

Layers: A New Approach to Locating Objects in Space

Maureen Donnelly¹ and Barry Smith^{1,2}

¹Institute for Formal Ontology and Medical Information Science, University of Leipzig
{maureen.donnelly,bsmith}@ifomis.uni-leipzig.de

²Department of Philosophy, University at Buffalo, NY

Abstract. Standard theories in mereotopology focus on relations of parthood and connection among spatial or spatio-temporal regions. Objects or processes which might be located in such regions are not normally directly treated in such theories. At best, they are simulated via appeal to distributions of attributes across the regions occupied or by functions from times to regions. The present paper offers a richer framework, in which it is possible to represent directly the relations between entities of various types at different levels, including both objects and the regions they occupy. What results is a layered mereotopology, a theory which can handle multiple layers (analogous to the layers of a lasagna) of spatially or spatiotemporally coincident but mereologically non-overlapping entities.

Keywords: Ontology, mereology, mereotopology, qualitative spatial reasoning, map layers, dynamic GIS

1. The Problem

Suppose you are a conservation biologist whose job it is to keep track of vulnerable species in a large national park. You develop a system of *map layers* representing not only the topography of the park region, the soil-types and vegetation zones, the system of roads and other man-made geographic objects, but also the most important habitats for each species. Your GISystem then needs to be able to use the results of your layer-making efforts in order to answer questions like: *Which of the four most vulnerable species live in riparian zones? Do Chiricahua leopard frogs live in the same vegetation zones as shovel-nosed snakes? How do climate changes in the elevated regions of the park affect long-term patterns of declining populations of scimitar-horned oryx?* And so forth.

Standard methods for dealing with map overlays in GIS were developed in the 1970s and have been incorporated into commercial GISystems software for more than two decades. Such systems can comfortably represent objects with crisp boundaries

2 Maureen Donnelly and Barry Smith

by describing the positions of their boundaries, and they can handle multiple objects which share the same footprint wherever the objects in question occupy the same regions of space. But such approaches are static in nature, and they have no direct facility for dealing with moving objects in space. With the increasing availability of temporally indexed geospatial information comes the ever more pressing need on the part of GISystems to solve the problems which arise when we need to deal with those complex interactions among different sorts of dynamic phenomena which arise in areas such as conservation biology, or predator tracking, or in the building of early warning systems for chemical hazards. For such purposes, however, we need a framework within which we can distinguish objects from the regions they occupy at successive points in time.

2. Region-Based Approaches

In the literature on spatial reasoning one dominant approach is in terms of region-based calculi of the sort that have been developed by Cohn and his associates. Such calculi have found favor among some GIScientists, not least because they correspond well not only to the field-based representations of geospatial phenomena used in scientific geography but also to our understanding of maps and map layers as consisting, in effect, of regions coloured differently according to the attributes which they instantiate. (Casati, Smith and Varzi 1998)

Region-based approaches are of interest to us here because they have served of late as the basis for serious attempts to create frameworks for reasoning about moving objects. The simplest such approach treats moving objects indirectly, in terms of the regions they occupy at successive intervals of time. Not only space itself is understood in terms of regions, but all the different kinds of entities located in space. We have not: *the talus snail spawns in the area of needle-spined pineapple cactus* or: *the tumamoc globeberry grows in the habitat of the Mexican long-tongued bat*, but rather, in each case: *(connected or scattered) region A is included as subregion within (connected or scattered) region B.*

How this region-based framework leads to problems becomes clear when we need to formulate a qualitative theory of motion. If we are to be able even to attempt to characterize movement, something more, for example temporal indexing, must be added to the mereotopology. But even then, regions-plus-attributes representations of organisms in their habitats must necessarily obscure what is involved when an enduring object is registered at different spatial locations in successive instants of time. For an adequate account of such registration data requires at least two independent sorts of spatial entities: one, the locations, which remain fixed, and the other, the objects, which move relative to them. Since the region-based approach admits only the first type of entity – the locations or regions – it must somehow simulate motion, for example via successive assignments of attributes to a fixed frame of locations.

That *a rufous-winged sparrow moves from one location (region A) to another (region B)* must then be cashed out, non-intuitively, as: *each member of this continuous sequence of sparrow-shaped regions, starting with region A and ending with region B, has at successive times, rufous-winged (etc.) attributes.* That is, instead

of talking about sparrows flying about in the sky, we talk rather of mappings of the form: Sparrow₁₅₂: time → regular closed subsets of \mathbf{R}^3 .

Notice that, besides failing to match our intuitions about what it is for birds to move through space and to inhabit different locations at different times, this picture also does less than adequate justice to regions themselves. For it leaves unexplained what it is for attributes to be correctly assigned to given region at given times. Intuitively, of course, this turns on the fact that there are corresponding *objects* which occupy those regions at the times in question. Appeal to this separate layer of objects is however precisely what is ruled out by the region-based approach, which, for reasons of mathematico-logical simplicity, attempts to explain what is cognitively more salient (objects) in terms of what is cognitively less salient (regions).

One side-effect of such reductionist treatments of physical objects and of their independent spatial properties – as well as of other spatial entities such as epidemics, wildfires, hurricanes and the like – is that we are unable to distinguish cases of true mereological overlap (i.e. the sharing of parts) from mere spatial co-location. We cannot, for example, distinguish the relation of a fish to the lake it inhabits (but is not a part of) from the relation of a genuine part of a lake (a bay, an inlet) to the lake as a whole. Both are represented in the regions-plus-attributes picture as the inclusion of a smaller in a larger region. Similarly, we cannot distinguish between genuine parts of the human body such as the heart or lungs, and foreign occupants such as parasites or shrapnel. This weakness in the region-based account is clear in the representation of phagocytosis and exocytosis offered in (Cui *et al.*, 1992), which offers no means of distinguishing between the relation of an amoeba to a portion of food which it has recently ingested and the relation of the amoeba to a genuine part, such as its nucleus.

A slightly different type of mereotopology-based analysis of motion is given in (Cohn and Hazarika, 2001). Here, the domain of the mereotopology is again restricted to regions, but this time to regions of a four-dimensional, spatial-temporal sort. This sort of account incurs the same sorts of problems as the three-dimensional region-based approach. Again, the relation between regions and the material entities (here: four-dimensional *processes*) that are supposed to inhabit them is left unexplained. Also, as in the three-dimensional region-based approach, this theory does not allow us to distinguish true overlap from spatio-temporal coincidence. Moreover, this approach is marked by an additional counter-intuitive feature in that we are forced to reduce all spatial objects to four-dimensional processes. Individuals, such as you and me, are thereby identified with their *lives* or with some totality of *histories* in which they are involved.

Bennett (2001) introduces a more sophisticated analysis, in which individuals are interpreted as distributions of matter with which count nouns are associated in a way which yields an account of motion of individuals in terms of the continuous changes in distributions of matter. Bennett thus distinguishes individuals from their spatial extensions; but he still defines individuals in terms of phenomena – distributions of matter – which are cognitively less salient than individuals themselves. Moreover, he does not take the further step of developing a framework in which entities on a plurality of levels might be recognized as coinciding spatially.

The vector methods commonly used in GIS for the manipulation of data about objects solve some of the problems at issue. From the vector point of view objects are identified with sets of points referenced to a spatial coordinate system. Moving

4 Maureen Donnelly and Barry Smith

objects are identified with sequences of sets of points indexed by times. Unfortunately however the vector approach involves a highly unrealistic understanding of what objects are. Moreover, embracing this approach means giving up many of the benefits for our understanding of spatial reasoning and of the structures of the spatial continuum which have accrued in recent years as a result of applications of mereotopology. In what follows, therefore, we seek a new sort of framework, building on mereotopology by adding a new conception of objects and their locations in space that is encapsulated in the notion of layer.

3. Mereotopology

The fruitfulness of mereotopology rests precisely on the fact that it can admit *extended* individuals such as regions, material objects, chunks of stuff, or spatio-temporally extended processes. In this way it can yield a more direct and realistic representation of the qualitative space of common sense than is available under standard reconstructions of the spatial continuum in terms of sets of points or vectors.

A mereotopology is a formal theory of parthood and connection relations. Several different mereotopologies have been proposed, including not only those of Cohn *et al.*, but also those of Asher and Vieu, Smith, and others. These theories are, it is clear, intended to be used for reasoning about spatial relations among material objects. When they are examined more closely, however, it becomes evident that – in keeping with what was said above – they assume that their immediate domains of application will be restricted to regions. The axioms formulated in (Smith 1996) are, it is true, neutral as between objects and regions; but even there no resources are provided for giving an account of the distinction between objects and the regions in which they are located. Because distinct *location* relations are not introduced into these mereotopologies, coincidence of spatial location collapses onto overlap. Where, as in (Cohn 2001), material objects are explicitly introduced, the mereotopological relations are still restricted to associated regions. Each object's spatial properties are determined by those of the region at which it is at any given time located.

Terminological clarity is important here. Note, first, that we are using 'overlap' to mean what some might prefer to call 'partial overlap'. That is, two entities will be said to overlap (share parts in common) even when they are identical. We shall say that two entities *coincide* when they occupy overlapping regions of space. Coincidence will then in general fall far short of overlap. We shall say that objects *coincide completely* with the spatial regions which, at any given instant of time, they occupy. The relation of coincidence holds not only between objects and their regions; it holds also between objects themselves. This first of all in a trivial sense: my hand and my arm (partially) coincide, and so also do the British Commonwealth and the European Union. In fact we have to deal in cases such as this with objects which do not merely coincide but also overlap (i.e. have parts in common). My hand is a part of both itself and my arm. Cyprus and Malta are both parts of the British Commonwealth and (will shortly be) parts of the European Union. If all entities (and thus all parts of entities) are spatial, then any two mereologically overlapping entities are, trivially, coincident in our sense. That is: their locations are identical at their intersections.

The relation of coincidence is however strictly broader than that of overlap. For there are pairs of coincident objects that do not share parts. The food that I am currently digesting coincides with, but does not overlap, my stomach cavity. The brain is located in the cranial cavity; but it is not a part of the cranial cavity. (Schulz and Hahn, 2001) The Great Plague of 1664 coincides with, but does not overlap, Holland.

The goal of this paper is to sketch a mereotopological framework that has the resources to deal with all types of coincident but non-overlapping entities, including not only material objects and their regions but also other types of entities such as *qualities*, *processes* and *holes*.

Qualities may coincide with material objects in the way in which, for example, the individual redness of a cube of coloured glass coincides with, but does not overlap, the glass itself. Qualities may also coincide with each other: the qualities of temperature and pressure of a given mass of air will coincide in this sense, and both will coincide with the mass of air itself.

Processes, too, may coincide in similar ways with the spatiotemporal regions they occupy. A process of deforestation may coincide with, but does not overlap, a specific geopolitical region over a given time. Processes may also coincide with each other, as when a process of absorption of a drug in a patient's body coincides with, but does not share parts with, the disease processes which the drug is designed to alleviate.

Holes may coincide with material objects in the way in which, for example, the chamber of a revolver coincides completely with, but does not overlap, the bullet which fills it (Casati and Varzi, 1994). If the revolver is moved around inside a moving railway carriage, then we can distinguish two levels of holes which are, in each temporal instant, coincident with each other but yet moving relative to each other and also relative to the territory through which the train is moving. The theory here presented allows reasoning about such multi-layered structures of coincident entities as they arise in domains such as mechanics (valves, pathways formed by piping) and medicine (body cavities and orifices). Here, however, we are concerned with layered structures in the domain of the geographic sciences.

4. Examples of Layers

Example 1

Suppose that you wish to represent the relations holding between a lake, the water in the lake, the fish swimming in the lake, and the mercury in the fish tissue. You might distinguish here four coincident three-dimensional layers:

- L1. a *region* layer, consisting of a regular spatial volume including in its interior the spatial region occupied by the lake,
- L2. a *lake* layer, consisting of a certain concave portion of the earth's surface together with a body of water,
- L3. a *fish* layer, consisting of a certain aggregate of fish,
- L4. a *mercury* (or *chemical contaminant*) layer, consisting of tiny deposits of organic mercury scattered through the lake and through the tissue of the fish.

6 Maureen Donnelly and Barry Smith

L1 serves as underlying reference-system for the edifice of layers taken as a whole. The mapping which takes every object to its spatial region is a projection of the whole domain onto the region layer L1.

Objects from the separate layers never overlap. Clearly, however, the corresponding *regions* may overlap or be parts of one another. For example, the spatial region occupied by the mercury in the fish is properly included within the spatial region occupied by the fish, which is in turn properly included within the spatial region occupied by the whole lake.

Summation may occur within layers. The mereological sum of any aggregate of fish in the lake is itself an aggregate of fish in the lake and thus included in layer L3. But there are no mixed sums, i.e. entities which have members of different layers as their parts. Thus there is no mereological sum of the fish, the mercury, and the lake. In tackling the relations between these different entities we must rather find ways to do justice to the fact that they exist on different levels.

Example 2

Suppose that you wish to represent the relations holding between topographical features, weather phenomena such as storms and winds, and the interactions of these with each other and with wind-borne nuclear, chemical or biological agents which have been released into the atmosphere. Here again you might distinguish four layers, as follows:

- L1. a region layer, consisting of a certain collection of spatial regions defined in relation to the relevant part of the surface of the earth,
- L2. a topographical layer, consisting of mountains, valleys, deserts, gullies, and other geographic features,
- L3. a storm-system and all its parts occupying sub-regions of this layer,
- L4: an airborne cloud of smallpox virus particles.

Here again the different layers are to be understood as mereologically exclusive. That is, while there can be (physical, chemical, biological) relations between the entities on neighboring layers, we stipulate that such entities never stand to each other in relations of part to whole.

Example 3

Suppose that you wish to represent the political systems of the United States, including the political systems of the constituent towns, counties and states, as well as the overarching system of the federal government, including also the governments of the other US territories (Guam, Puerto Rico, Palua, etc.). You might distinguish here five layers:

- L1. a *region* layer, defined in the obvious way,
- L2. a *local government* layer, consisting of collections of administrative geospatial units dealing with water, sewers, sanitation services, schools, etc.,
- L3. a *county government* layer, consisting of collections of administrative geospatial units dealing with property assessment and record keeping, administration of elections, consumer protection, planning and zoning, etc.,

L4. a *state government* layer, consisting of collections of administrative geospatial units dealing with judicial functions, welfare provision, public higher education, etc.,

L5. a *federal government* layer, consisting of collections of administrative geospatial units dealing with foreign relations, defense and the other functions specified in the Constitution.

Here again L1 serves as supporting framework for the edifice as a whole, and if we examine the results of projecting onto L1 the individual units in each successive layer then we see that they are organized in a system of nested honeycomb structures, with the largest cell in the structure comprehending the totality of the relevant territory, with successively finer spatial divisions as we move down through states and counties to individual towns. Some higher layers will when taken as a whole be spatially perfectly coincident with the region layer L1. Others will fall short of perfect coincidence (there is no layer of functioning county governments, for example, in Connecticut and Rhode Island).

When we construct an administrative layered hierarchy of the given sort – a hierarchy of fiat objects (Smith 2001) – then we make it true by stipulation that the objects from the separate layers do not overlap. We declare at the same time that summation is admissible (or admissible where certain conditions are satisfied) within any given layer. Thus for example when two county government bodies in L3 collaborate, e.g. in dealing with mercury in a lake, then the resultant joint administrative unit belongs also to this layer.

Example 4

Imagine that we wish to compare distinct modes of segmentation of topographic reality in a given area, for example the distinct segmentations of landforms effected by European and non-European peoples in rural Australia. Here (following Mark and Turk 2003) we might distinguish:

L1. a region layer, consisting of the middle part of the valley of the Fortescue River in northwestern Australia and of its adjacent uplands,

L2. a layer consisting of the segmentations effected by the Yindjibarndi people, employing terms such as *marnda*, *bargu*, *burbaa*, *bantha*, *gankala*, *thalungarn* and the various combinations thereof,

L3. a layer consisting of the segmentations effected by English speakers, employing terms such as *mound*, *hill*, *mountain*, *ridge*, *saddle*, *spur*, etc.

That what English-speakers call 'Pyramid Hill' corresponds to what the Yindjibarndi people call 'Googarana' signifies, in the terms of layered mereology, that the corresponding segments in L2 and L3 are centered on the peak at 21° 10' S latitude; 117° 22' E longitude in L1.

Example 5

Examples 1–4 are synchronic; they represent snapshots of reality at given times; the relative positions of the entities on certain layers, especially in the first two examples, may however vary greatly from one moment (snapshot) to the next. Layer ontologies

can however be developed also on the basis of spatiotemporal region layers, where we talk not of change but rather of an entity's having different qualities in different temporal parts (in just the way that they might have different qualities in different spatial parts). An example might be:

- L1. a four-dimensional region layer, defined as the product of a spatial component: the territory of Iraq, and a temporal component (shared by all the other layers): the interval between March 19 and April 14 2003,
- L2. a four-dimensional topographical layer, whose spatial component is defined as the sum of all topographical features (mountains, valleys, rivers, deserts, ...) within the locus of L1,
- L3. a four-dimensional geographical artifact layer, whose spatial component is defined as the sum of all geographical-scale artifacts (cities, oilwells, mine-fields, dams, ...) within the locus of L1,
- L4. a process layer, consisting of all movements of troops, materiel, and civilians within the locus of L1.

5. General Extensional Mereology

The formal theory presented here is taken from (Donnelly 2003), which also includes a semantics for the theory in terms of structures called Layered Models. The theory is formulated in standard first-order logic. It takes its starting point from mereology, the formal theory of the binary parthood relation P , and uses the following additional mereological relations, which are defined in terms of P :

- D1. $PPxy =: Pxy \ \& \ \sim Pyx$ (x is a proper part of y)
- D2. $Oxy =: \exists z(Pzx \ \& \ Pzy)$ (x and y overlap)

For the sake of simplicity we take P to be axiomatized via the familiar axioms of General Extensional Mereology (GEM) (Simons 1987), but with two small modifications, to be addressed below. The adoption of GEM axioms means that the relation of parthood is a partial ordering: it is reflexive (each entity is a part of itself), anti-symmetric (if two entities are part of each other then they are identical), and transitive. In addition it satisfies a remainder axiom, which states that, if x is not a part of y , then there is some part, z , of x that does not overlap y . Finally GEM employs an unrestricted comprehension or summation axiom (or better: axiom schema), which is formulated as follows:

$$\exists x\phi \rightarrow \exists z\forall w (Owz \leftrightarrow \exists x (\phi \ \& \ Owx))$$

Here ϕ stands for any first-order formula in which z does not occur free. The axiom schema states that, if any member of the domain satisfies the formula ϕ , then there must be in the domain a sum of all objects satisfying ϕ . More simply put: if any entity satisfies ϕ , then there is a sum of all entities satisfying ϕ . This unrestricted comprehension principle represents for many application purposes an unacceptable simplification. Modified versions of the comprehension principle can be formulated, for

example by restricting admissible sums to those consisting of topologically connected parts, and corresponding modified versions of layered mereology can then be constructed to fit.

Axioms and theorems for the base mereology will be labeled in what follows with 'P' for 'parthood'. (Here, as throughout this paper, initial universal quantifiers are suppressed unless confusion will otherwise result).

- P1. Pxx
 P2. $Pxy \ \& \ Pyx \rightarrow x = y$
 P3. $Pxy \ \& \ Pyz \rightarrow Pxz$
 P4. $\sim Pxy \rightarrow \exists z(Pzx \ \& \ \sim Ozy)$

The first three axioms require that P is a partial ordering. P4. is the remainder principle.

It follows from P1.-4. that overlap is extensional. This means that any two members of the domain that overlap the same entities are identical:

$$P-T1. \quad x = y \leftrightarrow \forall z(Oxz \leftrightarrow Oyz)$$

Because O is extensional, for any formula ϕ in which z does not occur free, if we can assign z to a member of the domain that satisfies

$$\forall w (Owz \leftrightarrow \exists x (\phi \ \& \ Owx)),$$

then this object is unique. Thus, the comprehension axiom of GEM tells us that, if any member of the domain satisfies ϕ , then there is a *unique* sum of all ϕ -ers.

6. Adding Layers

The summation axiom schema included in Layered Mereology must be weaker than GEM's comprehension schema, since we do not allow sums to cross layers. For example, let ϕ be the formula: $x = x$, then entities on different layers would satisfy ϕ , so there can in this case be no sum of all the ϕ -ers. What we need is a restricted version of the comprehension schema that requires sums to exist only if all summands are part of the same layer. To this end we introduce a new notion, that of *underlapping*, defined as follows:

$$D5. \quad Uxy =: \exists z (Pxz \ \& \ Pyz) \quad (x \text{ and } y \text{ underlap})$$

Two entities underlap whenever there exists some whole in which both are included as parts.

Our restricted version of GEM's comprehension schema then reads:

$$P5. \quad (\exists x\phi \ \& \ \forall xy(\phi \ \& \ \phi/y \rightarrow Uxy)) \rightarrow \exists z\forall w (Owz \leftrightarrow \exists x (\phi \ \& \ Owx))$$

where ϕ/y is the formula which results from ϕ when all free instances of x are replaced by instances of y (and using variable substitution as necessary so that y is free in ϕ/y where x is free in ϕ).

10 Maureen Donnelly and Barry Smith

P5. says that if there is some object that satisfies ϕ , and any two entities that satisfy ϕ underlap, then there is a sum of all entities that satisfy ϕ .

So far, we can say only that two entities share the same layer. We would like to be able to say that a certain entity *is a layer* or *is the layer of a particular object*. To this end we first of all define the relation that holds between z and y when z is y 's layer, as follows:

$$D6. \quad Lyz =: \forall w(Owz \leftrightarrow \exists x(Uxy \ \& \ Owx)) \quad (z \text{ is } y\text{'s layer})$$

D6 tells us that z is y 's layer if and only if z is the unique sum of all entities underlapping y .

Does it hold that every entity has some layer to which it belongs? This would follow from P5., but only if we knew that any two entities that underlap y must underlap each other (i.e. that U is transitive). The transitivity of U is not, however, derivable from P1.–P5., and so one final axiom needs to be added to layered mereology:

$$P6. \quad (Uxy \ \& \ Uyz) \rightarrow Uxz \quad (\text{underlap is transitive})$$

Thus if x and y are each parts of a single whole, and if the same holds for y and z , then there is some whole of which x and z are parts. From this it follows that every object has a layer, and that every object is part of its layer. We can prove further that two entities underlap if and only if they have the same layer, and we can introduce the unary predicate, L , which distinguishes those special entities called layers:

$$D7. \quad Lz =: \exists x Lxz \quad (z \text{ is a layer})$$

It is easily proved that an entity is a layer iff it is its own layer. In other words, the layers are just those members of the domain which are the mereological sums of all entities they underlap.

7. The Region Layer

Our framework thus far has been that of General Extensional Mereology with just two modifications: (i) we have restricted the comprehension principle by allowing sums to exist only if their constituents underlap; (ii) we have added an axiom to the effect that underlap is transitive. Since both layers and the relation of underlap are defined in terms of the usual mereological predicates, no new mereological primitives are needed.

For some applications, however, we may want a theory that is weaker than layered mereology in that the separate layers need not satisfy the full comprehension schema of GEM. In other words, for some applications we may not want to require arbitrary collections of entities from the same layer to have a mereological sum. Consider a modification of *Example 2*, above, which allows on the topographical layer only contiguous features to form wholes (including that maximal contiguous whole which is the given layer in its entirety). In cases such as this we can use a subtheory of layered mereology called *weak layered mereology*, in which axiom

schmema P5 is replaced by the following axiom, which merely requires that every entity has a layer:

$$P5'. \quad \forall x \exists z Lxz \quad (\text{every entity has a layer})$$

P5' tells us that for any x , there is a mereological sum of *all* entities underlapping x . But there need not be mereological sums corresponding to smaller collections of entities on x 's layer.

In what follows we shall presuppose the full resources of layered mereology, showing how these resources can be extended to account for the special role of the region layer and also for topological relations. It is then a simple matter to formulate analogous extensions within the framework of weak layered mereology.

We move beyond the strictly mereological framework first of all by adding the facility to specify the spatial locations of entities in a way that allows us to formulate relations between entities on different layers. We introduce a unary function r which assigns each object x to the region, $r(x)$, at which x is exactly located. (Compare Casati and Varzi 1999.)

Using r , we can define a one-place predicate R that distinguishes the subdomain of regions:

$$D11. \quad Ry =: \exists x(r(x) = y) \quad (y \text{ is a region})$$

Appropriate axioms can now be appended to the axioms P1.–6. of layered mereology, as follows:

$$\begin{aligned} R1. \quad & Ry \ \& \ Rz \rightarrow Uyz && (\text{all regions underlap}) \\ R2. \quad & Ry \ \& \ Uyz \rightarrow r(z) = z && (\text{every member of the region} \\ & & && \text{layer is its own region}) \end{aligned}$$

We can then prove theorems for example to the effect that arbitrary sums of regions are themselves regions.

The next axiom relates the region-function to parthood:

$$R3. \quad Pxy \rightarrow Pr(x)r(y) \quad (\text{if } x \text{ is part of } y, \text{ then } x\text{'s region} \\ \text{is part of } y\text{'s region})$$

From this we can infer that the mereological relations among objects on a given layer are mirrored in the mereological relations among their regions. Clearly, however, mereological relations among regions are not of necessity inherited by corresponding objects located at those regions – not every region is occupied by an object, and objects which do occupy regions may lie on distinct layers. We can, however, add the axiom:

$$R4. \quad Uxy \ \& \ O(r(x), r(y)) \rightarrow Oxy$$

which states that objects on the same layer which are located at overlapping regions themselves overlap.

We can finally define the relation of complete coincidence, which holds between two entities – for example a lake and the interior of its exact container – that occupy exactly the same region of space. (One might then explore the question of whether it is possible for layered mereotopology to make sense of philosophically problematic

candidate examples of exact coincidence such as you and your body, *Hamburg Stadt* and *Hamburg Land*, or Michelangelo's *David* and a certain lump of stone.)

Given axioms R1.–4. we can prove theorems to the effect that any object can completely coincide with at most one object in any given layer and that if two objects completely coincide and are on the same layer, then they are identical.

The relation of complete coincidence is reflexive, symmetric and transitive, which means that it partitions the domain of objects into equivalence classes of co-located entities. At the same time the equivalence relation of underlap partitions the domain, orthogonally, into layers, yielding a simple structure as in Figure 1.




Figure 1.

8. Layered Mereotopology

As is shown in (Donnelly, 2003), we can extend the theory of layered mereology by adding a binary relation of connectedness in the standard way, but with the addition of an axiom which restricts connection to a single layer.

Additional axioms are needed to tie the connection relation to the region function, namely axioms to the effect that connected objects occupy connected regions and objects on the same layer with connected regions are themselves connected objects.

In *Example 3* above, New York City cannot be connected to the state of New Jersey, because New York City belongs to the city government layer, while New Jersey belongs to the state government layer. They do however stand in the weaker relation of *abutment*, which holds among objects whenever their spatial regions are connected and do not overlap.

An entity is self-connected whenever all ways of dividing the entity into two parts yield parts which are themselves connected. For some purposes it will be useful to impose on the region layer the condition that it be self-connected. The idea here is that, in contrast to what is the case regarding layers housing other types of objects (for example the fish layer in *Example 1*), the region layer is characteristically a self-connected whole that extends beyond the scattered objects occupying other layers. We can impose further conditions on the region layer for other application purposes. Thus we may want to insist that the region layer be fixed across time. In the realm of cadastral ontology we may want to demand that region layers be not only self-connected but also subject to a Euclidean metric, so that they can serve as the coordinate-systems in relation to which parcels can be defined. To do this, we would need to extend layered mereotopology to include relations strong enough to specify the geometric properties of the region layer. (See Bennett 2001a)

Allowing direct descriptions of the spatial properties of material objects yields more generally the possibility of attributing different spatial structures to material objects and to the regions at which they are located at different times. For example, we may wish to represent material objects as having only closed, regular, divisible parts, but represent spatial region as sums of points; or we may wish to represent the fact that geospatial parcels and regions may possess *fiat* boundaries while the geographical features and artifacts located in them are marked by bona fide boundaries of a physical sort (Smith 1995, Smith and Varzi 2000, Bittner 2001).

9. Time and Change

There are two mutually complementary ways of introducing the factors of time and tense into layered mereotopology. The first is associated with what we shall call a ‘SNAP’ (for: snapshot) perspective on reality (Grenon and Smith 2003, Smith and Grenon 2003). This is designed to give us the facility to represent a series of inventories of what exists on a given system of layers at successive times. The conservation biologist takes samples at regular intervals of the populations of different species in given regions of the area for which he is responsible. At the same time he samples air and water quality, records meteorological data, and so forth. The machinery of the SNAP layered ontology then facilitates synchronic reasoning about those relations between populations and topographical and other features on different layers which are of ecological significance. In addition it supports the tracking of specific individuals and groups of organisms over time. This is because the SNAP ontological perspective recognizes *continuant entities* – which means objects together with their qualities, functions, roles, and powers –which endure identically through time even while undergoing certain sorts of changes.

On the SNAP perspective not only the region function but also all mereotopological and location relations are indexed by times. This allows us to represent in a natural way the results of our samplings of the positions of moving objects or populations over time. Movement is understood as change in an object’s position within a fixed location-structure from one time to the next. Hence:

rufous-winged sparrow #152 moves from region A to region B

can be represented quite naturally as:

$$(1) \quad r(\text{Sparrow}_{152}, t_1) = A \ \& \ r(\text{Sparrow}_{152}, t_2) = B.$$

We would like, however, to be able to draw on axioms which allow us to move from statements such as (1) to statements concerning the location of the sparrow during the intervening period, for example a statement to the effect that during the interval (t_1, t_2) there occurred a continuous process of spatial dislocation of Sparrow₁₅₂. To this end we need to appeal to something like the resources of the spatiotemporal process ontology ('SPAN'). Here the region layer consists of spatiotemporal volumes, and the entities on the higher layers include the occurrent processes which unfold themselves therein. When the conservation biologist needs to represent the properly dynamic features of the territory under his charge, he will want to reason (diachronically) about *processes* on different layers – about patterns of change and development. But such entities are invisible when all we have at our disposal are the successive inventories of the SNAP ontology. Such processes belong, rather, to the four-dimensional world of SPAN. According to the temporal granularity selected by the researcher the phenomena in the higher layers may be, for example, meteorological processes occurring over the span of a single day, or processes of climatic change occurring over periods of centuries. They may be cyclical processes of seasonal change in water quantity in lakes and rivers or irreversible processes of erosion stretching over millennia – and again, one of the virtues of layered mereology is that it allows us to represent such processes as spatio-temporally coincident even where they do not overlap.

10. Conclusions

The framework presented in this paper allows mereotopological relations to apply directly to all spatial entities, including spatial regions, material objects, and holes. We saw that the framework can be extended to yield a mereotopological treatment of four-dimensional entities such as motions of sparrows as well as of the spatiotemporal regions in which such processes occur. Our framework thus allows us to do justice to the expressive possibilities encapsulated in the idea of map layers while at the same time incorporating the resources of a rich and naturalistic ontology of geospatial dynamics.

Acknowledgement: This work was supported by the Alexander von Humboldt Foundation under the auspices of its Wolfgang Paul Program, and also by the National Science Foundation under grant BCS-9975557: Geographic Categories: An Ontological Investigation. Our thanks go to David Mark for helpful comments.

Bibliography

- (Asher and Vieu, 1995) N. Asher and L. Vieu. Towards a Geometry of Commonsense: A Semantics and a Complete Axiomatization of Mereotopology. *Proceedings of IJCAI '95*, 846–852, San Mateo, CA: Morgan Kaufmann.
- (Bennett, 2001) B. Bennett. Space, Time, Matter, and Things. In (Welty and Smith), 105–116.
- (Bennett, 2001a) B. Bennett. A Categorical Axiomatization of Region-Based Geometry. *Fundamenta Informaticae* 46, 1–2, 145–158.
- (Bittner 2001) T. Bittner. The Qualitative Structure of Built Environments. *Fundamenta Informaticae*, 46, 97–128.
- (Bittner and Smith, 2002) T. Bittner and B. Smith. A Theory of Granular Partitions. *Foundations of Geographic Information Science*, M. Duckham, M. F. Goodchild and M. F. Worboys, eds., London: Taylor & Francis, 117–151
- (Casati and Varzi, 1994) R. Casati and A. C. Varzi. *Holes and Other Superficialities*. Cambridge, MA: MIT Press.
- (Casati and Varzi, 1999) R. Casati and A. C. Varzi. *Parts and Places. The Structures of Spatial Representation*, Cambridge, MA: MIT Press.
- (Casati, Smith and Varzi 1998) R. Casati, B. Smith and A. C. Varzi. Ontological Tools for Geographic Representation. In: N. Guarino (ed.), *Formal Ontology in Information Systems*, Amsterdam: IOS Press, 77–85.
- (Cohn, 2001) A. G. Cohn. Formalizing Bio-Spatial Knowledge. In (Welty and Smith) 198–209.
- (Cohn and Hazarika 2001) A. G. Cohn and S. M. Hazarika. Continuous Transitions in Mereotopology. In *Commonsense 2001: 5th Symposium on Logical Formalizations of Commonsense Reasoning* (to appear).
- (Cohn and Varzi, 1998) A. G. Cohn and A. C. Varzi. Connection Relations in Mereotopology. In H. Prade (Ed.), *Proceedings of the 13th European Conference on Artificial Intelligence (ECAI '98)*, 150–154, New York: Wiley.
- (Cui et al., 1992) Z. Cui, A. G. Cohn, D. A. Randell. Qualitative Simulation Based on a Logical Formalism of Space and Time. *Proceedings of AAAI-92*.
- (Donnelly, 2003) M. Donnelly. Layered Mereotopology, forthcoming in *IJCAI 2003*.
- (Gotts et al., 1996) N. M. Gotts, J. M. Gooday, and A. G. Cohn. A Connection Based Approach to Commonsense Topological Description and Reasoning. *The Monist*, 79: 51–75, 1996.
- (Grenon and Smith, 2003) P. Grenon and B. Smith. SNAP and SPAN. A Prolegomena to Geodynamic Ontology. IFOMIS Technical Report, University of Leipzig.
- (Mark and Turk, 2003) D. M. Mark and A. Turk. Landscape Categories in Yindjibarndi: Ontology, Environment, and Language (in this volume).
- (Schulz and Hahn, 2001) S. Schulz and U. Hahn. Mereotopological Reasoning about Parts and (W)Holes in Bio-Ontologies. In (Welty and Smith), 210–221.

- (Simons, 1987) P. M. Simons. *Parts: A Study in Ontology*. Oxford: Oxford University Press.
- (Smith, 1995) B. Smith. On Drawing Lines on a Map. In: A. U. Frank and W. Kuhn (eds.), *Spatial Information Theory. A Theoretical Basis for GIS* (Lecture Notes in Computer Science 988), Berlin/Heidelberg/New York, etc.: Springer, 475–484.
- (Smith, 1996) B. Smith. Mereotopology: A Theory of Parts and Boundaries. *Data & Knowledge Engineering*, 20: 287–303.
- (Smith and Grenon, 2003) B. Smith and P. Grenon. The Cornucopia of Formal-Ontological Relations. Forthcoming in *Dialectica*.
- (Smith and Varzi, 2000) B. Smith and A. C. Varzi. Fiat and Bona Fide Boundaries. *Philosophy and Phenomenological Research*, 60: 2, 401–420.
- (Smith, 2001) B. Smith, Fiat Objects, *Topoi*, 20: 2, 131–148.
- (Welty and Smith, 2001) *Proceedings of the 2nd International Conference on Formal Ontology in Information Systems* (FOIS 2001), New York: ACM Press.