupied by } y.$$

If so, however, then the problems mentioned above resurface even for theories of this sort. For (74) will deliver an acceptable account if, and only if, spatial co-location is regarded as metaphysically impossible.

2.3.3 Self-connectedness. There are other ways of supplementing *MT* with bridging principles that go beyond (C.3). In particular, consider again Whitehead’s early attempts to characterize topological connection in terms of parthood, i.e. (53), with ‘ ψ ’ understood as ‘ C ’. As a *definition*, this was found defective. However, one may certainly consider adding the corresponding biconditional as an axiom—or at least adding one of the two conditionals:

$$(C.6) \quad Cxy \rightarrow \exists z(Ozx \wedge Ozy \wedge \forall w(Pwz \rightarrow Ow x \vee Ow y)) \quad \text{Left Join} \\ (C.7) \quad \exists z(Ozx \wedge Ozy \wedge \forall w(Pwz \rightarrow Ow x \vee Ow y)) \rightarrow Cxy \quad \text{Right Join}$$

Would this be a good way of tightening the conceptual link between the mereological and the topological ingredients of *MT*?

Consider (C.6). As it turns out, its status depends significantly on the underlying axioms for ‘ P ’. If our mereological theory is strong enough

to warrant the existence of a fusion for any pair of connected entities, i.e., if it contains the relevant instance of (P.11 _{ψ}) as an axiom or (as in *GEM*) as a theorem, then (C.6) itself is derivable as a theorem, since the fusion of x and y is sure to qualify as a z satisfying the consequent. If, however, our theory does not warrant all the relevant fusions, then (C.6) may still be regarded as a plausible addition to *MT*. Perhaps the fusion of any two connected entities turns out to be too large or gerrymandered to be acceptable, but we may still think of connection as sufficient for the existence of smaller fusions encompassing those portions of x and y that are sufficiently close to their common boundary. For example, with reference to Figure 1.11, left, suppose we are only willing to acknowledge the existence of entities that are composed of at most two disconnected parts. Then the fusion of x and y is out, but a fusion of x_2 and y_2 , as well as fusions of x and y_2 and of x_2 and y , would fit the bill. All of this speaks in favor of (C.6), though it may be argued that the existential import of this principle goes beyond the task of establishing a necessary conceptual link between P and C.

As for (C.7), the picture is different. Assuming this principle is virtually tantamount to excluding disconnected entities from the domain—not all of them, to be sure, but many of them. For example, if the underlying mereology is sufficiently weak, (C.7) is compatible with the existence of a disconnected composite such as c in Figure 1.11, right: on the assumption that c has no further proper parts besides a_1 , a_2 , and a_3 (and parts thereof), the antecedent is false so (C.7) is vacuously satisfied. But consider a bikini, or a printed inscription consisting of separate letter tokens. As we have already noted, one need not buy into unrestricted composition to appreciate the dignity of such things. Yet their existence would be banned by (C.7). If x and y are the two main parts of a bikini, then the consequent of (C.7) is false even though the antecedent is made true by the bikini as a whole. So, on the face of it, this direction of Whitehead's biconditional is definitely too strong as a general bridging principle and there is no philosophically neutral reason to add it to *MT*. One can, however, consider weaker versions, to the effect that the consequent of the conditional must hold whenever the antecedent is made true by the right sort of entity:

$$(C.7_\phi) \quad \exists z(\phi z \wedge Ozx \wedge Ozy \wedge \forall w(Pwz \rightarrow (Owx \vee Owy))) \rightarrow Cxy$$

In particular, one can take ' ϕ ' to express the property of being self-connected. After all, this was precisely the intended import of Whitehead's flawed definition. And we have seen that the flaw of the definition, in the *if* direction, was not conceptual but formal: it lied exclusively in the impossibility of expressing the relevant restriction in mereological




Figure 1.11. Left joining with partial fusions; right joining without connection.

terms. By making the restriction explicit, (C.7_φ) overcomes the difficulty and suggests itself as a natural bridging principle.

Surprisingly, it is not easy to express the property of self-connectedness even in the extended language of mereotopology. If the axioms on \mathbf{P} are strong enough, we can follow the ordinary set-theoretic definition—something is self-connected if it doesn't consist of disconnected parts:

$$(75) \quad \text{SC}x =_{df} \forall yz(\forall w(\text{O}wx \leftrightarrow (\text{O}wy \vee \text{O}wz)) \rightarrow \text{C}yz) \quad \text{Self-Conn.}$$

In particular, in $MT + GEM$ this becomes:

$$(76) \quad \text{SC}x \leftrightarrow \forall yz(x = y + z \rightarrow \text{C}yz)$$

This is a common definition in the literature, both among theorists subscribing to the converse monotonicity principle (C.4) (see e.g. Clarke 1981, Randell *et al.* 1992) and among theorists rejecting it (Tiles 1981, Varzi 1994, Smith 1996). However, if the axioms on \mathbf{P} do not secure the necessary composition patterns, the definition is inadequate. For example, the object c in Figure 1.11, right, is anything but self-connected, yet it (vacuously) satisfies the definiens of (75) unless we assume the existence of at least one sum consisting of a_1 and a_2 , a_2 and a_3 , or a_1 and a_3 . In the finitary case, the difficulty could be met by relying on the notion of mediate connection: any two parts of a self-connected entity must be at least n -connected for some n . More generally:

$$(77) \quad \text{PC}x =_{df} \forall yz(\text{P}yx \wedge \text{P}zx \rightarrow \text{TC}yz) \quad \text{Path Connectedness}$$

This, however, involves quantification over numbers (see (62)), which just confirms the expressivity limits in question. Moreover, (77) does not work in the infinitary case: the unit interval on the real line is connected, but we cannot account for this fact in terms of the relationships between the reals themselves; reference to subintervals is necessary, specifically reference to a subinterval and its relative complement. So, overall it appears that the notion of self-connectedness can be adequately grasped, via (75), only by theories that are at least as strong as $MT + (\text{P.6})$, the complementation principle, though this is at present an open question.




Figure 1.12. An implausible model of $MT + GEM$.

2.3.4 Fusions. We conclude this discussion of bridging principles by noting that even a theory as strong as $MT + GEM$ + any of the above axioms is incapable of capturing all the relevant links between mereology and topology. In particular, such a theory is consistent with the following implausible thesis (Tsai 2005: 137):

$$(78) \quad \exists z(Cz(x + y) \wedge \neg Cz x \wedge \neg Cz y)$$

A model is given in Figure 1.12, where the curve line indicates the relevant connection relationship (besides the obvious ones imposed by (C.3)).

Clearly, this is a sign that some additional bridging principle is on demand. The following option suggests itself:

$$(C.8) \quad z = \Sigma x \phi x \rightarrow \forall y(Cyz \rightarrow \exists x(\phi x \wedge Cyx)) \quad \text{Fusion Connection}$$

Whether this is enough to establish a good correlation between the mereological structure of composite objects and their topological behavior is a question that can hardly be addressed in general terms. The plausibility of (C.8), however, seems obvious. Since we have found good reasons to also accept (C.4) and (C.7_{5C}), the theory resulting by adding these three principles to $MT + GEM$ suggests itself as the natural topological extension of GEM . (Recall that (C.6) is already provable in $MT + GEM$.) For future reference, we shall call this theory $GEMT$, for *General Extensional Mereotopology*.

2.4 Extensions and restrictions

As it turns out, $GEMT$ does not officially appear in the literature, mostly due to the limited study of such principles as (C.7_{5C}) and (C.8). (The closest relatives are the axiomatic systems advocated by Smith 1996 and Casati and Varzi 1999.¹¹) A good thing about this theory, as

¹¹In Casati and Varzi (1999), $GEMT$ is identified with $MT + GEM$. Smith's (1996) version is based on a primitive 'P' for (possibly improper) interior parthood. See also Pianesi and Varzi (1996a, 1996b) for similar formulations based on the predicate 'B' and the operator 'c' defined below, respectively.

also about the core fragment afforded by $MT + GEM$, is that it makes it possible to supplement the mereotopological predicates and operators discussed so far with a number of additional operators that mimic the standard operators of point-set topology. For example:

$$\begin{array}{ll}
 (79) & ix =_{df} \Sigma z \forall y (Czy \rightarrow Oxy) & \text{interior} \\
 (80) & ex =_{df} i(\sim x) & \text{exterior} \\
 (81) & cx =_{df} \sim (ex) & \text{closure} \\
 (82) & bx =_{df} \sim (ix + ex) & \text{boundary}
 \end{array}$$

Like the mereological operators in (40)–(44), these operators are partially defined in view of the lack of a null entity that is part of everything. For instance, if x is a boundary, it has no interior, and if x is the universal entity \mathcal{U} , it has no exterior. Even so, in $GEMT$ all of these operators are rather well-behaved. In particular, we can get closer to standard topological theories by explicitly adding the mereologized analogues of the standard Kuratowski (1922) axioms for topological closure:

$$\begin{array}{ll}
 (C.9) & Px(cx) & \text{Inclusion} \\
 (C.10) & c(cx) = cx & \text{Idempotence} \\
 (C.11) & c(x + y) = cx + cy & \text{Additivity}
 \end{array}$$

(These axioms are to be read as holding whenever c is defined for its arguments. Here and below we omit the relevant existential conditions to improve readability.) Indeed, (C.9) and (C.11) turn out to be provable in $MT + GEM$; see Tsai (2005: 141).

The possibility of supporting such developments is of course a good indication of the strength of $GEMT$. Philosophically, however, this strength may be regarded with suspicion, and several complaints have been raised in the literature.

2.4.1 The open/closed distinction. The main sort of complaint concerns the very notion of connection that the theory is meant to characterize. So far we have worked mostly with an intuitive notion in mind but obviously more can and must be said—and $GEMT$ says a lot. In particular, the Kuratowski extension of $GEMT$ ($KGEMT$ for short) yields a full account of the intended interpretation of ‘C’: two things are (externally) connected if and only if they share (only) a boundary, i.e., if and only if the closure of one overlaps the other, or vice versa:

$$\begin{array}{ll}
 (83) & Cxy \leftrightarrow (Ox(cy) \vee O(cx)y) \\
 (84) & ECxy \leftrightarrow (Cxy \wedge \neg C(ix)(iy))
 \end{array}$$

Now, this shows in what sense the behavior of C in this theory closely approximates that of standard set-theoretic topological connection; just

let ‘ x ’ and ‘ y ’ range over sets of points and interpret ‘ O ’ as set intersection. On the other hand, one aspect in which ordinary point-set topology appears to conflict with common sense—an aspect that has been emphasized by authors interested in a mereotopological characterization of qualitative spatial reasoning, such as Randell *et al.* (1992) and Gotts *et al.* (1996)—is precisely the distinction between “open” and “closed” entities on which it rests, and which *GEMT* preserves *holus bolus*:

- (85) $OPx =_{df} x = ix$ Open
 (86) $CLx =_{df} x = cx$ Closed

This distinction goes at least as far back as Bolzano (1851: §66f). But already Brentano (1906: 146) regarded it as “monstrous”, and we have already seen that the sort of idealization it embodies does not sit well with the way we ordinarily speak. We may intuitively grasp the difference between an open and a closed interval on the real line—the objection goes—and we may even understand how this difference applies to ideal three-dimensional manifolds such as Euclidean space. But what does it mean to draw a similar distinction in the realm of concrete spatial entities, where the very notion of a boundary is the result of a conceptual idealization? What does it mean to say that some objects are closed and some are not, and that contact is only possible between objects of one type and objects of the other?

Besides, even if common sense and ordinary language were put aside, the open/closed distinction seems to yield genuine paradoxes as soon as we move from the realm of pure space to its worldly population: Consider (i) what happens when a solid body splits into two halves. Before the splitting the two halves were in contact, so we are to suppose that one was closed and the other open, at least in the relevant contact area. But then, after the splitting, only one of the two halves will have a complete boundary. This may not be “monstrous”, but it certainly seems implausible: the two halves—one should think—are perfectly indistinguishable. On the other hand, consider (ii) what happens when two bodies come into contact. We may imagine the same experiment performed twice. First we take an open cube and push it toward a closed cube until they touch. Then we do the same with two closed cubes. What reason can we offer to explain the fact that in the latter case the two cubes will *not* come into contact? As Zimmerman (1996a: 12) put it, what sort of “repulsive forces” can be posited to explain such deferential behavior?

There is no quarrel that these are pressing questions. (For more examples, see Kline and Matheson 1987.) Nonetheless, there are various things one can say in reply to such worries. Concerning (i), for example, one could say that the paradox is grounded on a questionable model of




Figure 1.13. Wrong (left) and right (right) topological models of splitting.

what happens when a process of “splitting” takes place (Varzi 1997). Topologically, this is no bloodstained business. Dissecting a solid body does not “bring to light” (Adams 1984: 400) a surface that was trapped inside and that must by necessity belong to one of two severed halves. Rather, the topological model is one of gradual deformation. Think of a splitting oil drop. The drop grows longer and, as it grows, the middle part shrinks and gets thinner and thinner. Eventually the right and left portions come apart and we have two drops, each with its own complete boundary. Ditto with any splitting object. A long, continuous process suddenly results in an abrupt topological change: there was one surface, one closed body, and now there are two (Figure 1.13). Of course, one can still raise a question about the last point of separation: Where does this one point belong—to the left half or to the right half? However, this is just a sign of the magic that surrounds any sort of topological change, as when you drill a hole through an object. The instantaneous event of a sphere turning into a torus is just as magical, and this magic has nothing to do with the open/closed opposition. It simply reflects the fact that topological change marks one point at which common sense reaches the limits of its theoretical competence, and a complete assessment would require a step beyond pure part-whole theorizing. It would require a step into the territory of qualitative kinematics, for example (as in Davis 1993), if not an account in terms of the microscopic analysis of matter. (What is—physically—the “last point of separation” involved in the splitting?)

This line of reply on behalf of *KGEMT* applies to (ii) as well—the “merging” puzzle (Casati and Varzi 1999). Surely the positing of repulsive forces to explain the peculiar behavior of the two closed cubes would be utterly and unagreeably *ad hoc*. But there are other possibilities. For instance, perhaps the two closed cubes *will* indeed come into contact. From the fact that two closed entities cannot *be* in contact it does not follow that they cannot *come into* contact, just as from the fact that two parts are connected it does not follow that they cannot be separated. Only, the coming into contact (just as the separation) determines a true topological catastrophe: there is a breaking through the relevant boundary parts and the two objects become one. (Think also of the two drops of oil merging into each other.) The two processes are dual: merging is the reverse of splitting. And both involve a seemingly magic moment that runs afoul of the confines of extensional mereotopology and calls for a thorough kinematic account.

One could still press the objection here by noting that the puzzles admit of perfectly *static* variants, where the appeal to kinematics would be out of place. Consider the dilemma raised by Leonardo in his *Notebooks*: What is it that divides the atmosphere from the water? Is it air or is it water?. Or consider Peirce’s puzzle: What color is the line of demarcation between a black spot and its white background? (1893: 98). More generally, given any object, x , does the boundary belong to x or to its complement? Does it inherit the properties—e.g., color properties—of x or of $\sim x$? There is no kinematic story to tell here. But how can one answer without selecting one candidate at random?

Here one might reply that figure/ground considerations will help. According to Jackendoff (1987, Appendix B), for example, normally a figure owns its boundary—the background is topologically open. This may well be the right thing to say *vis-à-vis* Peirce’s puzzle: the black spot is closed, so the line is black. But what is figure and what is ground when it comes Leonardo’s case? We do talk about the surface *of* the water, not of the air. But what goes on at the seashore? Three things meet—water, air, soil; how can figure/ground considerations help in such contexts? Perhaps such dilemmas are not real. Galton (2003: 167f), for example, argues that they arise as an artifact of the modeling process: surely properties like color or material constitution only apply to extended bodies, so it wouldn’t make sense to ask whether a boundary-like entity is air, water, or colored. There is, however, a less dismissive way to meet the challenge on behalf of *KGEMT*. For one may acknowledge that such dilemmas are real and yet insist on a friendly attitude towards the open/closed distinction. The actual ownership of a boundary—one might argue—is not an issue that a mereotopological theory must be

able to settle. The theory only needs to explain what it means for two things to be connected. *Which* things are open and which are closed is a metaphysical question that, plausibly enough, goes beyond the concerns of the theory. If the ocean is a closed body, then it can only touch the air if the latter is open. If it isn't, then it can only touch the air if the latter is closed. And if both water and air are open, they cannot truly touch, though they may touch a closed piece of land. That's all the theory says, and there is no reason to think that the theory is wrong just because it is difficult to classify actual things into open and closed.

One last problem is worth mentioning. Consider again the cutting of a solid object in half. We have said that this process does not bring to light a new surface. But, of course, we can *conceptualize* a new, potential surface right there where the cut would be. In fact, we can conceptualize as many boundaries as we like, even in the absence of any corresponding discontinuity or qualitative heterogeneity among the parts. Think of John's waist, the equator, the Mason-Dixon line between Maryland and Pennsylvania. As Smith (1995, 2001) has pointed out, such "fiat" boundaries are a major ingredient of our picture of the world. Even the surfaces of ordinary objects, as we have seen, may involve a certain degree of fiat owing to the microscopic scatterdness of the underlying stuff. Yet *no fact of the matter* can support the ownership of boundaries such as these by one side rather than the other, hence there is no point in deferring to a metaphysical theory of the extended entities at issue. Isn't this enough to give rise to the demarcation puzzle?

Once again the answer is in the negative. Fiat boundaries are not physical boundaries *in potentia*. They are not the boundaries that would envelop the interior parts to which they are associated in case those parts were actually cut off. To think so would take us back to the wrong topological model illustrated in Figure 1.13, top. On a better model, fiat boundaries are just *placeholders* for genuine physical boundaries and the demarcation puzzle need not, therefore, lead to ontological anxiety. We can say that a boundary of this sort stands for *two* boundaries, one for each side. Or we can say that in drawing it we leave the question of its belongingness (hence the open/closed distinction) indeterminate. We do so because it is a question of no practical relevance. But precisely for this reason the indeterminacy is innocuous: it is pragmatic, perhaps semantic, not ontological—just like the sort of indeterminacy that afflicts the vagueness of parthood (Section 1.5). (It is in this sense that the diagram in Figure 1.9 is partly indeterminate: in saying that x_2 is externally connected to y , for example, we did not specify which of these two regions owns the boundary in the relevant contact area. Ditto for all other cases of external connection, as in Figure 1.8, left.)

2.4.2 Connection by coincidence. All of this, of course, is subject to controversy. If the foregoing remarks are found compelling, then the strength of *KGEMT* is vindicated and such theorems as (85) and (86) deliver a full and correct understanding of ‘C’. If not, however, then *KGEMT* will be deemed inadequate and the intended interpretation of ‘C’ remains unsettled. Are there any other options? We may distinguish two main alternatives, depending on whether or not a rejection of the open/closed distinction is taken to be compatible with a realist attitude towards the ontological status of boundaries.

The realist option finds its best expression in those theories that attempt to provide a detailed reconstruction of the view Brentano put forward in reaction to Bolzano’s “monstrous doctrine”, as in Chisholm (1984, 1993) and Smith (1997). According to this view, boundaries are genuine denizens of the world of spatial entities, but their lack of proper interior parts makes them peculiar in two important respects (Brentano 1976, part I). First, they can never exist except as belonging to entities of higher dimension. There are, in other words, no isolated points, lines, or surfaces, for boundaries are, in Chisholm’s terms, *dependent* entities. Second, and more to the point, insofar as boundaries are not possessed of divisible bulk, they do not occupy any space and can therefore share the same location with other boundaries. They can *coincide*, and the topological relation of external connection is to be explained, not via the open/closed opposition, but in terms of genuine boundary coincidence. Thus, we can speak of *the* surface of an object. But this single surface is to be recognized as being made up of two parts, two perfectly coinciding boundaries bounding the object and its complement, respectively.

As is clear, a rigorous formulation of such theories is no straightforward business. For one thing, it is not immediately obvious how to formulate the dependence thesis, both because of the modal ramifications that a good theory of ontological dependence would require (see e.g. Correia 2005) and because the relevant notion of dimension is by itself hard to characterize mereotopologically (Chisholm 1984). Ignoring such complications, and assuming *GEM*, one can capture the gist of the thesis as follows (Smith 1996):

$$(C.12) \quad SCx \wedge Bxy \rightarrow \exists z(SCz \wedge BPxz \wedge \neg \exists wBzw) \quad \text{Dependence}$$

where

$$(87) \quad Bxy =_{df} Px(by) \quad \text{Boundary}$$

$$(88) \quad BPxy =_{df} Bxy \wedge Pxy \quad \text{Boundary Part}$$

In other words, every self-connected boundary is part of some self-connected entity which it bounds and which is not itself a boundary.

(The restriction to self-connected entities is to avoid that (C.12) be trivially satisfied by a scattered z containing x as an isolated proper part.) Without the full mereological support of *GEM*, however, things are significantly more complex, among other reasons because of the apparent elusiveness of the self-connectedness predicate ‘SC’ (Section 2.3.3.)

Secondly, and more to the point, a lot depends on how exactly one understands the relation of spatial coincidence invoked by such theories to explain the phenomenon of (external) connection. Chisholm and Smith treat it as an undefined primitive, suitably axiomatized so as to guarantee that coinciding entities have coinciding parts. (See also Smith and Varzi 2000 for a similar treatment of the relation of coincidence between *fiat* boundaries.) Alternatively, one can embed the theory of coincidence into a theory of spatial location broadly construed: to say that things coincide is to say that they literally share the same location. Clearly, the choice between these two options is not just a matter of taste. Treating coincidence as a primitive is in principle compatible with different metaphysical conceptions of the nature of space, whereas the second option is best understood within the framework of a substantialist (Newtonian) conception, i.e., a conception according to which space is an entity in its own right. In any event, it is apparent that both options yield theories that are not strictly mereotopological, since a third primitive—coincidence or location, respectively—needs to be brought into the picture to provide a full account of the connection relation. We shall not go into the details of the first option here, but we shall have more to say about the second option in Section 3. For the moment, let us just observe that construing coincidence explicitly in terms of spatial co-location amounts to a partial reduction of topology to mereology: connection between entities of a kind (space occupiers) is reduced to overlap between entities of a different kind (their spatial receptacles), as per the following principle:

- (89) x is connected to y if and only if the region occupied by x overlaps the region occupied by y .

This is by itself interesting, though we are obviously left with the task of providing an account of the topology of space as such. And if the account is to match the strength of a theory such as *KGEMT*, then the open/closed distinction will at least be partially preserved. It will be obliterated from the Brentanian realm of space occupiers, but space itself would be Bolzanian. (Compare the initial worry: we *can* grasp how the distinction applies to ideal manifolds such as the real line or Euclidean space; it is when it comes to the realm of ordinary objects that their classification into “open” and “closed” is problematic.)




Figure 1.14. Open, closed, and semi-closed discs with the same area.

2.4.3 Omitting boundaries. The alternative route is to avoid the puzzles raised by the open/closed distinction by dismissing boundary talk altogether. This is the anti-realist option.

Philosophically, this route is often motivated on its own grounds, for instance because of the dubious ontological status of boundaries *vis-à-vis* the microscopic analysis of the physical world (Stroll 1988), or because of their suspect nature *qua* lower-dimensional entities (Zimmerman 1996b). In the context of formal theories, however, the main motivation for doing away with boundaries is precisely the rejection of the open/closed distinction *vis-à-vis* common sense. To use an example from Gotts *et al.* (1996: 57), Figure 1.14 depicts a disc with and without its boundary, and with just part of its boundary. Of course, the depiction has to show the boundary as having some finite thickness, which strictly speaking it does not possess. But this is the very point that appears counter to common sense: all three discs, if superimposed, would cover exactly the same *area*; yet the second disc includes unextended parts that the others do not, while the third includes some that the first does not. Such discriminations—it is argued—are not warranted.

There are radical as well as moderate variants of this view. The radical variants are represented by those theories that follow Whitehead (1929) in doing away with *all* boundaries. This amounts to assuming the boundarylessness axiom (C.5) in its full strength: everything has interior proper parts. The moderate variants, by contrast, only assume some weaker version of the axiom in which the variable is suitably restricted so as to range over entities of a certain sort:

$$(C.5_\phi) \quad \phi x \rightarrow \exists zIPPyx \qquad \text{Restricted Boundarylessness}$$

In particular, relative to our present concerns it is natural to construe ‘ ϕ ’ as a distinguished property of all concrete spatial bodies. This would allow for the possibility that space as such include points and other boundary-like elements, which means that the open/closed distinction would be partially preserved. But as we have just seen, restricting the distinction in this way may be enough to bypass the intuitive puzzles that it raises, so this may well be a good compromise. For example, Cartwright (1975) holds that concrete spatial bodies are the material

content of (regular) open regions of space, connection relations between the former being explained in terms of overlap relations between the closures of the latter:

- (90) x is connected to y if and only if the closure of the region occupied by x overlaps the closure of the region occupied by y .

As a mereotopological theory, this is of course another hybrid—just as the theory behind (89)—for it requires an explicit treatment of locative relations. Still, there is no question that (90) allows for a systematic boundary-free account of the mereotopology of concrete spatial bodies. (The real challenge, rather, is to justify the claim that only *some* regions are receptacles, e.g., only open regular regions; see Hudson 2002).

Let us focus on the radical variants. We have seen that positing (C.5) is a necessary move for any reductive mereotopology based on the converse monotonicity axiom (C.4), and it is a fact that most theories that accept one axiom accept the other as well. But let us put that aside for a moment and let us just focus on (C.5). Where X is any theory including MT , let $\bar{B}X$ be the corresponding *boundaryless* extension obtained by adding this axiom. What sort of mereotopology do we get?

As it turns out, the number of options is significantly constrained, both mereologically and topologically. For example, surely X cannot be atomistic, since (C.5) implies the atomlessness axiom (P.7). So any model of $\bar{B}X$ is perforce infinitary. And surely the interaction between compositional and decompositional principles will have to be carefully re-examined. In particular, it is easy to verify that in $\bar{B}MT$ the weak supplementation principle (P.4) is incompatible with the unrestricted fusion axiom (P.13) and, more generally, with any version of the strong fusion axiom (P.13 $_{\xi}$) in which the condition ‘ ξ ’ is satisfied by all interior and tangential proper parts of any given thing. For suppose we allow for such fusions. Then every entity would have an interior as well as a closure, and the following would hold:

- (91) $PP(ix)(cx)$.

By (P.4), this would imply

- (92) $\exists z(Pz(cx) \wedge \neg Oz(ix))$,

which in turn would imply

- (93) $\exists zPz(bx)$,

contradicting (C.5). Thus, $\bar{B}MT + (P.4) + (P.13)$ is inconsistent, as is any theory $\bar{B}X$ including (P.4) along with (P.13 $_{\xi}$) with ‘ ξ ’ as indicated.

This is not surprising, since the whole point of going boundary-free is, in the present context, to avoid the open/closed distinction, hence the distinction between interiors and closures reflected in (91). However, this means that (C.5) prevents the formulation of any reasonably strong theory unless we are willing to give up weak supplementation, and this may certainly be regarded as a major drawback of the approach.

In fact, one may consider both options here. One may (i) regard the compositional weakness of the theory as a necessary price to pay to preserve mereological supplementation and avoid the topological conundrums surrounding the open/closed distinction. But one may also (ii) go for a stronger theory with generalized or even unrestricted fusions, dismissing the conundrums precisely by forgoing supplementation. After all, it could be argued that the open/closed distinction is problematic *only* insofar as it is cashed out in terms of a discrimination between entities that do and entities that do not possess their boundaries, and in the absence of boundaries the discrimination dissolves. In the literature the first option is more widespread, its closest representative being the greatly influential Region Connection Calculus (*RCC*) originated with Randell *et al.* (1992)—a reductive extension of $\bar{BMT} + (P.5)$ with binary sums and complements. But the second option, which is closer to Whitehead's original approach, is also well represented, as evidenced by Clarke (1981), Biacino and Gerla (1991), and Asher and Vieu (1995) *inter alia* (the first three admitting unrestricted fusion, the latter admitting fusions of interior parts). Indeed, we have seen in Section 1.4.2 that the idea of restricting mereological fusion in order to avoid undesired entities is by itself suspicious. If we agree with the thought that a fusion is nothing over and above the things that compose it, then the intuitive problems raised by the open/closed distinction are hardly solved by eschewing formal commitment to such things as interiors and closures. i.e., fusions of interior and of tangential parts. For such parts are *all* there already (and, of course, one needs IPP and TPP to be distinct in order to state (C.5) in the first place). In this sense, option (ii) might be regarded as preferable on philosophical grounds, though the failure of supplementation would remain a hindrance.

Unfortunately, all of these theories include the monotonicity axiom as well as its converse, (C.4), i.e., they are all of the reductive sort, which makes it difficult to assess their relative pros and cons *vis-à-vis* the two options in question. In fact, not only do such theories include (C.4) and (C.5); they also rest on a *sui generis* characterization of the fusion operator in which C takes over the role of O , which makes it difficult to compare them to *KGEMT*. For example, in Clarke's theory, which goes as far as to include analogues of the Kuratowski axioms,




Figure 1.15. The interior of a closed disc x is a fusion, but not a fusion^{*}, of its parts.

the unrestricted fusion principle does not equal (P.13) but, rather, the following schema:

$$(C.13) \quad \exists w \phi w \rightarrow \exists z \forall w (Cwz \leftrightarrow \exists v (\phi v \wedge Cwv)) \quad \text{Topological Fusion}$$

This leads to a correspondingly *sui generis* fusion operator:

$$(94) \quad \Sigma^* x \phi x =_{df} \iota z \forall w (Cwz \leftrightarrow \exists v (\phi v \wedge Cwv)) \quad \text{fusion}^*$$

(Recall that *RMT* treats C as extensional: see (73).) And it is easily checked that Σ^* does not coincide with the operator Σ defined in (39): if parthood reduces to enclosure, then the interior of a closed entity x qualifies as the fusion of x 's proper parts, but not as their fusion^{*}. (This is because in the absence of boundaries the interior is sure to overlap any y that x overlaps, and vice versa, though it will to be disconnected from any z to which x is *externally* connected; see Figure 1.15). This amendment is plausible enough. But it means that the mereotopological operators defined in (40)–(44) and in (79)–(82) must be revised accordingly, and at the moment there is no systematic comparison between the behavior of such revised operators in a boundary-free theory and the behavior of the original operators in a boundary-tolerant theory. Just to give an example, note that re-defining ‘ \sim ’ in terms of Σ^* :

$$(95) \quad \sim^* x =_{df} \Sigma^* z Dzx \quad \text{complement}^*$$

implies that nothing is connected to its own complement and, therefore, that the universe is bound to be disconnected. Of course we can rely on a different notion of complement (as suggested in Randell *et al.* 1992: 168), or one can change the definition of self-connectedness in such a way as to avoid at least the latter consequence (as suggested in Clarke 1985: 69). But this is playing with definitions. The “old” notions, when revisited in terms of (94), continue to make good sense no matter how we change the official nomenclature, so we can hardly leave it at that.

Perhaps the best way to assess the strengths and weaknesses of these theories is to note that the departure from the ordinary conception of fusion affects the very distinction between open and closed entities. Consider the following variants of (79) and (85):

$$(96) \quad i^*x =_{df} \Sigma^* z \forall y (Czy \rightarrow Oxy) \quad \text{interior}^*$$

$$(97) \quad OP^*x =_{df} x = i^*x \quad \text{Open}^*$$

It is easily checked that any theory at least as strong as $\bar{B}RMT + (C.13)$ has the following theorem:

$$(98) \quad OP^*x \rightarrow \neg ECxy.$$

Thus, open* entities never touch anything. It is only closed* and semi-closed* entities (defined similarly), that can touch something without sharing any parts.¹² And if the open/closed distinction is replaced by the open*/closed* distinction, then the intuitive import of the relevant misgivings is up for grabs, and the choice between a conservative attitude towards mereological supplementation (option (i)) and a liberal attitude towards mereological composition (option (ii)) calls for independent thinking. We are no longer dealing with a partitioning of the domain into entities that do and entities that do not possess a boundary; we are dealing with a partitioning into entities that do and entities that do not connect externally.

Be that as it may, all of this suggests that a thorough comparison between these two strategies for construing boundaryless mereotopologies is a challenging task (Cohn and Varzi 2003). What is clear is that the strategies are mutually incompatible in spite of their common motivation and this, in all fairness to *KGEMT*, is disturbing. No matter how one feels about subtracting or adding elements to the domain, there is something puzzling in the thought that a topological “monstrosity” should be cured through mereological surgery. Indeed, philosophically this puzzling feature is especially striking when it comes to explaining the intended *interpretation* of these theories. As it turns out, both can be modeled on a domain with a standard point-set topology, interpreting ‘C’ as in (99) for type-(i) theories, and as in (100) for type-(ii) theories:

$$(99) \quad x \text{ is connected to } y \text{ if and only if the closure of } x \text{ and the closure of } y \text{ have a point in common.}$$

$$(100) \quad x \text{ is connected to } y \text{ if and only if } x \text{ and } y \text{ have a point in common.}$$

(See e.g. Randell *et al.* 1992: 167, Gotts 1996b and Pratt and Schoop 2000 *vs.* Clarke 1981: 205, Biacino and Gerla 1991, and Asher and Vieu

¹²Actually, Clarke’s definition of the closure operator does not exactly parallel (81). In our notation it reads as follows:

$$c^*x =_{df} \Sigma^* z \neg Cz(i^*x).$$

However, this peculiarity does not affect the main point made in the text.

1995, respectively). There is, of course, nothing wrong with this sort of models when it comes to proving the consistency or even the completeness of such theories. And there would be nothing wrong with (99) and (100) as genuine models if we were dealing with boundaryless theories of the moderate sort, as seen above. It is disturbing, however, that one can hardly do any better when it comes to theories that are meant to be radically eliminativist—when it comes to explaining how contact relations may obtain in a world that is truly lacking the topological glue provided by points, lines, and surfaces even in the realm of pure space. (Whether one *can* do better is an open question. For example, with reference to type-(i) theories, Bennett 1996a suggests that *RCC* can be interpreted by encoding it into the bimodal propositional modal logic $\mathbf{S4}_u$, though the encoding is imperfect, as shown in Aiello 2000, and its natural canonical model is itself topological, as evidenced in Renz 1998 and Nutt 1999. Likewise, Stell and Worboys 1997, Stell 2000, and Düntsch *et al.* 2001 provide algebraic interpretations of *RCC* that dispense with any reference to point-based topologies, but the ontological transparency of such interpretations is itself a delicate matter.)

2.5 Expressivity and ontology

Let us conclude this philosophical excursus on topology with some general considerations concerning the delicate interplay between the expressive power of a theory and its ontological presuppositions. Regardless of whether we rely on the full strength of *KGEMT* or on theories of weaker import, we have seen that the move from mereology to mereotopology represents an important step towards the formulation of an adequate model of our spatial competence. The mereological distinction between a whole and its parts is crucial, but so is the distinction between interior and tangential proper parts, or the distinction between a connected whole and a scattered one, and such distinctions are intrinsically topological. This is not to say that mereotopological concepts exhaust the picture; geometric and morphological considerations also play a significant role when it comes to practical matters. But there is no question that a great deal of our spatial competence is grounded on our capacity to “parse” the world in terms of parthood and connection relationships. The interesting question, rather, is whether such relationships can be fully captured by the formal behavior of the binary predicates ‘P’ and ‘C’ when characterized by means of formal principles of the sort that we have been discussing up to now, and to what extent the answer depends on one’s specific views when it comes to matters of ontology. Here are some indicative examples.




Figure 1.16. Four patterns of external connection, from weakest to strongest.

2.5.1 Modes of connection. So far we have followed the familiar course of explaining connection in general terms, i.e., irrespective of the size (dimension) of the relevant contact area. In introducing that notion, however, we have mentioned the possibility of distinguishing connection ties of different strength—e.g., ties involving a single point of contact (as between Colorado and Arizona), an extended portion of a common boundary (France and Germany), or an entire boundary (Vatican and Italy). Even without bringing in boundaries, one may want to draw such distinctions to fully grasp, for example, the difference between a whole consisting of two spheres that barely touch from the whole consisting of two halves of a single sphere: both wholes are self-connected, but the second is surely more firmly connected than the first. And these distinctions have ramifications. For example, since tangential parthood is defined in terms of external connection, we may want to distinguish those proper parts that barely touch the exterior from those that firmly touch it—and so on. Moreover, the number of distinctions grows with the dimensionality of the entities we consider. Figure 1.16 illustrates the four main patterns of external connection (no overlap) that can be distinguished in 2D space. But in 3D we might want to further distinguish, for example, two cubes barely touching at a vertex, two cubes barely touching along an edge, two cubes touching along a face, and so on. Now, can all such distinctions be expressed in terms of the mereotopological primitives ‘P’ and ‘C’ (or just ‘C’, if one goes reductive)?

As it turns out, within a sufficiently rich mereotopological theory such as *KGEMT* the answer is in the affirmative. To illustrate, with reference to Figure 1.16 the difference between the first two cases can be explained as follows. In both cases, x and y are connected; but whereas in the first case this simply means that one can go from x to y without ever going through the exterior of the connected sum $x + y$, in the second case it is also possible to go from x to y without ever leaving the interior of $x + y$. More precisely, let us distinguish between the connectedness of a whole (‘SC’) and the connectedness of its interior, and let us say that an object is firmly self-connected just in case the latter condition holds:

$$(101) \quad \text{FSC}x =_{df} \text{SC}x \wedge \text{SC}ix \quad \text{Firm Self-Connectedness}$$

Then we can say that two entities are firmly connected when they have parts that add up to a firmly self-connected sum:

$$(102) \quad \text{FC}xy =_{df} \exists wz(\text{P}wx \wedge \text{P}zy \wedge \text{FSC}(w + z)) \quad \text{Firm Connection}$$

This captures the difference between the second case of Figure 1.16, where the relevant connection relationship is firm, and the first, where it isn't. The stronger connection patterns corresponding to the third and fourth cases can then be defined by reference to the complement of $x + y$:

$$(103) \quad \text{CC}xy =_{df} \text{FC}xy \wedge \neg \text{FC}x(\sim(x + y)) \quad \text{Complete Connection}$$

$$(104) \quad \text{PC}xy =_{df} \text{FC}xy \wedge \neg \text{C}x(\sim(x + y)) \quad \text{Perfect Connection}$$

(Strictly speaking, (103) and (104) would call for refinements, owing to the possibility that x and y have internal holes; see Cohn and Varzi 2003 for a more general picture.) On this basis, the generalization to spaces of higher dimensionality is not difficult. For instance, in 3D space the difference between two entities touching at a point and two entities touching along an edge can be described by further distinguishing two patterns of non-firm connection, depending on whether the common boundary is atomic or a self-connected composite (a line segment, or curve).

Now, all of this is easy in *KGEMT*. When it comes to weaker theories, however, things get more difficult. For the definitions above involve the self-connectedness predicate along with the interior, complement, and fusion operators, all of which may be absent in a theory deprived of the necessary compositional strength. For instance, in a boundaryless theory with no open/closed (or open*/closed*) distinction, we can supply for the lack of the interior operator by redefining firm self-connectedness as follows (Bennett 1996b: 345):

$$(101') \quad \text{FSC}x =_{df} \forall y(\text{IPP}yx \rightarrow \exists z(\text{IPP}zx \wedge \text{P}yz \wedge \text{SC}z))$$

This makes it possible to go ahead with definitions (102)–(104) and capture the relevant distinctions in the 2D case. (See also Borgo *et al.* 1996 for a boundaryless theory in which FC is treated as primitive.) Yet it is not clear how one can capture the further distinctions available in 3D and higher-dimensional spaces without appealing explicitly to the dimensionality of the relevant boundaries or to special assumptions concerning the structure of space (see Gotts 1994a). And of course things get worse in a theory that lacks the complementation principle (P.6), for in that case, as we have seen, the notion of self-connectedness is already problematic. In short, the expressive power of a theory depends crucially

on the underlying ontology, which in turn is reflected in the strength of the relevant compositional and decompositional principles.

Similar considerations apply to further conceptual distinctions that may be deemed relevant in the context of spatial reasoning. Consider again the common-sense notion of contact exemplified by such statements as (55): the table is touching the wall. We have said that this notion is not topological but metric. This is actually true in *KGEMT*, assuming that both entities—table and wall—are treated alike, i.e., as both closed or both open. For *KGEMT* has the following two theorems:

$$(105) \quad ECxy \rightarrow (CLx \rightarrow \neg CLy)$$

$$(106) \quad ECxy \rightarrow (OPx \rightarrow \neg OPy)$$

By contrast, in a boundaryless theory the picture is different. For example, insofar as *RCC* admits of models satisfying (99), it can treat the table and the wall as genuinely connected as long as their closures overlap. This may be at odds with physics, but it captures the common-sense intuition. (Ditto for a moderate variant such as Cartwright’s—see again (90).) So which of these accounts is better depends on the seriousness with which we handle the spatial ontology of common sense. On the other hand, none of this should be taken to imply that *KGEMT* lacks the resources to account for the loose notion of connection countenanced by common sense. Suppose we understand this notion in the following sense: the table is touching the wall insofar as nothing *can* be squeezed between them. The metric flavor of this notion lies in its modal ingredient: to say that nothing can be squeezed between two objects is to say that they are “vanishingly close” to each other—that their relative distance is arbitrarily small. We can, however, define a predicate that captures this ingredient in mereotopological terms: we can say that two objects are at least loosely connected in the relevant sense when one is connected to the closure of every open neighborhood of the other:

$$(107) \quad LCxy =_{df} \forall z(OPz \wedge Pyz \rightarrow Cx(cz)) \quad \text{Loose Connection}$$

(See Asher and Vieu 1995 for a similar definition.) This relation captures the intuition that nothing can lie between two entities that touch, even when those entities are closed. And surely enough the definition is consistent with *KGEMT*. So it is not that *KGEMT* lacks the conceptual resources to do justice to common sense. It is, rather, that the relation of loose connection is bound to be empty in those models of *KGEMT* where the open entities form a dense ordering, i.e., where the following holds:

$$(C.14) \quad OPx \wedge OPy \wedge PPxy \rightarrow \exists z(OPz \wedge PPyz \wedge PPzy) \quad \text{Open Density}$$




Figure 1.17. An interior tangential part?

And whether all models should satisfy this axiom is, on the face of it, a question about the ontological make-up of the world. (One can argue that as long as the open/closed distinction holds, common sense only requires a denial of (C.14) when the variables are restricted to the range of ordinary entities, as opposed to their spatial receptacles. For a full account of the mereotopology of discrete space, see Galton 1999 and 2000, §2.6, and the generalizations in Li and Ying 2004.)

2.5.2 Dimensionality. Consider a second example. We have seen that in a boundaryless theory everything must have interior proper parts, as per (C.5). However, the notion of an interior part is itself, in a way, a relative one, depending on the dimensionality of the space we are considering (Galton 2004). In 1D space, the middle portion of a line segment y , or any portion x that does not extend to y 's extremities, would qualify as an interior part of y , so y itself would satisfy the axiom. As soon as y is embedded in 2D (or higher) space, however, all of its parts would be tangential, as they can all be connected to things to which y itself would be externally connected—e.g. a disc z (Figure 1.17). Thus, in such higher spaces y would not satisfy (C.5).

Now, this is not by itself a disturbing fact if we take boundaryless theories to reflect a general intuition to the effect that all entities in the domain are of equal dimensionality. On the other hand, this very intuition is philosophically problematic. It is puzzling that whether something exists—a line segment, for instance—should depend on the dimension of space, for one may want to declare one's ontological commitments while remaining neutral with respect to the difficult question of the dimensionality of space. Indeed, this becomes a necessity if one does not have the resources to address such a question in the first place. In ordinary point-set topology, one can say that a domain is of dimension $< n$ if and only if every open cover O_1, \dots, O_k can be refined to a closed cover C_1, \dots, C_k such that every point occurs in at most $n + 1$ of the C_i s. This characterization is not available in mereotopology unless one goes second-order, and surely it cannot be mimicked in a boundaryless theory that eschews the open/closed distinction. Gotts (1994b) and Bennett

(1996b) suggest a way to bypass such difficulties by means of a different characterization, which only requires \mathbf{P} (or rather: \mathbf{E}) to be closed under the operations of complementation and binary sum, but the adequacy of such proposals is an open question. (By contrast, Galton 1996 shows that an adequate characterization is available in a boundary-based theory such as *KGEMT*, or weaker variants thereof; see also the layered mereotopologies of Donnelly and Smith 2003 and Donnelly 2004). So, again, we see here how the expressiveness of the theory depends crucially on the ontology it countenances, which in turn may be hard to specify within the theory itself.

2.5.3 Counting the holes. Finally, consider an example that lies somewhere between the previous ones—the notion of a simply connected whole, i.e., intuitively, a whole with no holes. Topologically, as also from the perspective of common sense, this notion is just as significant as the notion of a self-connected whole. Just as there is a big difference between an apple and a bikini, there is a big difference between an apple and a donut. Indeed, topology is often defined, intuitively, as a sort of rubber-sheet geometry that focuses precisely on these two differences, ignoring shape, size, and all sorts of other spatial properties that concern geometry broadly understood. Now, we have seen that self-connectedness is easily defined in mereotopological terms, though the adequacy of the definition presupposes the complementation principle (P.6) and is not, therefore, entirely neutral from an ontological standpoint. What about simple connectedness? More generally, the definition of self-connectedness can be extended so as to classify every object in terms of the (maximum) number of self-connected parts of which it consists—what is sometime called its “separation number”. This can be done by defining a corresponding sequence of predicates, one for each positive integer, as follows:¹³

$$\begin{aligned} (108) \quad \text{SN}_1x &=_{df} \text{SC}x && \text{Separation}_1 \\ (109) \quad \text{SN}_{n+1}x &=_{df} \exists yz(x = y + z \wedge \neg \text{C}yz \wedge \text{SC}y \wedge \text{SN}_nz) && \text{Sep}_{\cdot n+1} \end{aligned}$$

Can we also provide a mereotopological characterization of the genus of an object, so as to classify it in terms of the number of holes it has? As it turns out, the answer to this question is in the affirmative, again provided that we assume the complementation principle (P.6), but this affirmative answer has interesting ontological ramifications.

Here is how the basic account goes (Gotts 1994a). Let us say that something has dissectivity n (n a positive integer) just in case it is self-

¹³Here ‘ $n + 1$ ’ indicates arithmetical addition, as opposed to mereological summation.




Figure 1.18. Simple connectedness and genus classification in 2D.

connected and can be decomposed into $n+2$ self-connected, disjoint parts with the following property: two of them, y and z , are not connected, whereas all the others are connected to both of them but disconnected from one another. Formally, this amounts to the requirement that there be two disconnected parts y and z that are connected by a remainder whose separation number is n ('connected by' in the sense of (59)):

$$(110) \quad \text{DS}_n x =_{df} \text{SC}x \wedge \exists yzw(x - (y + z) = w) \\ \wedge \text{BC}yzw \wedge \neg \text{C}yz \wedge \text{SN}_n w \quad \text{Dissectivity}_n$$

Then we can say that something is simply connected just in case its maximum dissectivity equals 1:

$$(111) \quad \text{SSC}x =_{df} \text{DS}_1 x \wedge \neg \text{DS}_2 x \quad \text{Simple Connectedness}$$

With reference to Figure 1.18, for example, only the left pattern is simply connected, for the others have higher dissectivity numbers. This is how it should be, and it can be checked that the definition would yield the right classification also with reference to objects of different dimensionality. A donut has dissectivity 2, so it is not simply connected; a solid ball is. Indeed, we can now use (110) also to provide a mereotopological characterization of the genus of an object. Something is of genus n , i.e., has n holes ($n \geq 0$), if and only if its maximum dissectivity is $n + 1$:

$$(112) \quad \text{G}_n x =_{df} \text{DS}_{n+1} x \wedge \neg \text{DS}_{n+2} x \quad \text{Genus}_n$$

Again, Figure 1.18 illustrates the definition in the 2D case, but it can be checked that (112) yields the correct classification also for objects of higher dimensionality: a solid ball has genus 0, a donut has genus 1, a pretzel has genus 2, and so on.

So this is the basic account, which fully answers the questions above: provided we work with a mereotopology that is closed under complementation and binary sums, we can define simple connectedness and, more generally, classify any object in terms of the number of holes it has. There are, however, two additional questions one may ask at this point. The first is whether this basic account can be refined so as to do justice to further distinctions that could be drawn in view of the dimensionality




Figure 1.19. A donut, a punctured block, and other deviant solids of the same genus.

issues discussed above. We may, for example, want to tell a genuine donut from the deviant cases in Figure 1.19, all of which have the same genus. And here, as one might expect, the answer depends more heavily on the strength of the theory. In *KGEMT* we can go quite far; in weaker theories we may not, as we may not be able to distinguish the various kinds of non-firm connection that are needed to operate the relevant discriminations. For example, Gotts (1994b, 1996a) has shown that the issue can eventually be settled within the boundaryless framework of *RCC*, but this result rests on various assumptions on the topology and dimensionality of the entities in the domain that cannot themselves be expressed in the language of the theory. More importantly, it rests on the interpretation of ‘C’ given in (99): two bodies are connected if and only if their closures have a point in common. We have seen that such an interpretation is of dubious legitimacy in a boundaryless ontology (as Gotts himself laments in 1996b). So, once again, we reach a point where the expressive power of a theory depends crucially on the ontological commitments that one is willing to make.

The second question is whether the basic account can be refined so as to do justice to further distinctions that could be drawn in view of the various ways—and there are many—in which an object can be perforated. And here it appears that even a strong, boundary-based theory such as *KGEMT* may show its limits. In fact, there is a sense in which the limits in question are not just the limits of the *mereotopological* approach of which the theory is expression; they are the limits of topology as a general theory of space. Let us focus on the 3D case.

For one thing, we have been speaking of ‘holes’ in the sense of perforations, but we may also want to classify a self-connected object in terms of the number of its internal ‘cavities’. To some extent this is easy: on the assumption that the universe \mathcal{U} is self-connected, it is sufficient to identify the number of internal cavities with $n - 1$, where n is the separation number of the object’s complement. In 2D space, this coincides with the genus of the object—there is no difference between a 2D perforation and a 2D cavity. In 3D the numbers may diverge: a solid donut has genus 1 but 0 cavities, since its complement is self-connected. Dropping the assumption on \mathcal{U} , we can express this as follows:

$$(113) \quad \text{IC}_n x =_{df} \text{SC}x \wedge \forall y (\text{SC}y \wedge \text{IPP}xy \rightarrow \text{SN}_{n+1}(y - x)) \quad \text{Cavity}_n$$




Figure 1.20. Cavities, donut-cavities, and donuts with donut-cavities.

This definition works for every object in any dimension (except, of course, for \mathcal{U}). But this is just the beginning. A cavity may come in different forms. It may be a solid cavity, so to say, but it may itself be donut-shaped. It may also have the shape of an irregular donut of the sort illustrated in Figure 1.19. Or it may be “knotted” in various fashions—as a trefoil knot, for instance, or a granny knot. Clearly such distinctions are not covered by (113). Moreover, consider an object with two donut-shaped cavities. The cavities may lie next to each other, so to say, or they may be interlocked like the rings of a necklace. Or consider an object with a perforation—a donut—which also has an internal, donut-shaped cavity. The perforation may go through the “hole” in the cavity or it may lie next to it. All of these and many others are distinctions that are easily described in words just as they can easily be depicted (Figure 1.20), and reflect significant differences in the spatial structure of the objects in question. It is far from clear, however, whether one can capture them in mereotopological terms.

Secondly, a perforation may come in different forms, too. It can be straight or it can be knotted, and the knot may or may not wrap around another perforation, just as it may or may not go through the “hole” of an internal, donut-shaped cavity. It can also branch in the middle, so as to have more than two openings. Indeed, it can branch in many different fashions, as it can “merge” in various ways with internal cavities of various kinds. Again, all of these possibilities reflect significant distinctions that are easily described in words and can easily be depicted, but it is far from clear whether one can give a proper characterization in mereotopological terms—even with the full strength of *KGEMT*. In some cases it is not even clear to what extent such a characterization should just parallel the standard topological account of such patterns. Standardly, for example, a block with two parallel, straight perforations is equivalent to a block with a single, Y-shaped perforation—both have genus 2 (and can be transformed into each other by mere elastic deformation). This much *can* be said in mereotopological terms, using (112) above. But here is where standard topological considerations might be regarded as inadequate for a good description of the spatial structure of ordinary objects, and of the intuitions underlying our spatial reasoning broadly




Figure 1.21. Same genus, different holes.

understood. The topological equivalence between such patterns—and between such patterns and many others; see Figure 1.21—appears to deliver a partial account of the relevant spatial structures, for the genus of an object only captures the *intrinsic* topology of the object, not the way it relates to the environment. To get a better picture it seems necessary to keep an eye on the holes, not just on the object. And this is obvious from the fact that in describing such patterns we tend to do so by describing the mereotopology of the holes and the way *they* relate to each other; we do not describe the objects themselves. We tend to treat holes as objects in their own right, as “negative objects” about which we can say exactly the same sorts of thing we say about ordinary, “positive” objects. And we count both sorts of objects in the same way: we count *two* straight perforations and *one* Y-shaped perforation.

If this is correct, then there are two things one can say. One can say (and accept) that the limits of mereotopology *vis-à-vis* such fine-grained distinctions are just the limits of topology, *mutatis mutandis*. Or one can say that the limits in question reflect precise ontological assumptions concerning the domain of application of the theory, specifically a dismissive attitude towards the ontological status of holes. This is not to say that holes are left out of the picture. Surely any theory with unrestricted fusions has room for such things, for mereologically speaking a hole is nothing but part of the object’s complement. Rather, the point is that mereotopology by itself says nothing specific about *which* parts of the complement qualify as holes. The boundaries of a hole simply cannot be determined by purely mereological or topological considerations. Of course, we have seen that mereotopology says nothing about the boundaries of material objects either. But draw such boundaries as you like, chose the objects you like, unless you also draw the boundaries of their holes (if any) you cannot get a full picture of the mereotopological structure of the objects themselves. And to draw the boundaries of something is to confer ontological dignity to it.

In Casati and Varzi (1994, 1999) it is argued that this alternative way of construing the limits of mereotopology has far reaching consequences. Suppose we take holes seriously: a hole in an object is something with well-defined boundaries. Then the fine-grained distinctions mentioned




Figure 1.22. The objects of Figure 1.21 have topologically different internal skins.

above can be recovered by looking at the mereotopological interplay between matter and void, at the properties of the boundary where an object comes into contact with its holes. More precisely, let the interface between two entities x and y be the product of their boundaries:

$$(114) \quad x|y =_{df} \mathbf{b}x \times \mathbf{b}y \quad \text{Interface}$$

And let the internal skin of an object x be the interface between x and the fusion of its holes. Using ‘H’ for the binary relation ‘is a hole in’, this can be defined as follows:

$$(115) \quad \mathbf{s}x =_{df} \Sigma z \exists y (\mathbf{H}yx \wedge z = y|x) \quad \text{skin}$$

Then the distinctions in question are distinctions that reflect the mereotopology of an object’s skin. With reference to Figure 1.21, for example, it can be checked that the skin of the doubly perforated block on the left is the disconnected sum of two cylinders, i.e., topologically, two spherical surfaces with two punctures each. By contrast, the block with a Y-shaped hole has a connected skin that is equivalent to a spherical surface with three punctures, while the other blocks have skins equivalent to a torus with two punctures and to a bitorus with one puncture, respectively (Figure 1.22). (Note that a puncture is not a hole but a mere boundary. The surfaces of the objects in Figure 1.21 do not have boundaries, yet their internal skins do—and that makes all the difference.)

Now, in a boundaryless theory all of this is beyond reach. But in a sufficiently strong boundary-based theory such as *KGEMT* the notion of an internal skin is perfectly meaningful and well defined for every object of positive genus, and its mereotopological classification does not present any special challenge. This confirms once again the greater expressive power that comes with an ontological commitment to boundaries. It also shows, however, that such a commitment is not enough: the existential quantifier in definition (115) shows that an explicit commitment to holes is also needed. To the extent that the binary predicate ‘H’ is to be treated as a primitive, it is clear that this requires a step beyond *KGEMT* and its pure mereotopological extensions. (For an axiomatic treatment of ‘H’ and of its interplay with ‘P’ and ‘C’, see Casati and Varzi 1994, Appendix, and Varzi 1996b.)

3. Location Theories

Let us finally turn to the relation of spatial location. Intuitively, this is the relation that holds between an entity and the spatial region that it occupies, and we have already seen that this relation can hardly be reduced to a chapter of mereology and/or topology. Even if it were—as someone inclined to favor a Leibnizian, relationist conception of space against its Newtonian, substantivalist foes would urge—methodological prudence would suggest that we regard the reduction as a theorem, not as a starting point, hence that the relation of location be treated as an independent primitive next to parthood and connection. Exactly how this relation should be characterized, and how it should interact with the principles governing those other primitives, is precisely the sort of question that a good theory of location should aim to answer.

3.1 Varieties of Location

Before looking at the main options, the usual terminological caveats are in order. As with ‘part of’ and ‘connected to’, locative predicates have various meanings in ordinary language and it is important to be explicit.

For one thing, we often speak so as to specify the location of an object by reference to another object, as opposed to a spatial region. Consider:

- (116) The biceps muscle is located in the arm.
- (117) The parking area is located next to the stadium.
- (118) The elevator is located inside the main building.

Pretty clearly, such cases are of no special interest, as they reflect different ways of asserting mereotopological relations of the familiar sort: in (116) the locative predicate is just a variant of ‘part of’, in (117) it expresses the relation of external connection, and in (118) it stands for a relation of containment that can be cashed out in terms of interior proper parthood. Of course, establishing mereotopological relations may be an *indirect* way of specifying a genuine location: insofar as the biceps muscle is part of the arm, for instance, the muscle is bound to be located within the region occupied by the arm, though this is by itself an intuition that needs to be spelled out carefully (Section 3.3). Moreover, not every case of relative location can be explained in this fashion (Section 3.4). For the moment, let us just emphasize that the main concern of a theory of spatial location as we understand it here is with those cases in which an object’s location is specified *directly*, as in

- (119) The peak of Mount Everest is located at 27°59′ N 86°56′ E.
- (120) The new library will be located at this site.

It is in this sense that the theory presupposes an ontology that includes spatial regions as *bona fide* entities in their own right. Indeed, we shall assume that the location primitive is a relation whose second argument can *only* be a region of space—a “place”. (Never mind the question of what sort of linguistic expressions can serve the purpose of referring to places, as opposed to things that *have* a place. Statements such as

- (121) The bookcase is located in the living room.
 (122) The United Nations are located in Manhattan.

are somewhat ambiguous in this respect, but we may suppose that the context will always suffice to determine the intended meaning.)

Secondly, there are various ways in which an object may be said to be located at a region. In a very loose sense, I am located at any region that is not completely free of me (this room, or even the adjacent dining room if I am reaching a foot out of the doorway); in a stricter sense, I am only located at those regions that host me entirely (this room, if I am not reaching out of the doorway); and in a stricter sense still, I am only located at one region, namely the region that corresponds exactly to the volume of my body. In the following we shall use ‘located at’ as designating the last, strictest relation; the weaker relations can be introduced by definition. More precisely, suppose we use ‘L’ for the predicate of exact location. Then three additional predicates can immediately be defined as follows (from Parsons 2006¹⁴):

- (123) $GLxy =_{df} \exists z(Ozy \wedge Lxz)$ Generic Location
 (124) $ELxy =_{df} \exists z(Pzy \wedge Lxz)$ Entire Location
 (125) $ULxy =_{df} \exists z(Pyz \wedge Lxz)$ Ubiquitous Location

Thus, I am generically, in fact entirely located in Manhattan, but not ubiquitously (or exactly) located there; I am generically, in fact ubiquitously located at the region occupied by my left arm, but not entirely (or exactly) located there; and if I reach an arm in my neighbor’s window, then I am generically, but neither entirely nor ubiquitously (let alone exactly) located at the region corresponding to her living room. As we shall see, under suitable conditions these predicates are interdefinable, so the choice of ‘L’ as a primitive is ultimately immaterial.

Finally, it goes without saying that the location of an object may change over time. This would suggest treating ‘located at’, not as a binary predicate, but as a three-place predicate involving a spatial as well

¹⁴Parsons’s term for generic location is ‘weak location’, and his term for ubiquitous location is ‘pervasive location’. Our different terminology is dictated merely by notational convenience, in view of the predicates ‘whole location’ (WL) and ‘proper location’ (PL) introduced below.




Figure 1.23. Basic locative relations.

as a temporal argument (or as a temporally indexed binary predicate). However, we have seen that the same goes for parthood and connection: unless one accepts a radical form of mereotopological essentialism, an object may in principle change its parts or its topological relations without ceasing to exist. In the preceding sections we have tried to keep things simple by treating ‘P’ and ‘C’ as binary predicates, and we shall do the same with our location primitive ‘L’. In a way, this means that we are assuming the relevant time to be fixed throughout. But one may also consider an alternative reading, to the effect that the variables of the theory range over four-dimensional entities extended in space-time. (See again the brief discussion following (32), Section 1.3.2.) Not much of what follows depends on the strategy one favors, but for simplicity we shall continue to speak of the location of an object as a 3D region of space rather than—possibly—a 4D region of space-time.

3.2 Basic Principles

With these conventions in place, let us officially expand our formal language by adding the new binary predicate ‘L’, intuitively understood as the relation of exact location holding between an object and a region of space. To make this interpretation explicit, we may begin by assuming the following axiom:

$$(L.1) \quad Lxy \wedge Lxz \rightarrow y = z \quad \text{Functionality}$$

This guarantees that nothing can have more than one exact location, which is all that is needed to justify the definitions in (123)–(125) (Figure 1.23.). Indeed, it is easy to see that in the presence of an extensional mereology, (L.1) has the following corollaries:

- (126) $Lxy \leftrightarrow (ELxy \wedge ULxy)$
- (127) $ULxy \leftrightarrow (GLxy \wedge \forall z(Ozy \rightarrow GLxz))$
- (128) $ELxy \leftrightarrow (GLxy \wedge \forall z(GLxz \rightarrow Ozy))$

Thus, although we have settled on the strictest possible primitive, the predicate ‘L’, one could equally well settle on the weakest predicate, ‘GL’, and define the rest *via* (126)–(128). More precisely, if the mereological theory in the background is at least as strong as *EM*, it turns out that the L-based system defined by (L.1) plus (123)–(125) is equivalent to the GL-based system defined by (126)–(128) plus the following:

$$(L.2) \quad \text{GL}xy \rightarrow \exists z\text{L}xz \qquad \text{Exactness}$$

(See Parsons 2006; for a different choice of primitives, compare also Perzanowski 1993.) Note that the conjunct ‘GL xy ’ is redundant in (127) as long as O is reflexive. However, this extra conjunct is needed in (128) unless one assumes that everything is located somewhere:

$$(L.3) \quad \exists y\text{L}xy \qquad \text{Spatiality}$$

This is clearly an assumption that reflects a substantive thesis (a central tenet of most nominalistic ontologies), so it is fair to keep it separate.

To be sure, there is a sense in which (L.1) may also be read as a substantive thesis: functionality is reasonable only to the extent that we are thinking of so-called “particular” entities, entities such as material bodies or events, as opposed to “universal” entities such as properties. That a material body cannot be in two places at once was already a central thesis of Aristotle’s theory of location (Morison 2002). But Aristotle also held the view that universals, too, exist in space and time: they exist wherever and whenever they are exemplified. Wisdom, for example, exists whenever and wherever there are wise people—and whenever a wise person exists, wisdom exists in its entirety wherever that person is located. Wisdom can therefore be multi-located, and the same goes for all universals. Since this view is still very popular (Armstrong 1989), the functionality principle (L.1) would be objectionable. However, we can bypass this issue by taking ‘L’ to represent the location relation that is characteristic of particulars. In that sense, (L.1) is, if not a conceptual truth, a perfectly reasonable starting point.

What else is needed in order to fix the intended meaning of ‘L’ in this sense? Since the idea is that every object must be located at a *region*, some restriction must be imposed on the second argument of the relation. This would be a trivial task if the language contained an explicit predicate for regions. However, we can make do without such a predicate and try to characterize regionhood directly in terms of suitable axioms on L. There are two options here, depending on whether we think that spatial regions are themselves entities located somewhere. If we think so, then the obvious thing to say is that such entities can only be located at themselves (Casati and Varzi 1999: 121):

$$(L.4) \quad Lxy \rightarrow Lyy \qquad \text{Conditional Reflexivity}$$

This would immediately imply that no distinct regions can be exactly co-located, i.e., effectively, located at each other:

$$(129) \quad Lxy \wedge Lzw \wedge Lyw \rightarrow y = w$$

Moreover, given (L.1), conditional reflexivity would ensure that **L** is both antisymmetric and transitive:

$$(130) \quad Lxy \wedge Lyx \rightarrow x = y$$

$$(131) \quad Lxy \wedge Lyz \rightarrow Lxz$$

It follows that relative to the sub-domain of regions **L** would behave as a partial ordering. By contrast, if we think that regions do not have a location—they *are* locations—then the obvious option is given by:

$$(L.5) \quad Lxy \rightarrow \neg Lyz \qquad \text{Conditional Emptiness}$$

In this case, the restriction of **L** to the class of regions would again qualify as a partial ordering—a strict ordering—but only in a trivial sense: effectively, it would just collapse to the empty relation.

There is, arguably, no deep metaphysical issue behind these two options: both (L.4) and (L.5) are equally good stipulations, and the difference would disappear as soon as we focus on cases of *proper* location:

$$(132) \quad PLxy =_{df} Lxy \wedge \neg Lyx \qquad \text{Proper Location}$$

However, there are some differences that are worth mentioning. For one thing, given (L.1), the first option makes it possible to define regionhood in a perfectly straightforward way:

$$(133) \quad Ry =_{df} \exists x Lxy \qquad \text{Region}$$

(We are speaking of regions in a broad sense, including boundaries as limit cases.) By contrast, (L.5) would support this definition—and variants thereof—only on the assumption that there are no unoccupied regions, i.e., regions that fail to correspond to the location of some object or event. To put it differently, if all location is proper location, it is not possible to define regionhood unless the following principle is accepted:

$$(L.6) \quad Ry \rightarrow \exists x PLxy \qquad \text{Fullness}$$

And philosophically this principle is just as controversial as the spatiality principle (L.3). (Among other reasons, one might want to allow for

boundary-like regions while rejecting the existence of boundary-like objects, as seen in Section 2.4.3.) Secondly, it is also apparent that the two options differ with regard to (L.3) itself: (L.4) is compatible with this principle, (L.5) isn't. Thus, the second option makes it impossible to assert the thesis that everything is located somewhere—a thesis which, albeit controversial, is certainly not inconsistent. Again, this is a limitation that would dissolve if 'R' were available in the language, in which case the thesis in question could be reformulated as follows:

$$(L.3') \quad \neg Rx \rightarrow \exists y Lxy$$

But precisely because 'R' cannot be defined absent (L.6), the limitation is not immaterial.

For these reasons, in the following we shall favor the first option and assume the conditional reflexivity axiom (L.4). Together with the functionality postulate (L.1), this yields a minimal theory of (exact) spatial location, which we shall label *S*: this theory is incompatible with (L.5), but it includes the exactness principle (L.2) as a theorem and can be strengthened by adding the spatiality principle (L.3), the fullness principle (L.6), or both.

At this point, we could consider various ways of strengthening *S* (or its extensions) by imposing suitable axioms on the predicate 'R' defined in (133). For example, it seems reasonable to assume that regionhood is both dissective and cumulative, i.e., that any part of a region is itself a region, and that the sum of any regions (if it exists) is a region, too:

$$(L.7) \quad Ry \wedge Pxy \rightarrow Rx \quad \text{Dissectiveness}$$

$$(L.8) \quad z = \Sigma x \phi x \wedge \forall x (\phi x \rightarrow Rx) \rightarrow Rz \quad \text{Cumulativity}$$

It may also be reasonable to consider additional postulates concerning whether the class of all regions forms a dense domain, or whether it forms an atomless, possibly a boundaryless domain as opposed to an atomistic domain every element of which consists (intuitively) of points:

$$(L.9) \quad Rx \wedge Ry \wedge PPxy \rightarrow \exists z (Rz \wedge PPxz \wedge PPzy) \quad \text{R-Density}$$

$$(L.10) \quad Rx \rightarrow \exists y (Ry \wedge PPyx) \quad \text{R-Atomlessness}$$

$$(L.11) \quad Rx \rightarrow \exists y (Ry \wedge IPPyx) \quad \text{R-Boundarylessness}$$

$$(L.12) \quad Rx \rightarrow \exists y (Ry \wedge Ay \wedge Pyx) \quad \text{R-Atomicity}$$

More generally, it may be reasonable at this point to consider whether the domain of regions should be closed under various mereotopological principles, regardless of whether such principles hold of the entities that may occupy those regions. For example, even an anti-extensionalist about material objects will presumably deny that different regions may

consist of the same proper parts, and even those who have misgivings about strong composition principles for arbitrary objects might be happy to endorse unrestricted composition of spatial regions. All such extensions of S are obviously worth examining, and they are crucial if we want to fix the intended range of the relational predicate ‘L’, but there is no need here to review all the options: suffice it to say that the availability of ‘R’ makes it possible to examine them in a systematic fashion. (An interesting question, for instance, is whether one can provide a purely mereotopological characterization of Euclidean space; see Tsai 2005: §7.4, for a negative answer). Rather than focusing on the structure of space *per se*, let us see how S can be further extended by considering more closely the relationship between the two terms of the location relation—the structure of regions *and* the structure of their tenants.

3.3 Mirroring Principles

To this end, let us begin by noting that the four relations L, GL, UL, and EL do not exhaust all the options. Additional locative relations can be specified by replacing the plain mereological predicates in (123)–(125) with finer-grained mereotopological predicates—for example:

- | | | |
|-------|---|---------------|
| (134) | $\text{TEL}xy =_{df} \exists z(\text{TPP}zy \wedge \text{L}xz)$ | Tangential EL |
| (135) | $\text{IEL}xy =_{df} \exists z(\text{IPP}zy \wedge \text{L}xz)$ | Interior EL |
| (136) | $\text{TUL}xy =_{df} \exists z(\text{TPP}yz \wedge \text{L}xz)$ | Tangential UL |
| (137) | $\text{IUL}xy =_{df} \exists z(\text{IPP}yz \wedge \text{L}xz)$ | Interior UL |

More generally, given any mereotopological relation ψ , there is a corresponding locative relation L_ψ defined by:

$$(138) \quad \text{L}_\psi xy =_{df} \exists z(\psi zy \wedge \text{L}xz) \quad \psi\text{-Location}$$

(Thus, $\text{GL} = \text{L}_\text{O}$, $\text{EL} = \text{L}_\text{P}$, $\text{UL} = \text{L}_\text{P}^\text{S}$, etc.) Such generalizations are straightforward, but they bear out that the language of location can be as rich as the underlying mereotopological vocabulary. Indeed, this is only half of the story. According to (138), to be ψ -located at a region y is to be exactly located at some region, z , that is ψ -related to y . Thus, the resulting variety of locative relations is defined with reference to the mereotopological structure of the *range* of L—the structure of space. But one could also consider the obvious alternative, and characterize locative relations by reference to the mereotopological structure of the *domain* of L—the structure of space’s tenants. In this alternative sense, to be ψ -located at a region y is to be ψ -related to some object, z , that is exactly located at y :

$$(139) \quad \text{L}^\psi xy =_{df} \exists z(\psi xz \wedge \text{L}zy) \quad \psi\text{-Location (2)}$$

Now, the interesting question is whether these two ways of characterizing locative relations should coincide—whether ‘ L^ψ ’ and ‘ L_ψ ’ should always stand and fall together. This is trivially true when the relata are of the same kind, i.e., regions, for in that case locative relations collapse to mereotopological relations in view of the following S -theorems:

$$(140) \quad Rx \rightarrow (L_\psi xy \leftrightarrow \psi xy)$$

$$(141) \quad Rx \rightarrow (L^\psi xy \leftrightarrow \psi xy)$$

But what about the general case? This is not just an interesting question to ask if we want to get the map straight; it is also a question that calls for interesting philosophical decisions.

To address the matter properly, let us consider the two directions of the equivalence separately, corresponding to the following theses:

$$(L.13) \quad L^\psi xy \rightarrow L_\psi xy \qquad \text{Bottom Mirroring}$$

$$(L.14) \quad L_\psi xy \rightarrow L^\psi xy \qquad \text{Top Mirroring}$$

3.3.1 Bottom mirroring and co-location. Informally, the first of these theses says that the structure of space should be at least as rich as the structure of its tenants. For example, suppose I am exactly located at region r . Then my right foot, which is part of me, is L^P -located at r . But if my foot is part of me, then it is reasonable to suppose that its exact location is part of my exact location, which is precisely what (L.13) would imply: my foot is L_P -located (i.e., entirely located) at r . For another example, since my foot is connected to the rest of my body, which is located at a certain region r' , then it is reasonable to suppose that the location of my foot is connected to r' , too: my foot being L^C -located at r' implies its being L_C -located at r' . In general, adding (L.13) to S would secure that if two objects are ψ -related (where ψ is P , C , or any other mereotopological relation), then so are their locations:

$$(142) \quad Lxy \wedge Lzw \wedge \psi xz \rightarrow \psi yw.$$

Plausible as all this might sound, it is however easy to see that (L.13) does not generally hold. A simple counterexample is depicted in Figure 1.24, left. In this model there is just one (atomic) region, r , and a complex object, a , consisting of two (atomic) parts, b and c . Object a is located at r , hence its parts b and c are L^P -located at r . Yet they are not L_P -located at r because there is no part of r at which they are exactly located. Note that the mereotopological relations in this model can easily be extended so as to satisfy all the axioms of $KGEMT$, so the counterexample does not depend on the strength of the underlying mereotopological theory. It depends exclusively on the behavior of L .




Figure 1.24. Three violations of Bottom Mirroring. (Location relationships are represented by gray arrows; reflexivities omitted.)

Now, it might be observed that this model (or its *KGEMT* closure) would be ruled out if we assumed that every part of a spatially located entity had a spatial location—a restricted form of the spatiality principle (L.3). More generally, it might be thought that in order to justify (L.13) one must assume the following:

$$(L.15) \quad \psi xz \wedge Lzy \rightarrow \exists wLxw \quad \text{Conditional Spatiality}$$

Bottom mirroring should not hold *holus bolus*, but only when x and its ψ -relata are genuine spatial entities (in which case (L.13) turns out to be equivalent to (142)). However, it is easy to see that even $S + (L.15)$ would fail to warrant this claim. A counterexample is depicted in Figure 1.24, center. Here b and c are L^{PP} -located at r , since each is a proper part of a , which is located at r . Yet there is no proper part of r at which a or b is located, which is to say that neither is L_{PP} -located at r .

We begin here to see the hidden force of (L.13). By requiring that the structure of space mirror the structure of its tenants, this principle rules out the possibility that mereotopologically distinct entities be co-located, i.e., located at the same region. Of course, it may be difficult to imagine a concrete scenario corresponding to the model in question, and certainly very difficult to provide a less abstract representation of it. But it is not difficult to provide a good example if, for instance, we give up mereological extensionality. Consider again Tibbles, the cat, and the “mere” mereological sum of his tail with the rest of his body, Tib + Tail (Section 1.3.2). Obviously, even if one treated these entities as distinct, one would still like to say that they share one and the same location. Yet, if ‘ ψ ’ expresses the relevant mereotopological relationship of distinctness *cum* sameness of proper parts, it is obvious that the region to which Tibbles (or Tib + Tail) bears the relation L^ψ is not the region to which it bears the relation L_ψ (Figure 1.24, right). Moreover, even in the presence of mereological extensionality one may conceive of situa-

tions where spatial co-location seems possible. We have already seen, for instance, that according to Chisholm’s (1984) Brentanian theory, topological connection is explained precisely in terms of spatial coincidence of boundaries: boundaries are *located in* space but do not *occupy* space, hence they can coincide while being distinct (Section 2.4.2). For a scenario compatible with the full strength of *KGEMT*, consider Davidson’s (1969) example concerning event identity: arguably the rotation and the getting warm of a metal ball that is spinning fast are two events, yet they occur exactly in the same region and they share that location with the ball itself. Finally, if our ontology is rich enough to include immaterial or otherwise ethereal creatures for which genuine interpenetration is possible, then again co-location seems conceivable. Already Leibniz mentioned shadows as a case in point (*New Essays*, II-xxvii-1); other candidates include clouds (Shorter 1977), holes (Casati and Varzi 1994), ghosts (van Inwagen 1990: 81), and even angels (Lewis 1991: 75).

These are just some examples. But they are indicative of the many philosophical motives that may lie behind a rejection of the principle according to which proper spatial co-location is impossible—a principle that can be put thus:

$$(L.16) \quad \text{PL}xy \wedge \text{PL}zy \rightarrow x = z \qquad \text{Exclusiveness}$$

Giving up this principle involves giving up (L.13), at least in its general form. And the question of what special instances one should posit, i.e., what values of ‘ ψ ’ and ‘ x ’ satisfy bottom mirroring, calls for a detailed case-by-case investigation—and for explicit ontological decisions.

3.3.2 Top Mirroring. Consider now the converse of (L.13), namely the principle of top mirroring, (L.14). Informally, this says that the structure of space should be mirrored in the structure of those entities that inhabit it. For example, if the location of my body properly includes a region r , then it is reasonable to suppose that my body properly includes something located at r : L_{PE} -location, hence L^{PE} -location. Pretty clearly, however, there are numerous relations ψ for which this sort of implication appears problematic. The location of my body is a proper part of any region that includes this room, but there is no obvious reason to think that every such region is the location of some existing object—no reason to think that L_{PP} -location implies L^{PP} -location. (Compare Figure 1.25, left.) Similarly, my body is L_{EC} -located at many regions, viz. regions externally connected to my body’s current location; yet there is no obvious reason to think that my body is L^{EC} -located at those regions, since many of them may be (partly) empty (Figure 1.25, center.) Just as bottom mirroring appears to presuppose the spatiality




Figure 1.25. Three violations of Top Mirroring.

principle, or at least its restricted variant (L.15), top mirroring appears to presuppose fullness, or at least the following restricted version:

$$(L.17) \quad Ry \wedge \psi zy \wedge Lxz \rightarrow \exists wPLwy \quad \text{Conditional Fullness}$$

And this is a substantive presupposition that few might grant, regardless of their views concerning purely mereotopological matters.

Indeed, even when ψ is the seemingly innocent relation of proper extension, as in our first example, (L.17) appears problematic. Misgivings about this principle—hence about the corresponding instance of top mirroring—come in various forms. First of all, there are arguments that purport to show that the principle is empirically false (Parsons 2000). Second, the principle rules out *a priori* the possibility of spatially extended mereological atoms, as in Figure 1.25, right. To the extent that one can conceive of such things, it is argued, it should not be a conceptual truth that every region ubiquitously occupied by an object is exactly occupied by a part of that object. (See Markosian 1998a; other contemporary philosophers who endorse the possibility of extended simples include Parsons 2004 and Simons 2004, but the view goes back to Democritus’s claim that atoms come in an infinite variety of shapes and sizes.) Third, there are arguments to the effect that (L.17) sits ill with the thought that ordinary material bodies can gain or lose some parts (van Inwagen 1981). To illustrate, consider again Tibbles, the cat whose tail gets annihilated at t , and suppose we agree that it survives the accident. Prior to t , (L.17) would suggest that in addition to the whole cat there exist also two externally connected proper parts: Tail and Tib (the remainder). Now consider the following statements:

- (143) Tib (before t) = Tib (after t)
- (144) Tib (after t) = Tibbles without Tail (after t)
- (145) Tibbles without Tail (after t) = Tibbles with Tail (before t)
- (146) Tibbles with Tail (before t) \neq Tib (before t)

These four statements are jointly inconsistent, unless one is willing to give up the transitivity of identity, so something must give. (146), how-

ever, is trivially true: there is no way one could identify a whole cat with its tailless portion. And (145) is true by assumption: to give it up is to deny that Tibbles survives the accident, unless one is willing to construe cats as four-dimensional entities whose temporal parts are numerically distinct (Heller 1984). As for (144), its denial would obviously incur in a commitment to properly co-located entities, let alone a violation to mereological extensionality (Wiggins 1968). Thus—the argument goes—unless one is ready to accept such unpalatable consequences, the only option is to give up (143): that identity is false for the simple reason that prior to the accident Tib does not even exist; it only exists after the accident, and it exists as tailless Tibbles.¹⁵ (One might also say that before the accident Tib does not *actually* exist. The view that undetached parts are mere “potential entities” has been the focus of an intense debate in early modern philosophy; see Holden 2004. Brentano 1933 endorsed a similar view, too, and some authors have applied it explicitly to the puzzle in question—e.g., Smith 1994, §3.5.)

Whether any such arguments are found compelling is, of course, an open issue. Nonetheless, the obvious moral is that (L.14) can hardly be regarded as a conceptual truth about location. Even if we confine ourselves to its single, *prima facie* plausible instance in which ψ is the relation of (proper) extension, i.e., the converse of (proper) parthood, top mirroring is a substantive metaphysical thesis whose addition to S must be independently motivated.

3.3.3 Further locative relations. In discussing such matters, it is useful to keep in mind that the lack of a full correspondence between the mereotopology of space and the mereotopology of its tenants may find expression in the failure of other principles or equivalences that might otherwise suggest themselves. Consider, for instance, the following relation (from Parsons 2006, §4):

$$(147) \quad \text{WL}xy =_{df} \forall z(\text{P}zx \rightarrow \text{GL}zy) \qquad \text{Whole Location}$$

Intuitively, this says that an object is located in this room (for instance) if every part of the object is generically located in this room, i.e., if none of it is missing from the room. This might sound like a different way of saying that the object is entirely located in this room, in the original sense of (124) (corresponding to L_P -location) or in the alternative sense of (139) (L^P -location). In fact, however, all these notions are distinct.

¹⁵The puzzle raised by (143)–(146) has been introduced to contemporary discussion by Wiggins (1968), but it goes back at least to the Stoics; see e.g. Sorabji (1988: §1.6). For a detailed overview, see Simons (1987, §3.3) and the introduction to Rea (1997).

Not only do L_P -location and L^P -location come apart, as seen above. They also differ from whole location: the diagram in Figure 1.24, left, corresponds to a model in which an object, a , is both L_P - and L^P -located at a region, r , in spite of not being wholly located there; the diagram in Figure 1.25, right, depicts a model in which a (an extended atom) is wholly located at region r_1 in spite of being neither L_P - nor L^P -located there.¹⁶ (Note that both of these models are extensional, and would continue to exhibit these features even if closed under every *KGEMT*-axiom.)

The notion of whole location is just one example. In general, for any mereotopological relation ψ , one can define two additional locative relations besides L_ψ and L^ψ , obtained by switching to a universal quantifier and replacing ‘L’ by ‘GL’:

$$(148) \quad L_{\forall\psi}xy =_{df} \forall z(\psi zy \rightarrow GLxz)$$

$$(149) \quad L^{\forall\psi}xy =_{df} \forall z(\psi xz \rightarrow GLzy)$$

Here $WL = L_{\forall\check{p}}$, and it should be obvious from our single example that the equivalence between the new predicates in (148)–(149) and the old predicates in (138)–(139) is generally an open question. Indeed, at this point we get the full picture by further generalizing these four basic patterns in the obvious (recursive) way: if λ is *any* locative relation, then so are the following:

$$(150) \quad L_{\exists\psi\lambda}xy =_{df} \exists z(\psi zy \wedge \lambda xz)$$

$$(151) \quad L^{\exists\psi\lambda}xy =_{df} \exists z(\psi xz \wedge \lambda zy)$$

$$(152) \quad L_{\forall\psi\lambda}xy =_{df} \forall z(\psi zy \rightarrow \lambda xz)$$

$$(153) \quad L^{\forall\psi\lambda}xy =_{df} \forall z(\psi xz \rightarrow \lambda zy)$$

There are lots of redundancies and empty relations in this picture, whose complexity depends significantly on the mereotopological axioms governing ψ . Nonetheless, it is only through a careful study of such intricacies—and of the corresponding mirroring principles, at the moment vastly unexplored—that a reasonably complete theory of location can emerge.

3.4 Relative locations

The axiomatization of the region predicate ‘R’ and the positing of suitable mirroring principles constitute the two main directions in which theory *S* can be extended. Let us briefly mention a third direction, whose ramifications span philosophical and methodological issues alike. We have said that in ordinary language location is often understood as

¹⁶In Casati and Varzi (1999, §7.2), entire location is called ‘whole location’ and labelled ‘WL’.

a relation between two objects, as opposed to an object and a spatial region, and we have said that such understandings need not be taken to express any fundamental relationships: often, such “relative” locations are mereotopological relations in disguise, as in examples (118)–(120). There are, however, cases that resist this sort of explanation. Consider:

- (154) The brain is located inside the cranial cavity.
- (155) The swimming pool is located behind the house.
- (156) The bus stop is located right across the old oak tree.

Surely the truth conditions of these statements can hardly be explained in terms of mereotopological relations, so one can hardly leave it at that. Can S be strengthened so as to account for such cases as well?

In a way, the answer is straightforward. Consider (154). Although there is no direct mereotopological relationship between the brain and the cranial cavity (unless one thinks the latter literally surrounds the former), one can still explain the relevant truth conditions by reference to the mereotopology of the corresponding spatial locations: statement (154) is true if, and only if, the location of the brain is an interior proper part of the location of the cranial cavity. Equivalently, (154) is true if and only if the brain is located entirely in the interior of the spatial region occupied by the cranial cavity. This suggests that cases such as this can easily be accommodated within the present framework. When we say that a “target” object, x , bears a certain locative relation to a “reference” object, y , we mean to say that x bears that relation to y ’s place. To make this clear, let us introduce an explicit location functor, whose uniqueness follows directly from the functionality axiom (L.1):

$$(157) \quad \rho x =_{df} \iota y \mathbf{L}xy \quad \text{place}$$

(See Donnelly 2004 for a location theory with \mathbf{p} treated as a primitive.) Then, for any locative relation λ and any spatial objects x and y , we can define a corresponding predicate of relative location ‘ $\mathbf{R}\lambda$ ’ as follows:

$$(158) \quad \mathbf{R}\lambda xy =_{df} \lambda x(\rho y) \quad \text{Relative Location}$$

In the brain-cavity case, λ is the relation IEL defined in (135), i.e., $\mathbf{L}_{\mathbf{IPP}}$, but the same pattern would apply to a large variety of other cases, provided the reference object has a location somewhere. In fact, the same pattern can be applied to account for the initial examples in (116)–(118), too: the biceps muscle is $\mathbf{RL}_{\mathbf{P}}$ -located at the arm; the parking area is $\mathbf{RL}_{\mathbf{EC}}$ -located at the stadium; the elevator is $\mathbf{RL}_{\mathbf{IPP}}$ -located at the main building. In the presence of suitable bottom mirroring conditions, these three claims are equivalent to the pure mereotopological claims obtained by replacing the relation \mathbf{RL}_{ψ} with ψ itself.




Figure 1.26. Object-centered (left) and observer-centered (right) frames of reference.

Cases such as (155) and (156) are different. Here the difficulty does not just lie in the fact that the target object and the reference object do not stand in any mereotopological relation to each other. The difficulty is that the spatial relationships reported by such statements—corresponding to such prepositions as ‘behind’ and ‘across’, but also ‘above’, ‘underneath’, ‘left of’, etc.—have little or nothing to do with mereotopology. There is in fact a large literature devoted to the semantics of spatial prepositions (beginning with the classic work of Herskovitz 1986) and it is fair to say that their treatment requires a degree of sophistication that goes far beyond the conceptual apparatus developed above. Among other things, there are well-known complications owing to the fact that their treatment calls for a systematic distinction between object-centered frames of reference, as in (155), and observer-centered frames, as in (156) (see Figure 1.26), whereas mereotopological relations are completely independent of subjective or perspectival considerations. Still, this is not a limitation that speaks against the employment of a primitive such as ‘L’. It is, rather, an indication that the ensuing theory, S , may have to be matched with a more sophisticated background than a mereotopological theory can afford. Thus, suppose we allow the relational predicate ‘ ψ ’ in the definitions of Section 3.3 to stand for relations that are not purely mereotopological: relations such as ‘behind’, ‘across’, etc. (in each of their multiple uses). Then the idea illustrated with reference to the brain-cavity example can in principle be applied also to (155), (156), and the like. To say that the swimming pool is located behind the house is to say that it is RL_{behind} -located at the house. To say that the bus stop is located across the tree is to say that it is RL_{across} -located at the tree. And so on. Exactly how these relations should be formalized, i.e., what principles should be posited to fix the logical behavior of the relevant ψ , is where things get difficult and may require a lot of detailed work. (We have, after all, seen how difficult it is to do this when ψ is a mereotopological relation.) But that is not to say that relative locations require a *sui generis* treatment that a suitable extension of S could not accommodate.

On the other hand, here is where the main metaphysical assumption underlying S may be questioned. The accommodation comes with definition (158), which allows one to handle relative locations in terms of a primitive relation, L , whose range consists exclusively of spatial regions. In other words, it allows one to express a spatial relation between a target and a reference object as a relation between the target and the reference's *place*. But one might object that this has things the wrong way round. Relative locations are ontologically neutral with respect to the status of space, whereas the proposed treatment depends crucially on the assumption that places, and spatial regions generally, are entities of a kind. A relationist about space might therefore reject the account and take the opportunity to reverse the order of the explanation: to be RL_ψ -related to a place is to be ψ -related to the place's tenant, and talk about places is shorthand for talk about spatial objects. Also a substantivalist about space might think that when it comes to spatial reasoning, objects are conceptually prior to their locations, since we cannot identify the latter independently of the former (Strawson 1959, ch. 1.) Even a forerunning substantivalist such as Newton emphasized that absolute places, defined as in (157), are scarcely useful for locating things in the world: we do not locate an object on a moving ship with reference to an immobile environment but, rather, with reference to the ship itself (*Principia*, Definitions, Sch. 4). In short, there may be philosophical as well as methodological reasons for resisting the treatment of locative relations indicated above, and if these reasons are taken seriously, then cases such as (154)–(156) run afoul of the basic framework outlined here and call for independent treatment. An articulated proposal in this spirit may be found in Donnelly (2005), but much recent literature devoted to the formal representation of direction and other qualitative spatial relations (from the works of Mukerjee and Joe 1990, Frank 1992, and Hernández 1994 to more interdisciplinary works such as those collected in van der Zee and Slack 2003) may be viewed in this perspective.

3.5 Location, connection, and parthood

Let us conclude with a few remarks about the whole conceptual package that we have been putting together. We have at this point three main primitive notions: location, connection, and parthood. Some structural relationships among these notions have been examined, but the general question of their ontological intertwining is still open. Generally speaking, parthood and connection are independent from each other, unless one accepts the converse monotonicity principle (C.4) (Section 2.3.2). But what about location? Let us keep with the assumption that location

is a relation between a thing and its place. Is that relation completely autonomous or does it entail a mereotopological linkage of some sort?

A purely mereological linkage seems out of the question. There is no reason to think that I share any parts with the space I occupy, just as there is no reason to think that movement—change of location—is a form of mereological change. This is not to say that location implies disjointness, since the first argument of L may itself be a region, or a hybrid fusion including regions among other things. In general, however, it seems perfectly reasonable to assume that the implication holds for *ordinary* cases of location—a thesis that can be put as follows:

$$(L.18) \quad \text{PL}xy \rightarrow \exists z(\text{P}zx \wedge \text{L}zy \wedge \text{D}zy) \quad \text{Spatial Disjointness}$$

What about a topological linkage? In this case the picture seems different. If an object is located at some place, then it might be plausible to suppose that the object and its place are connected in some way:

$$(L.19) \quad \text{L}xy \rightarrow \text{C}xy \quad \text{Spatial Connection}$$

In what way would I be connected to my place? Since overlap is excluded, the relevant linkage would have to be one of external connection. However, this means that the plausibility of (L.19) depends crucially on the interpretation of ‘C’.

If we go along with the standard interpretation, corresponding to a boundary-based theory such as *KGEMT*, the prospects are slim. On that interpretation, two things can be externally connected only if one is open and the other closed, at least in the relevant contact area: two closed entities, or two open entities, can only connect through mereological overlap (Section 2.4.1). Now, suppose I am a closed body. Then (L.19), together with (L.18), imply that the region at which I am located—my place—must be open. That region, however, has a closure, and one should think that the closure of a spatial region is itself a region: its boundary is pure space, too. But my boundary is not pure space: whatever it is, it is part of me, and none of me is made of space. Thus, the closure of my place and I are disjoint—which is to say that we are not connected after all. A different way of putting this involves the thought that my place and I must have the same topology: if I am closed (for instance), then my place must be closed too, hence we cannot be externally connected. Strictly speaking, this involves an appeal to the principle of bottom mirroring, specifically the following instance of corollary (142):

$$(159) \quad \text{L}xy \wedge \text{L}zw \wedge \text{IPP}xz \rightarrow \text{IPP}yw$$

My interior (x) is a proper interior part of me (z), hence its place (y) must be an interior proper part of my place (w). We have seen that bottom mirroring may not generally hold, but this specific instance seems

fair. Yet (L.19) would require the opposite: it would require my place to be an interior proper part of the place of my interior—absurd.

By contrast, suppose we go along with a different interpretation. For instance, consider those interpretations that explain connection explicitly in terms of spatial location. We have seen three such interpretations, corresponding to the following necessary and sufficient conditions for two entities x and y to be connected: (i) the place of x is connected to the place of y (Section 2.3.2, thesis (74)); (ii) the place of x overlaps the place of y (Section 2.4.2, thesis (89)); or (iii) the closure of the place of x overlaps the closure of the place of y (Section 2.4.3, thesis (90)). With the help of ‘L’, and assuming the spatiality axiom (L.3), these three interpretations can now be formally stated as follows:

$$(160) \quad Cxy \leftrightarrow C(px)(py)$$

$$(161) \quad Cxy \leftrightarrow O(px)(py)$$

$$(162) \quad Cxy \leftrightarrow O(cpx)(cpy)$$

And, plainly, each of these statements is perfectly compatible with the spatial connection principle (L.19). In fact, since C and O are both reflexive, each statement *entails* (L.19) in view of the following S -theorem:

$$(163) \quad Lxy \rightarrow px = y = py$$

This is not surprising. After all, these theories establish an intimate relationship between topology and spatiality, and the claim that location implies connection is but one way of making that explicit. One might even go as far as to say that on these theories location *is* connection of a kind: it is the relation of connection that always holds in the special case where one of the relata equals the place of the other. On the other hand, it is fair to note that all of this depends on the assumption that L is conditionally reflexive (L.4). Should one decide to go for the alternative option (L.5), treating L as a relation whose domain does not contain any spatial regions, (163) would not hold and (L.19) would not follow from (160)–(162). Indeed, on that way of interpreting ‘L’, the picture would be perfectly reversed: nothing would be connected to its place because places would lack a place of their own, hence (L.19) would be just as unacceptable in these theories as it is in *KGEMT*.

These are just some examples, but they suffice to show that the ontological intertwining between locative and mereotopological concepts cannot be assessed without a clear stand on some very basic semantic issues concerning those concepts. Of course, this is also true of the interplay between mereological and topological concepts, but in the present case the stakes are higher. Some of the options leave room for no intertwining whatsoever; other options trivialize the intertwining by explaining

one sort of concept directly in terms of the other. There is, to be sure, some room for compromise. After all, *KGEMT* is a pretty strong theory and its weaker children are compatible with (L.19). Similarly, there is strictly speaking no constraint to supplement its main alternatives with the explicit equivalences in (160)–(162), so the status of (L.19) is in principle open for those theories, too. One final remark, however, is in order. For suppose we do accept (L.19) together with (L.8), that is, suppose we do establish an overt mereotopological linkage of external connection between a spatial object and its spatial location. Then it follows that no spatial entity will ever qualify as a interior proper part of anything, except for empty regions of space:

$$(164) \quad \text{PL}xy \rightarrow \neg\text{IPP}xz \wedge \neg\text{IPP}yz$$

This is an immediate consequence of the definition of ‘IPP’ (Section 2.3.1): my “interior” proper parts, for example, would immediately turn into tangential proper parts by virtue of being connected to something with which I have no parts in common—their places; and the “interior” proper parts of my place, at least those proper parts that are not empty, would immediately turn into tangential proper parts by virtue of being connected to something with which they have nothing in common—their material guests. This is bad news, for it means that our mereotopology would collapse altogether. Or rather, it means that it would have to be largely re-written by replacing throughout our topological primitive ‘C’ with the following impure connection predicate:

$$(165) \quad \text{RC}xy =_{df} Cxy \wedge (Rx \leftrightarrow Ry) \quad \text{Restricted Connection}$$

It is, of course, this notion of connection that we had in mind in setting up *KGEMT* and its variants. RC never cuts across levels: in order for two entities to be RC-related, both of them or neither of them must be regions of space. But then we have come around the circle, for the RC-variant of (L.19) is clearly false.

We face, here, the fundamental limit of topology—and mereotopology—as a general theory of space. One way or the other, the *analysis situs* cannot do proper justice to the fact that objects are *situated*, which is why the theory of location is independently needed. One way or the other, spatial reasoning must come to terms with the fundamental metaphysical mystery on which it depends—embedding in space.

Acknowledgments

Parts of this chapter draw on previous material. In particular, Section 1 has some overlap with Varzi (2003) while Section 2 has some overlap with Chapters 1, 4 and 5 of Casati and Varzi (1999).

References

- Aczel P., 1988, *Non-Well-Founded Sets*, Stanford: CSLI Publications.
- Aiello M., 2000, 'Topo-distance: Measuring the Difference between Spatial Patterns', *Proc. JELIA 00*, Berlin: Springer, pp. 73–86.
- Armstrong D.M., 1989, *Universals*, Boulder CO: Westview Press.
- Asher N. and Vieu L., 1995, 'Toward a Geometry of Common Sense: A Semantics and a Complete Axiomatization of Mereotopology', *Proc. IJCAI 95*, San Mateo: Morgan Kaufmann, pp. 846–52.
- Baker L. R., 1997, 'Why Constitution Is Not Identity', *J. of Philosophy* 94: 599–621.
- Baumgartner W. and Simons P.M., 1994, 'Brentano's Mereology', *Axiomathes* 5: 55–76.
- Baxter D., 1988, 'Identity in the Loose and Popular Sense', *Mind* 97: 575–82.
- Bennett B., 1996a, 'Modal Logics for Qualitative Spatial Reasoning', *Bulletin of the Interest Group in Pure and Applied Logic* 4: 23–45.
- Bennett B., 1996b, 'Carving Up Space: Steps Towards Construction of an Absolutely Complete Theory of Spatial Regions', *Proc. JELIA 96*, Berlin: Springer, pp. 337–53.
- Biacino L. and Gerla G., 1991, 'Connection Structures', *Notre Dame J. of Formal Logic* 32: 242–7.
- Bigelow J., 1996, 'God and the New Math', *Philosophical Studies* 84: 127–54.
- Bolzano, B., 1851, *Paradoxien des Unendlichen*, Leipzig: Reclam.
- Borgo S., Guarino N., Masolo C., 1996, 'A Pointless Theory of Space Based on Strong Connection and Congruence', *Proc. KR 96*, San Francisco: Morgan Kaufmann, pp 220–9.
- Bostock D., 1979, *Logic and Arithmetic, Vol. 2*, Oxford: Clarendon.
- Brentano F., 1906, 'Nativistische, empiristische und anoetistische Theorie unserer Raumvorstellung', pub. post. in his 1976, pp. 164–77.
- Brentano F., 1933, *Kategorienlehre* (ed. A. Kastil), Hamburg: Meiner.
- Brentano F., 1976, *Philosophische Untersuchungen zu Raum, Zeit und Kontinuum* (ed. S. Körner and R.M. Chisholm), Hamburg: Meiner.
- Bunge M., 1966, 'On Null Individuals', *J. of Philosophy* 63: 776–8.
- Bunt H.C., 1985, *Mass Terms and Model-Theoretic Semantics*, Cambridge: Cambridge U. Press.
- Burkhardt H. and Dufour C.A., 1991, 'Part/Whole I: History', in H. Burkhardt and B. Smith (eds.), *Handbook of Metaphysics and Ontology*, Munich: Philosophia, pp. 663–73.
- Cartwright R., 1975, 'Scattered Objects', in K. Lehrer (ed.), *Analysis and Metaphysics*, Dordrecht: Reidel, pp. 153–71.

- Casati R. and Varzi A.C., 1994, *Holes and Other Superficialities*, Cambridge MA: MIT Press.
- Casati R. and Varzi A.C., 1999, *Parts and Places*, Cambridge MA, MIT Press.
- Chisholm R.M., 1978, 'Brentano's Conception of Substance and Accident', *Grazer Philosophische Studien* 5: 197–210.
- Chisholm R.M., 1984, 'Boundaries as Dependent Particulars', *Grazer philosophische Studien* 10: 87–95.
- Chisholm R.M., 1987, 'Scattered Objects', in J.J. Thomson (ed.), *On Being and Saying: Essays for Richard Cartwright*, Cambridge MA: MIT Press, pp. 167–73.
- Chisholm R.M., 1993, 'Spatial Continuity and the Theory of Part and Whole. A Brentano Study', *Brentano Studien* 4: 11–23.
- Clarke B.L., 1981, 'A Calculus of Individuals Based on "Connection"', *Notre Dame J. of Formal Logic* 22: 204–18.
- Clarke B.L., 1985, 'Individuals and Points', *Notre Dame J. of Formal Logic* 26: 61–75.
- Cohn A.G., Bennett B., Gooday J.M., Gotts N., 1997, 'Qualitative Spatial Representation and Reasoning with the Region Connection Calculus', *Geoinformatica* 1: 1–44.
- Cohn A.G. and Varzi A.C., 2003, 'Mereotopological Connection', *J. of Philosophical Logic* 32: 357–90.
- Copeland B.J., 1995, 'On Vague Objects, Fuzzy Logic and Fractal Boundaries', *Southern J. of Philosophy* 33 (Suppl.): 83–96.
- Correia F., 2005, *Existential Dependence and Cognate Notions*, Munich: Philosophia.
- Cruse D.A., 1979, 'On the Transitivity of the Part-Whole Relation', *J. of Linguistics* 15: 29–38.
- Davidson D., 1969, 'The Individuation of Events', in N. Rescher (ed.), *Essays in Honor of Carl G. Hempel*, Dordrecht: Reidel, 216–34.
- De Laguna T., 1922, 'Point, Line, and Surface as Sets of Solids', *J. of Philosophy* 19: 449–61.
- Doepke F.C., 1982, 'Spatially Coinciding Objects', *Ratio* 24: 45–60.
- Donnelly M., 2004, 'A Formal Theory for Reasoning about Parthood, Connection, and Location', *Artificial Intelligence* 160: 145–72.
- Donnelly M., 2005, 'Relative Places', *Applied Ontology* 1: 55–75.
- Donnelly M. and Smith B., 2003, 'Layers: A New Approach to Locating Objects in Space', *Proc. COSIT 03*, Berlin: Springer, pp. 46–60.
- Düntsch I., Wang H., and McCloskey S., 2001, 'A Relation-Algebraic Approach to the Region Connection Calculus', *Theoretical Computer Science* 255: 63–83.
- Eberle R.A., 1968, 'Yoes on Non-Atomic Systems', *Noûs* 2: 399–403.

- Eberle R.A., 1970, *Nominalistic Systems*, Dordrecht: Reidel.
- Engel R. and Yoes M.G., 1996, 'Exponentiating Entities by Necessity', *Australasian J. of Philosophy* 74: 293–304.
- Fine K., 1975, 'Vagueness, Truth and Logic', *Synthese* 30: 265–300.
- Fine K., 1982, 'Acts, Events, and Things', *Proc. 6th Wittgenstein Symposium*, Vienna: HPT, pp. 97–105.
- Forrest P., 1996, 'From Ontology to Topology in the Theory of Regions', *The Monist* 79: 34–50.
- Frank A.U., 1992, 'Qualitative Spatial Reasoning about Distances and Directions in Geographic Space', *J. of Visual Languages and Computing* 3: 343–71.
- Gallois A., 1998, *Occasions of Identity*, Oxford: Clarendon.
- Galton A., 1996, 'Taking Dimension Seriously in Qualitative Spatial Reasoning', *Proc. ECAI 96*, Chichester: Wiley, pp. 501–5.
- Galton A., 1999, 'The Mereotopology of Discrete Space', *Proc. COSIT 99*, Berlin: Springer, pp. 251–6.
- Galton A., 2000, *Qualitative Spatial Change*, Oxford: Oxford U. Press.
- Galton A., 2003, 'On the Ontological Status of Geographical Boundaries', in M. Duckham *et al.* (eds.), *Foundations of Geographic Information Science*, London: Taylor and Francis, pp. 151–71.
- Galton A., 2004, 'Multidimensional Mereotopology', *Proc. KR 04*, Menlo Park: AAAI Press, pp. 45–54.
- Geach, P.T., 1949, 'On Rigour in Semantics', *Mind*, 58: 518–22.
- Gerstl P. and Pribbenow S., 1995, 'Midwinters, End Games, and Body-parts: A Classification of Part-Whole Relations', *International J. of Human-Computer Studies* 43: 865–89.
- Giritli M., 2003, 'Who Can Connect in RCC?', *Proc. KI 03*, Berlin: Springer, pp. 565–79.
- Goodman N., 1951, *The Structure of Appearance*, Cambridge MA: Harvard U. Press (1977³, Dordrecht: Reidel).
- Goodman N., 1956, 'A World of Individuals', in J. Bochenski *et al.*, *The Problem of Universals*, Notre Dame: Notre Dame U. Press, pp. 13–31.
- Goodman N., 1958, 'On Relations that Generate', *Philosophical Studies* 9: 65–6.
- Gotts N., 1994a, 'How Far Can We 'C'? Defining a 'Doughnut' Using Connection Alone', *Proc. KR 94*, San Mateo: Morgan Kaufmann, pp. 246–57.
- Gotts N., 1994b, 'Defining a 'Doughnut' Made Difficult', in C. Eschenbach *et al.* (eds.), *Topological Foundations of Cognitive Science*, U. of Hamburg: Cognitive Science Program, Rep. No. 37, pp. 105–29.
- Gotts N., 1996a, 'Topology from a Single Primitive Relation', U. of Leeds: School of Computer Studies, Rep. 96.23.

- Gotts N., 1996b, 'An Axiomatic Approach to Topology for Spatial Information Systems', U. of Leeds: School of Computer Studies, Rep. 96.25.
- Gotts N., Gooday J.M., Cohn A.G., 1996, 'A Connection-Based Approach to Commonsense Topological Description and Reasoning', *The Monist* 79: 51–75.
- Heller M., 1984, 'Temporal Parts of Four Dimensional Objects', *Philosophical Studies* 46: 323–34.
- Hempel C.G., 1953, 'Reflections on Nelson Goodman's "The Structure of Appearance"', *Philosophical Review* 62: 108–16.
- Henry D., 1991, *Medieval Mereology*, Amsterdam: Grüner.
- Hernández D., 1994, *Qualitative Representation of Spatial Knowledge*, Berlin: Springer.
- Herskovits A., 1986, *Language and Spatial Cognition*, Cambridge: Cambridge U. Press.
- Holden T., 2004, *The Architecture of Matter*, Oxford: Clarendon.
- Hudson H., 2001, *A Materialist Metaphysics of the Human Person*, Ithaca NY: Cornell U. Press.
- Hudson H., 2002, 'The Liberal View of Receptacles', *Australasian J. of Philosophy* 80: 432–9.
- Hudson H. (ed.), 2004, *Simples*, special issue of *The Monist* 87: 303–451
- Husserl E., 1900/1901, *Logische Untersuchungen. Zweiter Band*, Halle: Niemeyer (1913²).
- Iris M.A., Litowitz B.E., Evens M., 1988, 'Problems of the Part-Whole Relation', in M. Evens (ed.), *Relations Models of the Lexicon*, Cambridge: Cambridge U. Press, pp. 261–88.
- Johansson I., 2004, 'On the Transitivity of Parthood Relations', in H. Hochberg and K. Mulligan (eds.), *Relations and Predicates*, Frankfurt: Ontos, pp. 161–81.
- Johnston M., 1992, 'Constitution Is Not Identity', *Mind* 101: 89–105.
- Keefe R., 2000, *Theories of Vagueness*, Cambridge: Cambridge U. Press.
- Kline A.D. and Matheson C.A., 1987, 'The Logical Impossibility of Collision', *Philosophy* 62: 509–15.
- Kuratowski C., 1922, 'Sur l'opération A^- de l'Analysis Situs', *Fundamenta Mathematicae* 3: 182–99.
- Leonard H.S. and Goodman N., 1940, 'The Calculus of Individuals and Its Uses', *J. of Symbolic Logic* 5: 45–55.
- Leśniewski S., 1916, *Podstawy ogólnej teorii mnogości. I*, Moskow: Prace Polskiego Koła Naukowego w Moskwie.
- Lewis D.K., 1986, *On the Plurality of Worlds*, Oxford: Blackwell.
- Lewis D.K., 1991, *Parts of Classes*, Oxford: Blackwell.
- Li S. and Ying M., 2004, 'Generalized Region Connection Calculus', *Artificial Intelligence* 160: 1–34.

- Lowe E.J., 1989, *Kinds of Being*, Oxford: Blackwell.
- Lowe V., 1953, 'Professor Goodman's Concept of an Individual', *Philosophical Review* 62: 117–26.
- Markosian N., 1998a, 'Simples', *Australasian J. of Philosophy* 76: 213–28.
- Markosian N., 1998b, 'Brutal Composition', *Philosophical Studies* 92: 211–49.
- Martin R.M., 1965, 'Of Time and the Null Individual', *J. of Philosophy* 62: 723–36.
- Masolo C. and Vieu L., 1999, 'Atomicity vs. Infinite Divisibility of Space', *Proc. COSIT 99*, Berlin: Springer, 1999, pp. 235–50.
- McGee V., 1997, "'Kilimanjaro'", *Canadian J. of Philosophy* 23 (Suppl.): 141–95.
- McGrath M., 1998, 'Van Inwagen's Critique of Universalism', *Analysis* 58: 116–121.
- Merricks T., 1999, 'Composition as Identity, Mereological Essentialism, Counterpart Theory', *Australasian J. of Philosophy* 77: 192–5.
- Moltmann F., 1997, *Parts and Wholes in Semantics*, Oxford: Oxford U. Press.
- Morison B., 2002, *On Location: Aristotle's Concept of Place*, Oxford: Clarendon.
- Mukerjee A. and Joe G., 1990, 'A Qualitative Model for Space', *Proc. AAAI 90*, Menlo Park: AAAI Press, pp. 721–7.
- Needham P., 1981, 'Temporal Intervals and Temporal Order', *Logique et Analyse* 24: 49–64.
- Noonan H.W., 1993, 'Constitution Is Identity', *Mind* 102: 133–46.
- Nutt W., 1999, 'On the Translation of Qualitative Spatial Reasoning into Modal Logics', *Proc. KI 99*, Berlin: Springer, pp. 113–24.
- Parsons J., 2000, 'Must a 4-Dimensionalist Believe in Temporal Parts?', *The Monist* 83: 399–418.
- Parsons J., 2004, 'Distributional Properties', in F. Jackson and G. Priest (ed.) *Lewisian Themes*, Oxford: Clarendon, pp. 173–80.
- Parsons J., 2006, 'Theories of Location', *Oxford Studies in Metaphysics* 3: 201–232.
- Parsons T. and Woodruff P., 1995, 'Worldly Indeterminacy of Identity', *Proc. Aristotelian Society* 95: 171–91.
- Peirce C.S., 1893, 'The Logic of Quantity', pub. post. in his *Collected Papers*, Vol. IV (ed. C. Hartshorne and P. Weiss), Cambridge MA: Harvard U. Press, 1933, pp. 85–152.
- Perzanowski J., 1993, 'Locative Ontology', *Logic & Logical Philosophy* 1: 7–94.
- Pianesi F. and Varzi A.C., 1996a, 'Events, Topology, and Temporal Relations', *The Monist* 78: 89–116.

- Pianesi F. and Varzi A.C., 1996b, 'Refining Temporal Reference in Event Structures', *Notre Dame J. of Formal Logic* 37: 71–83.
- Polkowski L. and Skowron A., 1994, 'Rough Mereology', *Proc. ISMIS 94*, Berlin: Springer, pp. 85–94.
- Pratt I. and Schoop D., 2000, 'Expressivity in Polygonal, Plane Mereotopology', *J. of Symbolic Logic* 65: 822–38.
- Pyle A., 1995, *Atomism and Its Critics*, Bristol: Thoemmes.
- Randell D.A., Cui Z., Cohn A.G., 1992, 'A Spatial Logic Based on Regions and Connections', *Proc. KR 92*, Los Altos CA: Morgan Kaufmann, pp. 165–76.
- Rea M. (ed.), 1997, *Material Constitution*, Lanham MD: Rowman & Littlefield.
- Rea M., 1998, 'In Defense of Mereological Universalism', *Philosophy and Phenomenological Research* 58: 347–360.
- Renz J., 1998, 'A Canonical Model of the Region Connection Calculus', *Proc. KR 98*, San Francisco: Morgan Kaufmann, pp. 330–41.
- Rescher N., 1955, 'Axioms for the Part Relation', *Philosophical Studies* 6: 8–11.
- Ridder L., 2002, *Mereologie*, Frankfurt: Klostermann.
- Roeper P., 1997, 'Region-Based Topology', *J. of Philosophical Logic* 26: 251–309.
- Rosen G. and Dorr C., 2002, 'Composition as a Fiction', in R. Gale (ed.), *The Blackwell Guide to Metaphysics*, Oxford: Blackwell, pp. 151–74.
- Sanford D., 1993, 'The Problem of the Many, Many Composition Questions, and Naive Mereology', *Noûs* 27: 219–28.
- Schuldenfrei R., 1969, 'Eberle on Nominalism in Non-Atomic Systems', *Noûs* 3: 427–30.
- Sharvy R., 1983, 'Mixtures', *Philosophy and Phenomenological Research* 44: 227–39.
- Shorter J.M., 1977, 'On Coinciding in Space and Time', *Philosophy* 52: 399–408.
- Sider T., 1993, 'Van Inwagen and the Possibility of Gunk', *Analysis* 53: 285–9.
- Sider T., 2001, *Four-Dimensionalism*, Oxford: Clarendon.
- Simons P.M., 1987, *Parts*, Oxford: Clarendon.
- Simons P.M., 1991a, 'Faces, Boundaries, and Thin Layers', in A.P. Martinich and M.J. White (eds.), *Certainty and Surface in Epistemology and Philosophical Method. Essays in Honor of Avrum Stroll*, Lewiston: Edwin Mellen Press, pp. 87–99.
- Simons P.M., 1991b, 'Free Part-Whole Theory', in K. Lambert (ed.), *Philosophical Applications of Free Logic*, Oxford: Oxford U. Press, pp. 285–306.

- Simons P.M., 1991c, 'Part/Whole II: Mereology Since 1900', in H. Burkhardt and B. Smith (eds.), *Handbook of Metaphysics and Ontology*, Munich: Philosophia, pp. 209–10.
- Simons P.M., 2003, 'The Universe', *Ratio* 16: 237–50.
- Simons P.M., 2004, 'Extended Simples: A Third Way between Atoms and Gunk', *The Monist* 87: 371–84.
- Smith B., 1994, *Austrian Philosophy*, La Salle IL: Open Court.
- Smith B., 1995, 'On Drawing Lines on a Map', *Proc. COSIT 95*, Berlin: Springer, pp. 475–84.
- Smith B., 1996, 'Mereotopology: A Theory of Parts and Boundaries', *Data & Knowledge Engineering* 20: 287–304.
- Smith B., 1997, 'Boundaries: An Essay in Mereotopology', in L.H. Hahn (ed.), *The Philosophy of Roderick Chisholm*, La Salle IL: Open Court, pp. 534–61.
- Smith B., 2001, 'Fiat Objects', *Topoi* 20: 131–148.
- Smith B. and Mark D., 2003, 'Do Mountains Exist? Towards an Ontology of Landforms', *Environment & Planning B*, 30: 411–27.
- Smith B. and Varzi A.C., 2000, 'Fiat and Bona Fide Boundaries', *Philosophy and Phenomenological Research* 60: 401–20.
- Sorabji R., 1988, *Matter, Space, and Motion*, Ithaca NY: Cornell U. Press.
- Stell J.G., 2000, 'Boolean Connection Algebras: A New Approach to the Region-Connection Calculus', *Artificial Intelligence* 122: 111–36.
- Stell J.G. and Worboys M.F., 1997, 'The Algebraic Structure of Sets of Regions', *Proc. COSIT 97*, Berlin: Springer, pp. 163–74.
- Strawson P.F., 1959, *Individuals*, London: Methuen.
- Stroll A., 1988, *Surfaces*, Minneapolis: U. of Minnesota Press.
- Tarski A., 1929, 'Les fondements de la géométrie des corps', *Annales de la Société Polonaise de Mathématique* 7 (suppl.): 29–33.
- Tarski A., 1935, 'Zur Grundlegung der Booleschen Algebra. I', *Fundamenta Mathematicae* 24: 177–98.
- Thomson J.J., 1983, 'Parthood and Identity Across Time', *J. of Philosophy* 80: 201–20.
- Thomson J.J., 1998, 'The Statue and the Clay', *Noûs* 32: 149–73.
- Tiles J.E., 1981, *Things That Happen*, Aberdeen: Aberdeen U. Press.
- Tsai H-C., 2005, *The Logic and Metaphysics of Part-Whole Relations*, PhD Dissertation, Columbia U.: Department of Philosophy.
- Tversky B., 1989, 'Parts, Partonomies, and Taxonomies', *Developmental Psychology* 25: 983–95.
- van Benthem J., 1983, *The Logic of Time*, Dordrecht: Reidel (1991²).
- van Cleve J., 1986, 'Mereological Essentialism, Mereological Conjunctivism, and Identity Through Time', *Midwest Studies in Philosophy* 11: 141–56.

- van der Zee E. and Slack J. (eds.), 2003, *Representing Direction in Language and Space*, Oxford: Oxford U. Press.
- van Inwagen P., 1981, 'The Doctrine of Arbitrary Undetached Parts', *Pacific Philosophical Quarterly* 62: 123–37.
- van Inwagen P., 1990, *Material Beings*, Ithaca NY: Cornell U. Press.
- van Inwagen P., 1993, 'Naive Mereology, Admissible Valuations, and Other Matters', *Noûs* 27: 229–34.
- van Inwagen P., 1994, 'Composition as Identity', *Philosophical Perspectives* 8: 207–20.
- Varzi A.C., 1994, 'On the Boundary Between Mereology and Topology', *Proc. 16th Wittgenstein Symposium*, Vienna: HPT, pp. 423–42.
- Varzi A.C., 1996a, 'Parts, Wholes, and Part-Whole Relations: The Prospects of Mereotopology', *Data & Knowledge Engineering* 20: 259–86.
- Varzi A.C., 1996b, 'Reasoning about Space: The Hole Story', *Logic & Logical Philosophy* 4: 3–39.
- Varzi A.C., 1997, 'Boundaries, Continuity, and Contact', *Noûs* 31: 26–58.
- Varzi A.C., 2000, 'Mereological Commitments', *Dialectica* 54: 283–305.
- Varzi A.C., 2001, 'Vagueness in Geography', *Philosophy & Geography* 4: 49–65.
- Varzi A.C., 2003, 'Mereology', in E.N. Zalta (ed.), *Stanford Encyclopedia of Philosophy*, Stanford: CSLI (on line publication).
- Varzi A.C., 2005, 'A Note on the Transitivity of Parthood', *Applied Ontology* 1: 141–146.
- Varzi A.C., 2006, 'The Universe among Other Things', *Ratio* 19: 107–20.
- Whitehead A.N., 1919, *An Enquiry Concerning the Principles of Human Knowledge*, Cambridge: Cambridge U. Press.
- Whitehead A.N., 1920, *The Concept of Nature*, Cambridge: Cambridge U. Press.
- Whitehead A.N., 1929, *Process and Reality*, New York: Macmillan.
- Wiggins D., 1968, 'On Being in the Same Place at the Same Time', *Philosophical Review* 77: 90–5.
- Wiggins D., 1980, *Sameness and Substance*, Oxford: Blackwell.
- Winston M., Chaffin R., Herrmann D., 1987, 'A Taxonomy of Part-Whole Relations', *Cognitive Science* 11: 417–44.
- Yi B.-U., 1999, 'Is Mereology Ontologically Innocent?', *Philosophical Studies* 93: 141–60.
- Yoes M., 1967, 'Nominalism and Non-Atomic Systems', *Noûs* 1: 193–200.
- Zadeh L., 1965, 'Fuzzy Sets', *Information and Control* 8: 338–53.
- Zimmerman D.W., 1996a, 'Could Extended Objects Be Made Out of Simple Parts?', *Philosophy and Phenomenological Research* 56: 1–29.
- Zimmerman D.W., 1996b, 'Indivisible Parts and Extended Objects', *The Monist* 79: 148–80.