from *Foundations of Geographic Information Science*, M. Duckham, M. F. Goodchild and M. F. Worboys (eds.), London: Taylor & Francis, 2003, 117–151

CHAPTER 1

A Theory of Granular Partitions

Thomas Bittner and Barry Smith
Institute for Formal Ontology and Medical Information Science
University of Leipzig, Leipzig, 04107, Germany

1 INTRODUCTION

Imagine that you are standing on a bridge above a highway checking off the makes and models of the cars that are passing underneath. Or that you are a postal clerk dividing envelopes into bundles; or a laboratory technician sorting samples of bacteria into species and subspecies. Or imagine that you are making a list of the fossils in your museum, or of the guests in your hotel on a certain night. In each of these cases you are employing a certain grid of labeled cells, and you are recognizing certain objects as being located in those cells. Such a grid of labeled cells is an example of what we shall call a *granular partition*. We shall argue that granular partitions are involved in all naming, listing, sorting, counting, cataloguing and mapping activities. Division into units, counting and parceling out, mapping, listing, sorting, pigeonholing, cataloguing are activities performed by human beings in their traffic with the world. Partitions are the cognitive devices designed and built by human beings to fulfill these various listing, mapping and classifying purposes.

In almost all current work in areas such as common-sense reasoning and natural language semantics it is the naïve portion of set theory that is used as basic framework. The theory of granular partitions as it is developed in this paper is intended to serve as an alternative to set theory both as a tool of formal ontology and as a framework for the representation of human cognition. Kinds, sorts, species and genera are standardly treated as sets of their instances; subkinds as subsets of these sets. Set theory nicely does justice to the granularity that is involved in our sorting and classification of reality by giving us a means of treating objects as *elements* of sets, i.e. as single whole units within which further parts are not recognized. But set theory also has its problems, not the least of which is that it supports no distinction between natural totalities (such as the species *cat*) and such *ad hoc* totalities as, for example, {the moon, Napoleon, justice}.

Set theory has problems, too, when it comes to dealing with time, and with the fact that biological species and similar entities may remain the same even when there is a turnover in their instances. For sets are identical if and only if they have the same members. If we model the species *cat* as the set of its instances, then this means that cats form a different species every time a cat is born or dies. If, similarly, we identify an organism as the set of its cells, then this means that it becomes a different organism whenever cells are gained or lost.

Set theory has problems also when it comes to dealing with the relations between granularities. An organism is a totality of cells, but it is also a totality of molecules, and it is also a totality of atoms. Yet the corresponding *sets* are distinct, since they have distinct members.

More recently, attempts have been made to solve some of these problems by using mereology or the theory of part and whole relations (Smith, 1998) as a framework for ontological theorizing. Mereology is better able to do justice in realistic fashion to the relations between wholes and their constituent parts at distinct levels of granularity. All the above-mentioned totalities (of cells, molecules, atoms) are, when treated mereologically, one and the same. Mereology also has the advantage over set theory when it comes to serving as a tool for the sort of middle-level ontological theorizing which the study of common-sense reasoning requires. For mereology does not require that, in order to quantify over wholes of given sorts, one must first of all specify some level of ultimate parts (the *Urelemente* of set theory) from out of which all higher-level entities are then constructed.

But mereology, too, has its problems. Thus it, too, has no way it has no way of dealing with entities which gain and lose parts over time and it has no way of distinguishing intrinsically unified wholes from *ad hoc* aggregations. Above all, its machinery for coping with the phenomenon of granularity brings problems of its own, for if we quantify over wholes, in a mereological framework, then we thereby quantify over all the parts of such wholes, at all levels of granularity. The selectivity of intentionality means however that we are often directed cognitively to coarse-grained wholes whose finer-grained parts are traced over: when I think of Mary I do not think of all the molecules in Mary's arm. Mereology has no means of mimicking the advantages of set-theory when it comes to dealing with such phenomena, and it is no small part of our project here to rectify this defect. The theory of granular partitions is the product of an effort to build a more realistic, and also a more general and flexible, framework embodying the strengths of both set theory and mereology while at the same time avoiding their respective weaknesses.

2 TYPES OF GRANULAR PARTITIONS

Some types of granular partitions are flat: they amount to nothing more than a mere list. Others are hierarchical: they consist of cells and subcells, the latter being nested within the former. Some partitions are built in order to reflect independently existing divisions on the side of objects in the world (the subdivision of hadrons into baryons and mesons, the subdivision of quarks into *up*, *down*, *top*, *bottom*, *charm*, *strange*). Other partitions—for example the partitions created by nightclub doormen or electoral redistricting commissions—are themselves such as to create the corresponding divisions on the side of their objects, and sometimes they create those very objects themselves. Quite different sorts of partitions—having cells of different resolutions and effecting unifyings and slicings and reapportionings of different types—can be applied simultaneously to the same domain of objects. The people in your building can be divided according to gender, social class or social security number. Or they can be divided according to tax bracket, blood type, current location or Erdös number. Maps, too, can impose subdivisions of different types upon the same domain of spatial reality, and the icons which they employ represent objects in granular fashion (which means that they do not represent the corresponding object parts). Maps will turn out to be important examples of granular partitions in the sense intended here.

The theory of partitions is, as will by now be clear, highly general, and this generality brings with it a correspondingly highly general reading of the term 'object'. Here we take an object to be any portion of reality: an individual, a part of an individual, a class of individuals (for example a biological species), a spatial region, a political unit (county, polling district, nation), or even (for present purposes) the universe as a whole. An object in the partition-theoretic sense is everything (existent) that can be recognized by some cell of a partition.

Objects can be either of the bona fide or of the fiat sort (Smith, 2001a). Bona fide objects, for example the moon, your armchair, this piece of cheese, are objects which exist (and are demarcated from their surroundings) independently of human partitioning activity. Fiat objects are objects which exist (and are demarcated from their surroundings) only because of such partitioning activity. Examples are: census tracts, your right arm, the Western Hemisphere.

In some cases partition cells recognize pre-existing fiat objects, in other cases the latter are created through the very projection of partition cells onto a corresponding portion of reality. Examples are the partitions creating the States of Wyoming and Montana, or the partitions of a population into persons belonging to distinct tax brackets created by tax legislation. Once fiat objects have been created in this way subsequent partitions may simply recognize them (without any object-creating effect), just as there are partitions which simply recognize bona fide objects..

Our notion of granular partition is only distantly related to the more familiar notion of a partition defined in terms of equivalence classes. Our partitions can include more structure in the form of hierarchically arranged subcells and supercells. Moreover, it is possible to define a partition in terms of an equivalence relation only where the relevant domain has already been divided up into units (the elements of the set with which we begin). The very process of division into units—for example through the imposition of fiat subdivisions or fiat discretizations upon continuous gradations—is however one of the things which our present theory is designed to illuminate.

In Smith and Brogaard (2002b) the notion of granular partition was introduced as a generalization of David Lewis's (1991) conception of classes as the mereological sums of their constituent singletons. Granular partitions, too, can in first approximation be conceived as the mereological sums of their constituent *cells*. The cells within a granular partition may however manifest a range of properties which the singletons of set theory lack. This is because, where a singleton is defined in the obvious way in terms of its member, each cell of a granular partition is defined by its *label*, and this means: independently of any object which might fall within it. The cells of a partition are what they are independently of whether there are objects located within them. A map of Middle Earth is different from a map of the Kingdom of Zenda, even though there is in both cases precisely nothing on the side of reality upon which these maps would be projected. 'The Morning Star' and 'The Evening Star' were for a long time used as labels for two distinct cells in astronomers' partitions of the heavenly bodies, even though, as it later turned out, it is one the same object that is located in each.

If one thinks that there are dodos, then one makes a different sort of error from the error which one makes if one thinks that there is an intra-Mercurial planet. (Set theory, it almost goes without saying, lacks the machinery to deal with such different sorts of error.)

Just as when we point our telescope in a certain direction we may fail to find what we are looking for, so when we point our partition in a certain direction it may be that there are no objects located in its cells. There may, in this sense, be empty cells within a partition (and even, in the most general version of our theory, partitions all of whose cells are empty). But this does not mean that the theory of partitions recognizes some counterpart of the set theorist's empty set (an entity that is contained as a subset within every set). For the empty set is empty by necessity; a cell in a partition, in contrast, is at best empty *per accidens*; it is empty because of some failure on our part in our attempts to partition the reality beyond.

The theory of partitions is thus more powerful than set theory, in that it is better able to do justice to the various ways in which human beings are related, cognitively, to objects in reality. In many ways, however, partition theory is also much weaker than set theory. For the axioms of set theory imply the existence of an entire hierarchy of sets, sets of sets, and so on, *ad infinitum*, reflecting the fact that they were designed to yield an instrument of considerable mathematical power. Partition theory, in contrast, is like mereology in that it is attuned to purposes other than those of mathematics. More specifically, it is designed to do justice to the sometimes *ad hoc* ways in which cognitive classificatory instruments are constructed by human beings for specific human purposes.

Partition theory differs from set theory also in this: that it puts partitions and objects in two entirely separate realms. Partitions themselves are never objects, and there are no partitions of partitions. Thus partition theory has no counterpart of sets of sets or of the distinction between two ways in which one set can be contained within another (on the one hand as element, on the other hand as subset). Partition theory can thus provide a framework for theorizing about the relations between cognitive artifacts such as lists and maps and the reality to which such artifacts relate in such a way that debates for example concerning the status of the hierarchy of transfinite sets can be avoided.

3 GRANULAR PARTITIONS AS SYSTEM OF CELLS

3.1 A bipartite theory

In the present paper we present the basic formal theory of granular partitions, leaving for a later work the presentation of the theory of cell-labeling (and, more generally, of the cognitive aspects of partitions as we here understand them). Our formal theory has two orthogonal and independent parts: (A) a theory of the relations between cells, subcells, and the partitions in which they are contained; (B) a theory of the relations between partitions and objects in reality. The counterpart of (A) in a set-theoretic context would be the study of the relations among subsets of a single set; the counterpart of (B) would be the study of the relations between sets and their members. These set-theoretical counterparts of (A) and (B) are, be it noted, not independent. This is because the standard subset relation of set theory is itself defined in terms of the set-membership relation (x is a subset of y means: all the members of x are members of y). In the context of partition theory, in contrast, the corresponding relations are defined independently of each other. Partition theory thus departs from the extensionalism of set theory (i.e. from the assumption that each set is defined exclusively by its members). A cell is defined by its position within a partition and by its relations to other cells, and it is this which gives rise to the relations treated of by theory (A). What objects in reality are located in a cell—the matter of theory (B) —is then a further question, which is answered, in different ways from case to case. Briefly, we can think of cells as being projected onto objects in something like the way in which flashlights are projected upon the objects which fall within their purview.

Consider the left part of Figure 1. Theory A governs the way we organize cells into nesting structures and the way we label cells. Theory B governs the way these cell-structures project onto reality indicated by the arrows connecting the left and the right parts of the Figure. Our strategy in what follows will be first of all to define a series of master conditions belonging to theory (A) and theory (B) respectively, and which—for the purposes of the present paper—all partitions will be assumed to satisfy. In later sections we will add further conditions, satisfied by some partitions but not by others.

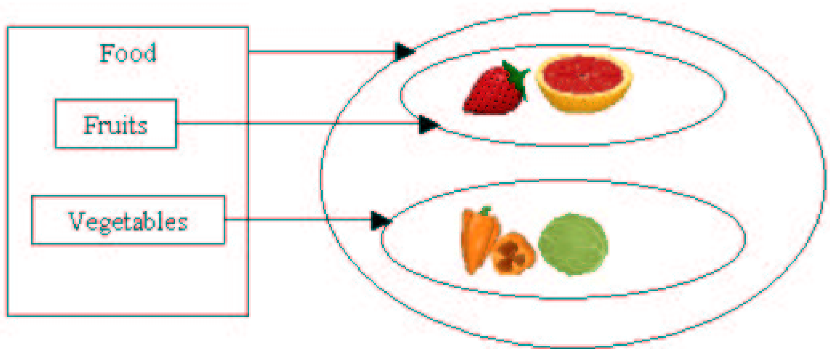

Figure 1: Relationships between cells and objects

3.2 The subcell relation

Theory (A) is effectively a theory of well-formedness for partitions; it studies properties partitions have in virtue of the relations between and the operations performed upon the cells from out of which they are built, independently of any linkage to reality beyond. Cells in partitions may be nested one inside another in the way in which, for example, the species *crow* is nested inside the species *bird*, which in turn is nested inside the genus *vertebrate* in standard biological taxonomies. When one cell is nested inside another in this way we say that the former is a sub-cell of the latter. Note that the subcell relation can hold between two cells independently of whether there are any objects located in them (as for example in relation to the cells labeled ‘male dodos’ and ‘dodos’ in a classification of extinct animals).

We use $z, z_1, z_2, \ldots$ as variables ranging over cells and $A, A_1, A_2, \ldots$ as variables ranging over partitions ('cell' is '*Zelle*', partition is '*Aufteilung*' in German). We write $z_1 \subseteq_A z_2$ in order to express the fact that z_1 stands in a sub-cell relation to z_2 within the partition A. (Where confusion will not result we will drop the explicit reference to the partition A and write simply '$\subseteq$'). We can then state the first of several master conditions on all partitions as follows:

MA1: The subcell relation $\subseteq$ is reflexive, antisymmetric, and transitive.

This means that within every partition: each cell is a subcell of itself; if two cells are subcells of each other then they are identical; and if cell z_1 is a sub-cell of z_2 and z_2 a sub-cell of z_3, then z_1 is in its turn a sub-cell of z_3. We can think of the sub-cells of a cell within a given partition as special sorts of *parts* of the cell; they are those parts which are included within this same partition as cells in their own right.

3.3 Existence of a maximal cell

We define a *maximal cell* of a partition A as a cell satisfying:

DMax: $M(z_1, A) \equiv Z(z_1, A)$ and $\forall z : Z(z, A) \rightarrow z \subseteq z_1$.

Here '$Z(z, A)$' means that z is a cell in the partition A. (Again: we shall normally omit the condition $Z(z, A)$ where confusion will not result.) We now demand as a further master condition that

MA2: Every partition has a unique maximal cell in the sense of DMax.

The motivation for MA2 is very simple: it turns on the fact that a partition with two maximal cells would either be in need of completion by some extra cell representing the result of combining these two maximal cells together into some larger whole; or it would not be one partition at all, but rather two separate partitions, each of which would need to be treated in its own right within the framework of our theory.

We also call the unique maximal cell of a partition its root, $r(A)$. The maximal cell of a partition is such that all the cells in the partition are included in it as subcells. MA2 implies that there are no partitions which are empty *tout court* in that they have no cells at all.

3.4 Finite chain condition

The transitivity of $\subseteq$ generates a nestedness of cells inside a partition in the form of chains of cells satisfying $z_1 \supset z_2 \supset \ldots \supset z_n$, with z_1 as root. We shall call the cells at the ends of such chains *minimal cells* or *leaves*, and define:

DMin: $Min(z_1, A) \equiv Z(z_1, A)$ and $\forall z : Z(z, A) \rightarrow (z \subseteq z_1 \rightarrow z = z_1)$

Another important aspect of a partition is then:

MA3: Each cell in a partition is connected to the root via a finite chain of immediate succeeding cells.

A cell z_2 is the immediate successor of the cell z_1 if and only if $z_1 \subseteq z_2$ and there does not exist a cell z_3 such that $z_1 \subset z_3 \subset z_2$ holds.

MA3 does not rule out the possibility that a given cell within a partition might have infinitely many immediate subcells (also called daughter cells). Enforcing finite chains thus leaves open the issue as to whether partitions themselves are finite.

If, in counting off the cars passing beneath you on the highway, your checklist includes one cell labeled *red cars* and another cell labeled *Chevrolets*, we will rightly feel that there is something amiss with your partition. One problem is that you will almost certainly be guilty of double counting. Another problem is that there is no natural relationship between these two cells, which seem rather to belong to distinct partitions. As a step towards rectifying such problems we shall insist that all partitions must satisfy a condition according to which every pair of distinct cells within a partition stand to each other either in the subcell relation or in the relation of disjointness. In other words:

MA4: If two cells within a partition overlap, then one is a subcell of the other.

Or in symbols:

$$\exists z : (z \subseteq z_1 \text{ and } z \subseteq z_2) \rightarrow z_1 \subseteq z_2 \text{ or } z_1 \supset z_2.$$

(Here and in what follows initial universal quantifiers are taken as understood.) From MA3 and MA4 we can prove by a simple *reductio* that the chain connecting each cell of a partition to the root is unique.

3.5 Partition-theoretic sum and product of cells

The background to all our remarks in this paper is mereology. We take the relation $\leq$ meaning 'part of' as primitive, and define the relation of overlap between two entities simply as the sharing of some common part. $\leq$ is like $\subseteq$ in being reflexive, anti-symmetric and transitive, but the two differ in the fact that $\subseteq$ is a very special case of $\leq$.

The subcells of a cell are also parts of the cell (just as, for David Lewis, 1991, each singleton is a part of all the sets in which it is included). What happens when we take the mereological products and sums of cells existing within a partition? In regard to the mereological product, $z_1 * z_2$, of two cells matters are rather simple. This product exists only when the cells overlap mereologically, i.e. only when they have at least one subcell in common. This means that the mereological product or intersection of two cells, if it exists, is in every case just the smaller of the two cells.

In regard to the mereological *sum* of cells $z_1 + z_2$, in contrast, it is a more difficult situation which confronts us. Given any pair of cells within a given partition, the corresponding mereological sum does indeed exist—simply in virtue of the fact that the axioms of mereology allow unrestricted sum-formation. (This is a trivial matter, for the mereologist: if you got the parts, whatever they are, then you got the whole.) But only in special cases will this mereological sum be itself a cell within the partition in question. This occurs for example when cells labeled 'male rabbit' and 'female rabbit' within a partition have as their sum the cell labeled 'rabbit'. There is, in contrast, no cell in our standard biological partition of the animal kingdom labeled *rabbits and jellyfish*, and there is no cell in our standard geopolitical partition of the surface of the globe labeled *Hong Kong and Algeria*.

To make sense of these matters we need to distinguish the mereological sum of two cells from what we might call their partition-theoretic sum. We can define the former as just the result of taking the two cells together in our thoughts and treating the result as a whole. We can define the latter as follows. The partition-theoretic sum $z_1 \cup z_2$ of two cells in a partition is the smallest subcell within the partition containing both, z_1 and z_2, as subcells; i.e., it is the least upper bound of z_1 and z_2 with respect to $\subseteq$. (By MA2 and MA4 we know that this is always defined and that it is unique.) This partition-theoretic sum is in general distinct from the mereological sum of the corresponding cells. (The partition-theoretic sum of the cells labeled *rabbit* and *lion* is the cell labeled *mammal* in our partition of the animal kingdom.) The best we can say in general is that $z_1 + z_2$ is at least part of $z_1 \cup z_2$ (Smith, 1991). Note, on the other hand, that if we analogously define the partition-theoretic product, $z = z_1 \cap z_2$, of two cells within a given partition as the largest subcell shared in common by z_1 and z_2, i.e., as their greatest lower bound with respect to $\subseteq$, then it turns out that this coincides with the mereological product already defined above.

Mereological sum and product apply to both cells and objects; partition-theoretic sum applies only to cells. Here we use the symbols for the two groups of relations as shown in Table 1:

	Partition-theoretic (for cells)	Mereological (for cells and for objects)
Sum	$\cup$	$+$
Product	$\cap$	$*$
Inclusion	$\subseteq$	$\leq$
Proper Inclusion	$\subset$	$<$

Table 1: Partition-theoretic and mereological relations and operations.

When restricted to cells within a given partition $\subseteq$ and $\leq$ coincide, and so also do $\cap$ and $*$. We can think of $\subseteq$ as the result of restricting $\leq$ to the *natural units* picked out by the partition in question. We can think of set theory as amounting to the abandonment of the idea that there is a distinction between natural units and arbitrary unions. Set theory, indeed, derives all its power from this abandonment.

3.6 Trees

Philosophers since Aristotle have recognized that the results of our sorting and classifying activities can be represented as those sorts of branching structures which mathematicians nowadays called trees. Trees are directed graphs without cycles. They consist of nodes or vertices and of directed edges that connect the nodes. That the edges are directed means that the vertices connected by an edge are related to each other in a way that is analogous to an ordered pair. Here we are interested specifically in rooted trees, which is to say: trees with a single topmost node to which all other vertices are connected, either directly or indirectly, via edges. In a rooted tree, every pair of vertices is connected by one and only one chain (or sequence of edges). We shall think of the directedness of an edge as proceeding down the tree from top to bottom (from ancestors to descendants). That a tree is without cycles means that, if we move along its edges, then we will always move down the tree and in such a way that, however far we travel, we will never return to the point from which we started.

The connection between partitions and trees will now be obvious: it is a simple matter to show that every finite partition can be represented as a rooted tree of finite depths and vice versa (Mark, 1978). To construct a tree from a finite partition we create a graph by mapping the cells z_i of the partition onto nodes v_i within the graph and by introducing a directed edge from vertex v_i to v_j if and only if the cell z_i has cell z_j as an immediate subcell. That this is always possible follows from the fact that the subcell relation is well defined (by MA1) and from the fact that chains of immediate cells are always finite (MA3). We can easily show also that the resulting graph is a rooted tree, which follows from MA2; that the graph structure is connected (from MA2), and acyclical (from MA4); and that there is a unique path between any two vertexes (from MA2, MA3 and MA4). The complementary reconstruction of a partition from its tree representation is no less trivial.

We can represent a partition not only as a tree but also as a simple sort of Venn diagram. In a Venn diagram partition cells are represented as topologically simple and regular regions of the plane. Our partitions are Venn diagrams within which regions do not intersect. (Conversely every array of non-intersecting, possibly nested regions in the plain can be transformed into a tree in such a way that each region is represented by a node in the tree, and each directed link in the tree represents an *immediately contains* relation between a corresponding pair of nested regions.) In the remainder we will often think of partitions as such planar maps (that is as Venn diagrams without overlapping), and the minimal cells correspond to the smallest regions within such diagrams.

Tree and Venn-diagram representations of granular partitions are not equivalent. To see this consider Figure 2. Even if we ignore the labeling it is obvious that the two Venn-diagrams represent two distinct partitions. The mammal-partition contains 'empty space' and the first-couple-partition is full in the sense that it does not contain 'empty space'. This distinction, however, can not be made in terms of the corresponding tree representations. In order to represent it in the tree we needed consider labeled trees with nodes labeled *full* or *not-full*. We will discuss these issues in more detail in our section on fullness and cumulativeness of granular partitions.

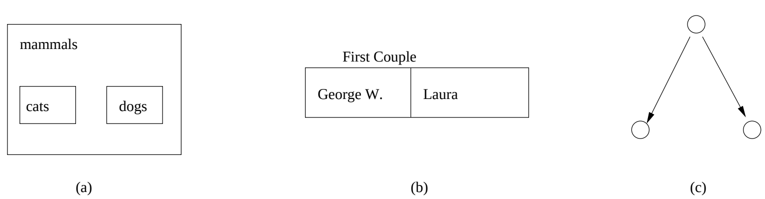

Figure 2: Venn-diagram and tree representations of granular partitions.

4 GRANULAR PARTITIONS IN THEIR PROJECTIVE RELATION TO REALITY

4.1 Projection

Partitions are more than just systems of cells. They are constructed to serve as inventories or pictures or maps of specific portions of reality, and in this they are analogous to *windows*, or to the latticed grills purported to have been used by Renaissance artists as aids to the faithful representation of objects in reality (Smith, 2001b). They are analogous also to propositions (*Elementarsätze*) as described by Wittgenstein in the *Tractatus* (1961). A proposition, for Wittgenstein, is built out of simple signs (names) arranged in a certain order. Each name, Wittgenstein tells us, stands in a projective relation to a corresponding object in the world: it cannot fail to strike its target. If a proposition is true, then its simple signs stand to each other within the proposition as the corresponding objects stand to each other in the world. It is in this sense that a true atomic proposition is a picture, as Wittgenstein puts it, of a state of affairs in reality. That a proposition is a complex of names arranged in a certain order is in our present context equivalent to the thesis that a partition is a complex of cells arranged in a certain order.

A partition is a complex of cells in its projective relation to the world (compare *Tractatus*, 3.12). This relation may be effected either directly by the user of the partition—for example in looking through the cells of the grid and recording what objects are detected on the other side—or indirectly, with the help of proper names or other referring devices such as systems of coordinates or taxonomic labels.

For Wittgenstein it is guaranteed a priori for every name that there is some unique object onto which the name is projected. From the perspective of the theory of granular partitions, in contrast, projection may fail. That is, a partition may be such that—like the partition cataloguing Aztec gods—there are no objects for its cells to project onto. Works of fiction and also not yet realized plans may be conceived as involving partitions of this kind.

In this paper, however, we are interested primarily in partitions which do not project out into thin air in this way. We write '$P(z,o)$' as an abbreviation for: cell z is projected onto object o. We can also, if the context requires it, write '$P_A(z,o)$' to indicate that the projection of z onto o obtains in the context of partition A. In what follows we shall assume that a unique such projection is defined for each partition. In a more general theory we can weaken this assumption, for example by allowing projections to vary with time while the partition remains fixed (Smith and Brogaard, 2002a). Such variation of projection for a fixed partition is involved in all sampling activity. Consider, for example, what happens when we use a territorial grid of cells to map the presence of one or more birds of given species in given areas from one moment to the next.

4.2 Location

If projection is successful, then the object upon which a cell is projected is located in that cell. The use of the term 'location' reflects the fact that one important inspiration of our work is the study of location relations in spatial contexts. One motivating example of a location relation within our theory is the relation between a spatial object such as a railway station and an icon on a map. Other motivating examples are of a non-spatial sort: they include the relation between an instance (Tibbles) and its kind (cat) or the relation between a customer and the corresponding record in a database. Indeed they include the relation between you and your name.

We can compare a partition with a rig of spotlights projecting down onto an orchestra during the performance of a symphony. Each cell of the partition corresponds to some spotlight in the rig. Some cells (spotlights) will project upon single players, others onto whole sections of the orchestra (string, wind, percussion, and so forth). One cell (spotlight) will project upon the orchestra as a whole. Note that the spotlights do not hereby create the objects which they cast into relief. When once the rig has been set, and the members of the orchestra have taken their places, then it will be an entirely objective matter which objects (individuals and groups of individuals) are located in which illuminated cells.

In what follows we make the simplifying assumption that objects are exactly located at their cells (that spotlights never partially illuminate single players or sections). Compare the way in which Wyoming is exactly located at the cell 'Wyoming' in the partition of the US into States or the way in which your brother Norse is exactly located at the cell 'Norse' in your partition (list) of your family members. In a more general theory we liberalize the location relation in such a way as to allow also for partial or rough location (Casati and Varzi, 1995; Bittner and Stell, 1998).

4.3 Transparency

When projection succeeds, then location is what results. Projection and location thus correspond to the two directions of fit—from mind to world and from world to mind—between an assertion and the corresponding truthmaking portion of reality (Searle, 1983; Smith, 1999). Projection is like the relation which holds between your shopping list and the items which, if your shopping trip is successful, you will actually buy. Location is like the relation which obtains between the items you have bought and the new list your mother makes after your return, as she checks off those items which you have in fact succeeded in bringing back with you.

The formula '$L(o, z)$' abbreviates: object o is located at cell z. (And again where this is required we can write '$L_A(o, z)$' for: o is located at z in partition A.) Location presupposes projection: an object is never located in a cell unless the object has already been picked out as the target of the projection relation associated with the relevant partition. But *successful* projection—by which is meant the obtaining of the projection relation between a cell and an object—also presupposes location, so that where both L and P obtain they are simply the converse relations of each other. We have now reached the point where we can formulate the first of our master conditions on partitions from the perspective of theory (B):

MB1: $L(o, z) \rightarrow P(z, o)$
MB2: $P(z, o) \rightarrow L(o, z)$

(Successful) projection and (successful) location are simple converses of each other. (We formulate this principle as two separate conditions in order to leave room for a more general theory in which these two relations are teased apart.)

MB1 and MB2 tell us that a partition projects a given cell onto a given object if and only if that object is located in the corresponding cell. Very many partitions—from automobile component catalogues to our maps of states and nations—have this quality without further ado.

We shall call partitions which satisfy MB1 and MB2 transparent partitions, a notion which we can define in the obvious way as follows:

DTr: $Tr(A) \equiv \forall z \forall o : P_A(z, o) \leftrightarrow L_A(o, z)$

MB1 and MB2 jointly ensure that objects are actually located at the cells that project onto them. Notice however that a transparent partition, according to our definition, may still have empty cells. Such cells may for example be needed in the context of scientific partitions in order to leave room for what, on the side of the objects, may be discovered in the future. (Compare the cells labeled Ununnilium, Unununium and Ununbium in the Periodic Table of the Elements.) Empty cells may similarly be needed to cover up for temporary lapses in memory. You are attempting to account for the people at your party last night. Your partition consists of six cells labeled: *John*, *Mary*, *Phil*, *Chris*, *Sally*, and *anyone else* (for people you might have forgotten). Assume that John, Mary, Phil, Chris and Sally is a complete listing of all the people at the party. Your *anyone else* cell is then empty.

4.4 Functionality constraints (constraints pertaining to correspondence to objects)

4.4.1 Projection is functional: the confused schoolboy

The property of transparency is still rather weak. Thus transparency is consistent with ambiguity on the side of the cells in relation to the objects they target, that is with the case where one cell projects onto two distinct objects. An example of the sort of problem we have in mind is the partition created by a lazy schoolboy studying the history of the Civil War in England. This partition has one cell labeled 'Cromwell'—and so it does not distinguish between Oliver and his son Richard. Another example might be the partition utilized by those who talk of 'China' as if the Republic of China and the People's Republic of China were one single object.

To eliminate such ambiguity we lay down a requirement to the effect that each partition must be such that its associated projection is a *functional* relation:

MB3: $P(z, o_1) \text{ and } P(z, o_2) \rightarrow o_1 = o_2$

For partitions satisfying MB3, cells are projected onto single objects (one rather than two). Consider the left part of Figure 3. The dotted arrow can occur in partitions satisfying merely MB1–2 but not in partitions also satisfying MB3. Notice, though, that projection might still be a partial function, since MB3 does not rule out the case where there are empty cells.

To impose the functionality of projection on all partitions is in one respect trivial. For we can very easily convert a partition A which does not satisfy MB3 into one which does. If z is a cell in A which does not satisfy MB3, then we create this new partition A' by adjusting A in such a way that z now projects upon the mereological sum of the objects

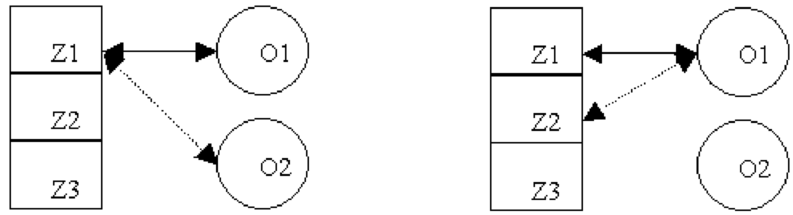

Figure 3: Transparent partitions in which projection is not functional (left); location is not functional (right)

its projects upon in A. This account seems, indeed, to do justice to what is involved in the confused schoolboy case, namely that Richard and Oliver are run together, somehow, into one composite human being.

In the remainder of this paper we use the notation $o = p(z)$ instead of $P(z, o)$ whenever we assume that projection is functional.

4.4.2 Location is functional: the Morning Star and the Evening Star

Consider a partition having root cell labeled ‘heavenly bodies’ and three subcells labeled: ‘The Morning Star’, ‘The Evening Star’, and ‘Venus’, respectively. As we know, all three subcells project onto the same object. This partition is perfectly consistent with the conditions we have laid out thus far. Its distinct subcells truly, though unknowingly, project onto the same object. It is not unusual that we give different names (or class-labels) to things in cases where we do not know that they are actually the same. A good partition, though, should clearly be one in which such errors are avoided.

Partitions manifesting the desired degree of correspondence to objects in this respect must in other words be ones in which location, too, is a *functional* relation:

MB4: $L(o, z_1)$ and $L(o, z_2) \rightarrow z_1 = z_2$

In partitions that satisfy MB4, location is a function, i.e., objects are located at single cells (one rather than two). Consider the right part of Figure 3. The dotted arrow can occur in partitions satisfying MB1–2, not however in partitions also satisfying MB4. As MB3 rules out co-location (overcrowding), so MB4 rules out co-projection (redundancy). Note that natural analogues of co-location and co-projection are not even formulable within a set-theoretic framework.

5 CORRESPONDENCE OF MEREOLOGICAL STRUCTURE

MB1 and MB2 are, even when taken together with MB3 and MB4, still very weak. They tell us only that, if a cell in a partition projects upon some object, then that object is indeed located in the corresponding cell. They do not tell us what happens in case a cell fails to project onto anything at all. MB1–4 thus represent only a first step along the way towards an account of correspondence to reality for partitions. Such correspondence will involve the two further dimensions of *structural mapping* and of *completeness*.

5.1 Recognizing mereological structure

An object o is *recognized* by a partition if and only if the latter has a cell in which that object is located (Smith and Brogaard, 2002b). Intuitively, recognition is the partition-theoretic analogue of the standard set-membership relation. Partitions embody the selective focus of our mapping, classifying, and listing activities. To impose a partition on a given domain of reality is to *foreground* certain objects and features in that domain and trace over others. Note hereby that we trace over not only the objects which *surround* that which is foregrounded, as according to the usual understanding of the foreground/background structure; for we also trace over those *parts* of the foregrounded object which fall beneath the threshold of our concerns. Partitions are granular in virtue precisely of the fact that a partition can recognize an object without recognizing all its parts.

Partitions—think again of Venn diagrams—are designed to reflect the part-whole structure of reality through the fact that the cells in a partition are themselves such as to stand in relations of part to whole. Given the master conditions expressed within the framework of theory (A) above, partitions have at least the potential to reflect the mereological structure of the domain onto which they are projected. And in felicitous cases this potential is realized.

That we distinguish between the recognition (foregrounding, selection) of objects on the one hand and the reflection of mereological structure on the other hand is not an arbitrary matter. In Tractarian semantics we distinguish between projection and isomorphism. In set theory we distinguish, for any given set, between a domain of elements and the set-theoretic structure imposed on this domain. Just as it is possible to have sets consisting entirely of *Urelemente* (together with a minimal amount of set-theoretic packaging), so it is possible to have partitions built exclusively out of minimal cells (and one root cell). Such partitions amount, simply, to lists of the things that are recognized by their cells, with no mereological structure on the side of these objects being brought into account.

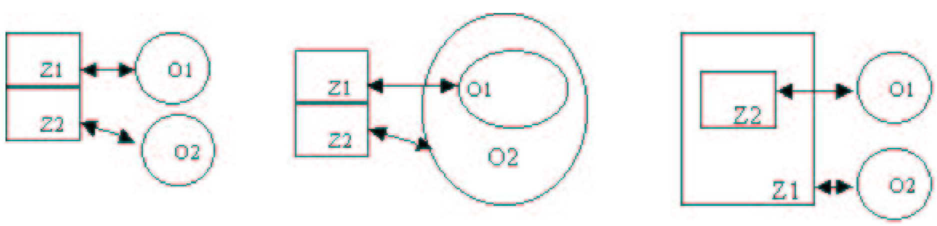

(a) (b) (c)

Figure 4: Transparent partitions with more or less desirable properties

Figure 4(a) and 4(b) represent partitions consisting of two minimal cells z_1 and z_2 projecting onto objects o_1 and o_2. Case (a), a simple list, is unproblematic. Case (b) we shall also allow. This is in keeping with the notion that minimal cells are the (relative) atoms of our system, and we take this to mean that they should be neutral with regard to any mereological structure on the side of their objects. An example of type (b) would be a list of regions represented at a conference to discuss measures against terrorism, a conference including representatives from both Germany and Bavaria.

Cases like (c), in contrast, represents projections in which, intuitively, something has gone wrong. All three cases satisfy the master conditions we have laid down thus far, for the latter allow both for disjoint cells to be projected onto what is not disjoint (b) and also for disjoint objects to be located in cells which are not disjoint (c). Cases like (c) on the other hand seem to fly in the face of a fundamental principle underlying the practice of hierarchical classification, namely that objects recognized by species lower down in a hierarchical tree should be included as parts in whatever is recognized by the genera further up the tree. To exclude cases like (c) we shall impose a condition to the effect that mereological structure within a partition should *not misrepresent the mereological relationships* between the objects which the corresponding cells are projected onto. We first of all define the following relation of *representation of mereological structure* between pairs of cells:

DS1: $RS(z_1, z_2) \equiv \forall o_1, o_2 : (L(o_1, z_1) \text{ and } L(o_2, z_2) \text{ and } z_1 \subseteq z_2) \rightarrow o_1 \leq o_2$

If z_1 is a subcell of z_2 then any object recognized by z_1 must be a part of any object recognized by z_2. A partition is then mereologically structure-preserving if and only if each pair of cells within the partition satisfies DS1:

DS2: $RS(A) \equiv \forall z_1, z_2 : (Z(z_1, A) \text{ and } Z(z_2, A)) \rightarrow RS(z_1, z_2)$.

We can now impose a new master condition:

MB5: All partitions are *mereologically structure-preserving* in the sense of DS2.

Note that even MB5 is still very weak. Its effect is entirely negative, since it merely ensures that partitions do not misrepresent the mereological relationships between their objects. Partitions might still be entirely blind to (trace over) such relationships. Two minimal cells might project onto objects which stand to each other in any one of the possible mereological relations (identity, proper parthood, disjointness, overlap), and all pairs of cells are likewise neutral as to the mereological relations between the objects onto which they are projected provided only that they do not stand to each other in the subcell relation. This means that, given such cells, we are entitled to infer nothing at all about the mereological relations among the corresponding objects.

Consider, for example, a partition that contains cells, z_1 and z_2, that recognize John and his arm, respectively, so that $L(John, z_1)$ and $L(John's\ arm, z_2)$. Then cell z_1 need not be a proper subcell of the cell z_2, for the partition may not know that the object located in z_2 is properly designated as John's arm. Or consider a partition containing two cells that recognize, respectively, mammals and whales. Suppose that this is a partition constructed at a time when the status of whales as mammals was not yet recognized. The cell labeled *whales* is not, then, included as a subcell of the cell labeled *mammals*. But the partition can still satisfy our conditions laid down so far. This is so, for example, if the cell that recognizes whales is a subcell of the cell recognizing animals but not a subcell of any other subcell of the cell recognizing animals (Partition A_1 in Figure 5). If the cell that recognizes whales were also a subcell of the cell that recognizes fish, for example, then the partition would misrepresent the mereological relationship between these two species and so violate MB5 (Partition A_2 in Figure 5).

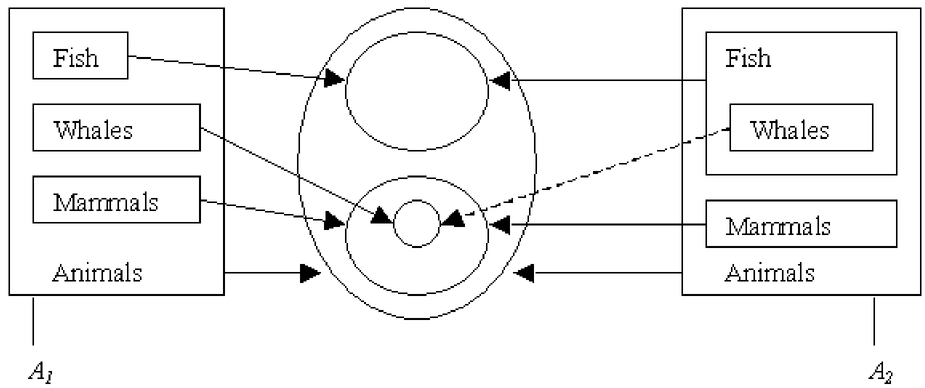

Figure 5: Partition A_1 does not misrepresent the mereological structure of the underlying domain. Partition A_2 places whales incorrectly in relation to fish and mammals

Partitions may trace over mereological relationships between the objects they recognize, but MB5 is strong enough to ensure that, if a partition tells us something about the mereological relationships on the side of the objects which it recognizes, then what it tells us is true. Notice that partition A_2 still satisfies MB1–4.

Consider a domain of objects consisting of two regions, x and y, that properly overlap in the region v, so that $x * y = v$ with $v < x$ and $v < y$. Consider now a partition with cells z_1 and z_2 recognizing x and y, respectively, so that $L(x, z_1)$ and $L(y, z_2)$. Assume further that z_1 and z_2 do not stand in any subcell relation to each other, i.e., their partition-theoretic intersection is empty. Only four possibilities regarding the representation of v now remain: (1) our partition does not recognize v at all; (2) it recognizes v but traces over its mereological relationships to x and y; (3) it recognizes v through a subcell of z_1 but it traces over the fact that v could equally well be recognized by a subcell of z_2; (4) it recognizes v through a subcell of z_2 but it traces over the fact that v could equally well be recognized by a subcell of z_1. The fifth possibility—of allowing sub-cells of both z_1 and z_2 to recognize v is excluded by the tree structure of granular partitions.

Let x and y be two neighboring countries which disagree about the exact location of their common boundary and let v be the disputed area. The inhabitants of country x consider v to be part of x, the inhabitants of country y consider v to be part of y. Possibility (1) then corresponds to the view of some third country at the other side of the globe who recognizes the countries x and y but does not care about their border dispute. (2) corresponds to the view of an observer who recognizes that there is a disputed area but who is neutral about the status of the disputed area. (3) corresponds to the view of country x and (4) to that of country y.

Another example of case (2) is provided by Germany and Luxemburg, which overlap at their common border on the River Our. The river is part of both countries. Mapmakers normally have no facility to represent cases such as this, and so they either adopt the policy of not representing such common regions at all (the border is represented as a line which we are to imagine as being without thickness), or they recognize the region constituted by the river on the map but trace over its mereological properties. Larger-scale maps often embrace a third alternative, which is to *misrepresent* the relations between Germany and Luxenburg by drawing the boundary between the two countries as running down the center of the river.

5.2 The domain of a partition

That upon which a partition is projected is a certain domain of objects in reality (the term 'domain' being understood in the mereological sense). We shall conceive the domain of a partition as the mereological sum of the pertinent objects. It is, as it were, the total mass of stuff upon which the partition sets to work: thus it is stuff conceived as it is prior to any of the divisions or demarcations effected by the partition itself. The domains of partitions will comprehend not only individual objects and their constituents (atoms, molecules, limbs, organs), but also groups or populations of individuals (for example biological species and genera, battalions and divisions, archipelagos and diasporas) as well as their constituent parts or members. Domains can comprehend also extended regions (continua) of various types. Spatial partitions, for example maps of land use or soil type (Frank et al., 1997), are an important family of partitions with domains of this sort. There are also cases where partitions impose upon continuous domains a division into discrete units for example by creating temperature or frequency bands.

We are now able to specify what we mean by 'domain of a partition.' Our representation of partitions as trees and our condition on reflection of structure (MB5) ensure that all partitions trivially reflect the fact that the objects recognized by their cells are parts of some mereological sum. For MB5 is already strong enough to ensure that everything that is located at some cell of a partition is part of what is located at the corresponding root cell. If any cell pointed outside of what is located at the root cell it would misrepresent the mereological structure of the corresponding domain.

We can thus define the domain of a partition simply as the object (mereological whole) onto which its root cell is projected. By functionality of projection and location there can be only one such object.

DD: $D(A) = p(r(A))$

We now demand as a further master condition that every partition has a non-empty domain in the sense of DD:

MB6: $\exists x : x = D(A)$

We then say that a partition *represents its domain correctly* if and only if MA1–5 and MB1–6 hold. Note that this condition of *correctness* is still rather easily satisfied. (It is achieved already in every simple list, provided only that the list involves no double counting and no ambiguous reference of the sort involved in the Oliver and Richard Cromwell case.)

If there is a single maximal object (the whole universe), then one correct representation thereof is provided by a partition consisting of just one cell labeled ‘everything’ (we might call this the Spinoza partition). A partition with just three cells: a root cell, labeled *animals*, and two subcells, labeled *dogs* and *cats*, represents its domain correctly; it just falls far short of a certain desirable completeness. Correct representations, as we see, can be highly partial.

5.3 The granularity of granular partitions

A correct representation, as we see, is not necessarily a complete representation. Indeed, since partitions are cognitive devices, and cognition is not omniscient, it follows that no partition is such as to recognize all objects. There is no map of all the objects in the universe. The complexity of the universe is much greater than the complexity of any single cognitive artifact. This feature of partiality is captured already by our terminology of *granular* partitions. Partitions characteristically do not recognize the proper parts of the whole objects which they recognize; for example they do not recognize parts which fall beneath a certain size.

It is the cells of a partition which carry with them this feature of granularity. Because they function like singletons in set theory, they recognize only single whole units, the counterparts of set-theoretic elements or members. If a partition recognizes not only wholes but also one or more parts of such wholes, then this is because there are additional cells in the partition which do this recognizing job. Consider, for example, a partition that recognizes human beings and has cells that project onto John, Mary, and so forth. This partition does not recognize parts of human beings—such as John’s arm or Mary’s shoulder—unless we add extra cells for this purpose. Even if a partition recognizes both wholes and also some of their parts, it is not necessarily the case that it also reflects the mereological relationships between the two. Imagine we are forensic scientists examining photographs taken at a crime scene and that these photographs generate a partition with cells recognizing John, Mary, and an arm. It may then be the case that the state of our knowledge is such that the cell recognizing the arm is not a subcell of the cell recognizing either John or Mary. Or let the arm be Kashmir and let John and Mary be India and Pakistan, respectively.

In relation to this granularity of partitions, we can once more call in the aid of Wittgenstein:

> In the proposition there must be exactly as many things distinguishable as there are in the state of affairs, which it represents. They must both possess the same logical (mathematical) multiplicity ... (4.04)

Wittgenstein himself takes care of the issue of granularity by insisting that the world is made up of discrete simples, and by insisting further that all partitions (for Wittgenstein: propositions) picture complexes of such simples. (A similar simplifying assumption is proposed by Galton, 1999.) This is a simplifying assumption, which our present theory of granular partitions will enable us to avoid. For the latter admits partitions of arbitrary granularity including partitions which reflect distinct cross-cuttings of the same domain of reality (Smith and Brogaard, 2002a). The theory of granular partitions enables us moreover to remain neutral as to the existence of any ultimate simples in reality from out of which all other objects would be constructed via summation. This is due to the fact that partitions are by definition *top-down* structures. The duality with trees puts special emphasis on this aspect: we trace down from the root until we reach a leaf. A leaf has no further parts within the partition to which it belongs. But it need not necessarily project upon something that itself has no parts. The fact that there are leaves simply indicates that a partition does not care about (traces over) what lies beneath a certain level of granularity on the side of its objects. An object located at a minimal cell is an atom only relative to the partition involved.

Partitions are cognitive devices which have the built-in capability to recognize objects and to reflect certain features of the latters' mereological structure and to ignore (trace over) other features of this structure. We can now see that they can perform this task of tracing over in two ways: (1) by tracing over mereological relations between the objects which they recognize; (2) by tracing over (which means failing to recognize) parts of those objects. (2) is (unless atomism is true) a variety of tracing over that must be manifested by every partition. A third type of tracing over arises in reflection of the fact that partitions (we leave to one side here the Spinoza partition) are partial in their focus. In foregrounding some regions of reality each partition thereby traces over everything that lies outside its domain.

Consider a simple biological partition of the animal kingdom including a cell projecting on the species dog (*Canis familiaris*). Our definition of the domain of a partition and our constraint on functionality of projection implies that, besides the species dog also your dog Fido, and also Fido's DNA-molecules, proteins, and atoms are parts of the domain of this partition. *But the latter are of course not recognized by the partition itself.* It is cases such as this which illustrate why mereology requires supplementation by a theory like the one presented here. Partition theory allows us to define a new, restricted notion of parthood that takes granularity into account (compare Degen et al., 2001). This restricted parthood relation is an analogue of partition-theoretic inclusion, but on the side of objects:

DRP: $x \leq_A y \equiv \exists z_1, z_2 : L_A(x, z_1) \text{ and } L_A(y, z_2) \text{ and } z_1 \subseteq z_2$

This means that x is a part of y relative to partition A if and only if: x is recognized by a subcell of a cell in A which recognizes y. From this we can infer by MB5 that x is a part of y also in the unrestricted or absolute sense.

The usual common-sense (i.e., non-scientific) partition of the animal kingdom contains cells recognizing dogs and mammals, but no cells recognizing DNA molecules. Relative to this common-sense partition, DNA molecules are not parts of the animal kingdom in the sense defined by DRP, though they are of course parts of the animal kingdom in the usual, non-relativised sense of 'part'.

6 STRUCTURAL PROPERTIES OF CORRECT REPRESENTATIONS

In this section we discuss some of the more fundamental varieties of those partitions which satisfy the master conditions set forth above. We classify such partitions according to: (1) degree of structural fit; (2) degree of completeness and exhaustiveness; (3) degree of redundancy.

6.1 Mereological monotony

We required of partitions that they at least not misrepresent the mereological structure of the domain they recognize. This constraint is to be understood in such a way that it leaves room for the possibility that a partition is merely neutral about (traces over) some or all aspects of the mereological structure of its target domain. Taking this into account, we can order partitions according to the degree to which they actually do represent the mereological structure on the side of the objects onto which they are projected. At the maximum degree of structural fit we have those partitions which completely reflect the mereological relations holding between the objects which they recognize.

Such a partition satisfies a condition to the effect that if o_1 is part of o_2, and if both o_1 and o_2 are recognized by the partition, then the cell at which o_1 is located is a subcell of the cell at which o_2 is located. Such partitions satisfy the weak converse of MB5. Formally we can express this constraint on mereological structure (CM) as follows:

CM: $L(o_1, z_1) \text{ and } L(o_2, z_2) \text{ and } o_1 \leq o_2 \rightarrow z_1 \subseteq z_2$

A partition satisfying CM is *mereologically monotonic*. This means that it recognizes all the restricted parthood relations obtaining in the pertinent domain of objects. A very simple example is given by a flat list (a partition having only minimal cells together with a root) projected one-for-one upon a collection of disjoint objects.

6.2 Completeness

So far we have allowed partitions to contain empty cells, i.e., cells that do not project onto any object. We now consider partitions which satisfy the constraint that every cell recognizes some object:

CC: $Z(z, A) \rightarrow \exists o : L(o, z)$

We say that partitions that satisfy CC *project completely*. Notice that this condition is independent of the functional or relational character of projection and location. Of particular interest, however, are partitions that project completely and in such a way that projection is a total function (partitions which satisfy both MB3 and CC). An example is a map of the United States representing its constituent states. There are no no-man's lands within the territory projected by such a map and every cell projects uniquely onto just one state.

6.3 Exhaustiveness

So far we have accepted that there may be objects in our target domain that are not located at any cell. This feature of partitions is sometimes not acceptable: governments want *all* their subjects to be located in some cell of their partition of taxable individuals. They want their partitions to satisfy a completeness constraint to the effect that every object in the domain is indeed recognized. In this case we say that location is *complete*. Alternatively we say that the partition *exhausts* its domain. Unfortunately, we cannot use

(*) $o \leq D(A) \rightarrow \exists z : Z(z, A) \text{ and } L(o, z)$

to capture the desired constraint. The tax authorities do not (as of this writing) want to tax the separate molecules of their subjects. Trivially, we have:

$$o \leq_A D(A) \rightarrow \exists z : Z(z, A) \text{ and } L(o, z)$$

but this is much too weak, since it asserts only that every object within a given domain that is recognized by a partition is indeed recognized by that partition. It will in fact be necessary to formulate several restricted forms of exhaustiveness, each one of which will approximate in different ways to the (unrealizable) condition expressed in (*).

One such exhaustiveness condition might utilize a sortal predicate (schema) φ that singles out the kinds of objects our partition is supposed to recognize (for example, in the case of the partition of taxable individual human beings, rather than proper parts of human beings). We now demand that the partition A recognize all of those objects in its domain which satisfy φ:

CE$_\varphi$: $o \leq D(A) \text{ and } \varphi(o) \rightarrow \exists z : Z(z, A) \text{ and } L(o, z)$

Let Δ be some domain and let A be a partition such that $\Delta = D(A)$. Since we can very simply use any predicate to define a partition over any domain – by setting

$$L_A(o, z) \equiv o < \Delta \text{ and } \varphi(o)$$

– we can also think of CE$_\varphi$ as asserting the completeness of one partition *relative to* another. Note that the idea underlying CE$_\varphi$ is closely related to the idea of granularity. Thus for some purposes we might find it useful to formulate condition φ as a restriction on object size.

The tax office probably does not care too much about empty cells in its partition, nor is it bothered too much by the idea of charging you twice. The main issue is to catch everything above a certain resolution at least once. This is the intuition behind constraints like CE$_\varphi$. If you are a law-abiding citizen, you will accept CE$_\varphi$(where 'φ' stands for 'is a citizen'), but you will insist that the partition not locate you in two separate cells, i.e., that you are not charged twice. This means that you want the tax partition to satisfy CE$_\varphi$ and MB4. There might be a pedantic clerk in the tax office who does not rest until he has made sure that all empty cells have been removed. Partitions that will satisfy you, the government, and the clerk in the tax office must satisfy CC, CE$_\phi$, and MB1–5. Projection and location are then total functions (relative to a selected predicate ϕ) and one is the inverse of the other. Under those circumstances projection and location are bijective functions. Notice that neither of the following holds:

(**) if MB4 and CE_ϕ and CC then MB3
(***) if MB3 and CE_ϕ and CC then MB4

Counterexamples are given in Figure 6 (a) and (b), respectively, where each depicted object is assumed to satisfy φ.

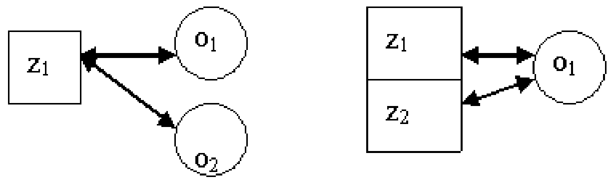

Figure 6: Functionality of projection and location are independent of completeness and exhaustiveness

6.4 Comprehension axioms

The following is the partition-theoretic equivalent of the unrestricted set-theoretic comprehension axiom. For each predicate φ there is a partition $A(\varphi)$ whose location relation is defined as follows:

$$\exists z : L_{A(\varphi)}(o, z) \text{ iff } \varphi(o)$$

Under what conditions on φ can this be allowed?

One type of restriction that is relevant to our purposes would allow φ to be unrestricted but affirm additional restrictions on objects, for example in terms of spatial location. Thus we might define a family of spatial partitions $A(\varphi, r)$, where r is some pre-designated spatial region, in such a way that

$$\exists z : L_{A(\varphi,r)}(o, z) \text{ iff } \varphi(o) \text{ and } o \text{ is spatially located in } r.$$

Something like this is in fact at work in the taxation partition (the tax office is interested in human beings bearing a special relation to a specific geographic location), as also in the partitions used by epidemiologists, ornithologists and others who are interested in (types of) objects at specific sites.

6.5 Redundancy

Partitions are natural cognitive devices, for example they are lists, maps, and so forth, used by human beings to serve various practical purposes. This means that partitions will normally be called upon to avoid certain sorts of redundancy. Here we distinguish what we shall call correspondence redundancy and structural redundancy. Necessarily empty cells (cells whose labels tell us *ex ante* that no objects can be located within them) represent one type of correspondence redundancy, which is excluded by condition CC.

Another type of correspondence redundancy we have addressed already in our discussion of the functionality of location. This occurs in a partition with two distinct cells whose labels would tell us, again *ex ante*, that they must necessarily project upon the very same object. Clearly, and most simply, a partition should not contain two distinct cells with identical labels.

The following case is not quite so trivial. Consider a partition with a cell labeled *vertebrates* which occurs as a subcell of the cell labeled *chordates* in our standard biological classification of the animal kingdom. Almost all chordates are in fact vertebrates. Suppose (for the sake of argument) that biologists were to discover that all chordates must be vertebrates. Then such a discovery would imply that, in order to avoid structural redundancy, they would need to collapse into one cell the two cells (of chordates and vertebrates) which at present occupy distinct levels within their zoological partitions.

A constraint designed to rule out such structural redundancy would be:

CR: A cell in a partition never has exactly one immediate descendant.

This rules out partition-theoretic analogues of the set theorist's $\{\{a\}\}$.

7 FULLNESS AND CUMULATIVENESS

There are a number of different sorts of knowledge shortfall which should be considered in any complete theory of granular partitions. One type arises when there are missing levels within a hierarchy. A partition of the United Kingdom which mentions regions, counties, towns, etc., but leaves out the cells *England*, *Scotland*, *Wales* and *Northern Ireland* is an example of this sort of incompleteness.

More important however are those types of shortfall which have to do with relations between existing levels. We have distinguished thus far *completeness*, which has to do with the absence of empty cells, and *exhaustiveness*, which has to do with the successful capturing of all pertinent objects in a given domain. But shortfalls arise also in relation to a third type of completeness, which has to do with ensuring that successive levels of a partition relate to each other in the most desirable way. We can initially divide this third type of completeness into two sub-types: fullness and cumulativeness. Fullness, intuitively, is a requirement to the effect that each cell z has enough daughter (immediate descendant) cells to fill out z itself. Cumulativeness is a requirement to the effect that these daughter cells are such that the objects onto which they are projected are sufficient to exhaust the domain onto which the mother cell is projected. Fullness, accordingly, pertains to theory (A), cumulativeness to theory (B). (We shall henceforth assume for the sake of simplicity that there are no redundancies in the sense of CR.)

Non-fullness and non-cumulativeness represent two kinds of shortfall in the *knowledge* that is embodied in a partition. Non-fullness is the shortfall which arises when a cell has insufficiently many subcells within a given partition (for instance it has a cell labeled *mammal*, but no subcells corresponding to many of the species of this genus). Non-cumulativeness is the shortfall which arises when our projection relation locates insufficiently many objects in the cells of our partition, for example when I strive to make a list of the people that I met at the party yesterday but leave out all the Welshmen. Fullness and cumulativeness are rarely satisfied by our scientific partitions of the natural world.

They are satisfied primarily by artificial partitions of the sort which are constructed in database environments.

7.1 Fullness

Consider a partition consisting of three cells, labeled *people*, *Laura* and *George W.* Or consider a partition with three cells labeled: *mammals*, *cats* and *dogs*. Both of these partitions are transparent, by our definition (DTr); but both are, again, such as to fall short of a certain sort of ideal completeness, which we can express by asserting that the mereological sum of the cells *Laura* and *George W.* (or of the cells *cats* and *dogs*) falls far short of the corresponding partition-theoretic sum.

If a collection of subsets of some given set forms a partition of this set in the standard mathematical sense, then these subsets are (1) mutually exhaustive and (2) pairwise disjoint. An analogue of condition (2) holds for minimal cells in our present framework, since minimal cells are always mereologically disjoint (they cannot, by definition, have subcells in common). Condition (1) however does not necessarily hold within the framework of partition theory. This is because, even where the partition-theoretic sum of minimal cells is identical to the root cell, the minimal cells still do not necessarily exhaust the partition as a whole. The mereological sum ($+$) of cells is, we will recall, in general smaller than their partition-theoretic sum ($\cup$).

We call a cell *full* relative to its descendant cells within a given partition if these descendants are such that their mereological sum and their partition-theoretic sum coincide. Formally we define:

DFullcell: $\mathit{Fullcell}(z_1) \equiv (+_{z \subset z_1} z) = (\cup_{z \subset z_1} z)$

where $(+_{z \subset z_1} z)$ and $(\cup_{z \subset z_1} z)$ symbolize respectively the operations of applying mereological and partition-theoretic sum for all proper subcells z of the cell z_1. Since we can easily prove that $(\cup_{z \subset z_1} z) = z_1$, DFullcell could be reformulated as asserting that a cell is full if and only if it is identical to the mereological sum of its descendants:

$$\mathit{Fullcell}(z_1) \leftrightarrow +_{z \subset z_1} z = z_1$$

DFullcell does not suffice to capture the intended notion of fullness for partitions however. To see the problem, consider the partition consisting of

top row: one maximal cell, labeled *first couple*
middle row: two intermediate cells, labeled *George W.* and *Laura*
bottom row: four minimal cells, labeled *George W.'s left arm*, *George W.'s right leg*, *Laura's left arm*, *Laura's right leg*.

The cell in the top row satisfies DFullcell, i.e. it is full relative to the second row; but it is not full relative to all of its descendants, since the mereological sum of the cells *George W.'s arm* and *George W.'s leg* is not identical to the cell *George W.* (and analogously for *Laura*). The problem arises because if $x \leq y$ then $x + y = y$ and if $x \subseteq y$ then $x \cup y = y$. From this it follows that only the immediate descendants of a given cell z_1 contribute to its mereological and partition-theoretic sums.

This, however, tells us what we need to take into account in defining what it is for a cell to be full relative to all its descendant cells within a given partition A, namely that each of its constituent cells must be full relative to its immediate descendents. This yields:

DFullcell*: $Fullcell{*}(z_1, A) \equiv \forall z : z \subseteq z_1 \rightarrow (Fullcell(z) \text{ or } Min(z, A))$

Here minimal cells have been handled separately because they do not have subcells. One can see that, while $Fullcell(\mathit{first\ couple})$ holds in the mentioned partition, $Fullcell{*}(\mathit{first\ couple})$ does not, because the cells *George W.* and *Laura* are neither full nor minimal.

We can now define what it means for a *partition* to be full, as follows:

DFull2: $Full(A) \equiv \forall z : Z(z, A) \rightarrow (Fullcell(z) \text{ or } Min(z, A))$

A partition is full if and only if all its non-minimal cells are full (or, equivalently, all its non-minimal cells are full*). Notice that full partitions might in principle contain empty cells, which may or may not have subcells. (Consider, for example, the cell *dodo* with subcells *male dodo* and *female dodo*.)

7.2 Empty space

When a cell falls short of fullness, this means that, while the cell successfully projects onto some given domain, its subcells do not succeed in projecting onto the entirety of this domain. It is then as if there is some extra but invisible component in the cell, in addition to its subcells. We shall call this additional component 'empty space' (noting that the term 'empty' here has a quite different meaning from what it has in the phrase 'empty cell'). Consider the partition depicted in Figure 2.

Here the empty space is that part of the cell *mammals* that is not occupied by *cats* and *dogs*. Notice that this empty space is a component of the cell *mammals* but it is not itself a cell nor is it made up of cells. Empty space is that part of a cell that is not covered by its subcells. The notion of empty space is then similar to the notion of hole in the sense of Casati and Varzi (1994). A hole requires an object which serves as its host. A hole is in every case a hole *in* something. In the same sense empty space requires a cell within which it can exist as empty space. As a hole is a concavity of a host object where no parts of this host object are to be found, so empty space is a zone within a cell where no subcells are to be found.

Empty space within a cell is like a hole also in the sense that there must be something that potentially fills it. In our case this means: more subcells. If all the empty space in a cell is filled, then the cell itself is full and the empty space has been eliminated. Empty space is inert in the sense that it does not project onto anything. Empty space is normally hidden to the user of the partition in which it exists, for otherwise this user would surely have constructed a fuller partition. In some cases however a user might deliberately accept empty space in order to have means to acknowledge the fact that something has been left out. Alternatively, the existence of empty space in a given partition might be brought to the attention of the user. We point in a certain direction and ask: What is *there*? The theory of empty space thereby serves as the starting-point for an ontology of questions (Schuhmann and Smith, 1987): empty space corresponds to a hole in our knowledge.

7.3 Empty space and knowledge

The presence or absence of empty space is a dimension of a granular partition that is skew to the dimension pertaining to the existence of empty cells. An empty cell is a cell that fails to project. Empty space is that which leaves room for the addition of new knowledge. It is a zone within a cell that is not (but given advances in knowledge, could be) occupied by further subcells. Notice, however, that there cannot be a partition that is full and yet consists entirely of empty cells. This is because by MB6 every partition projects onto some object.

Figure 2 (a) depicts a partition of the animal kingdom consisting of three cells, where z_3 recognizes the animal kingdom as a whole, z_1 recognizes dogs, and z_2 recognizes cats. In terms of partition-theoretic union we have $z_1 \cup z_2 = z_3$, but clearly $p(z_1) + p(z_2) < p(z_3)$. New cells can be inserted into the partition if new species are discovered (e.g., the species indicated by o_3).

Empty space from one point of view reflects the potential for adding new knowledge. From another point of view it can be seen as a matter of *hidden* knowledge. From this perspective it is as if we start from partitions which reflect the way God sees reality. Empty space then covers up what we humans do not yet know. The latter view was developed by Mislove et al. (1990) and resulted in the theory of partial sets. What we call empty space Mislove compares to packaging material. If we remove packaging material we potentially discover new things that were previously hidden from view. Our distinction between full and non-full partitions corresponds very closely to the distinction between 'clear' and 'murky' sets drawn in the theory of partial sets.

7.4 Fullness and emptiness

Since we have $\mathit{Fullcell}(z_1) \leftrightarrow (+_{z \subset z_1} z) = z_1$, we also have $\neg \mathit{Fullcell}(z_1) \leftrightarrow (+_{z \subset z_1} z) < z_1$. Consequently we can define what it means for x to be the empty space of the cell z_1 as follows. We first of all define x fills z_1:

DFills: $\mathit{Fills}(x, z_1) \equiv \neg \mathit{Fullcell}(z_1)$ and $x + (+_{z \subset z_1} z) = z_1$

The empty space in z_1 is then its smallest filler and we define that x is the empty space in z_1 if and only if x fills the space not occupied by the subcells of z_1 and x is disjoint from all subcells of z_1:

DES: $ES(x, z_1) \equiv \mathit{Fills}(x, z_1)$ and $\forall y : \mathit{Fills}(y, z_1) \rightarrow x \leq y$

We note in passing that minimal cells, on the basis of the definitions above, are either empty (they do not project) or they are completely made up of empty space. ('Minimal' means: there is no further knowledge available, within a given partition, as concerns the objects onto which minimal cells are projected.)

ES determines the empty space of a cell uniquely. To prove this, assume $ES(x, z)$ and $ES(y, z)$. For simplicity we use 'c' as an abbreviation for '$(+_{z_1 \subset z} z_1)$'. We can then rewrite $ES(x, z)$ as: $x + c = z$ and $\forall u : u + c = z \rightarrow x \leq u$, and rewrite $ES(y, z)$ as: $y + c = z$ and $\forall v : v + c = z \rightarrow y \leq v$. Since $z = u + c = v + c = x + c = y + c$ we have for all u, v: if $v + c = u + c = x + c = y + c$ then ($x \leq u$ and $y \leq v$). In particular this holds in the case of $\forall u, v : u = v$. Now assume that $x \neq y$. This means that for all u: if $u + c = x + c = y + c$ then $(x < u$ and $y \leq u)$ or $(x \leq u$ and $y < u)$. Without loss of generality consider for all u: if $u + c = x + c = y + c$ then $(x < u$ and $y \leq u)$. Now without loss of generality consider for all u: if $u + c = x + c$ then $x < u$. Since for all u: $u + c = x + c$ holds by assumption it in particular this holds for $u = x$. Consequently we have $\exists u : u + c = x + c$ and $u = x$ and $u < x$ leading to a contradiction of the form $\exists u, x : u = x$ and $u \neq x$. Consequently we have $x = y$ as desired.

7.5 Cumulativeness

We can now define the notion of *cumulativeness*, which plays the same role in theory (B) which fullness plays in theory (A). The intuitive idea is as follows: a cell is cumulative relative to its immediate descendant cells if the mereological sum of the *projections* of these immediate descendants is identical to the projection of their partition-theoretic sum. For non-empty and non-minimal cells with at least two immediate descendants we define:

DCu1: $Cu(z_1) \equiv +_{z \subset z_1} p(z) = p(\cup_{z \subset z_1} z)$

One can see that $p(\cup_{z \subset z_1} z) = p(z_1)$ holds under the given conditions. Consequently:

$$Cu(z_1) \leftrightarrow +_{z \subset z_1} p(z) = p(z_1)$$

Again, $Cu(z_1)$ ensures that z_1 is cumulative relative to its immediate descendants. In order to ensure cumulativeness of a cell with respect to all its subcells, we define

DCu2: $Cu{*}(z_1, A) \equiv \forall z : z \subseteq z_1 \rightarrow (Cu(z) \text{ or } Min(z, A))$

A partition is cumulative if and only if all its cells are cumulative.

DCu3: $Cu(A) \equiv \forall z : Z(z, A) \rightarrow (Cu(z) \text{ or } Min(z, A))$

Another way of expressing this is as follows: a partition is cumulative if and only if it has a *basis in objects* (the objects projected by its minimal cells), and is then built up in stages in such a way that each non-minimal cell z projects onto the mereological sum of the objects projected by z's immediate descendants.

Recall that the notions of fullness and cumulativeness are intended to characterize partitions that have no redundancies of the sort defined in CR. A cumulative partition A is also exhaustive (CE$_\varphi$) with $\varphi = \exists z : Z(z, A)$ and $P(z, o)$.

7.6 Classes of partitions regarding fullness and cumulativeness

We can now distinguish four classes of partitions regarding their fullness and cumulativeness:

Full and cumulative. Consider a list of the 50 US States, divided into two sub-lists: the contiguous 48, the non-contiguous 2. Here the objects towards which this partition is directed are the States themselves under the obvious 'Utah'–Utah projection relation.

Full and non-cumulative. Consider Figure 7 (a). This partition represents the belief of some child who thinks that cats and dogs are the only animals there are.

Non-full and cumulative. Consider Figure 7 (b). This partition represents the way a child sees the world who does not understand the concepts of the Northern and Southern hemisphere and who thinks that there are places on Earth that are neither in the Northern, nor in the Southern Hemisphere, nor overlapping both, but are rather in some secret and wonderful land that has not yet been discovered.

Non-full and non-cumulative. Consider Figure 7 (c). Imagine that you have a terrible hangover, and your accounting of the people at the party last night consists of a root with three subcells: *you*, *John*, and *Mary*. You know that you are missing somebody, but you cannot remember who.

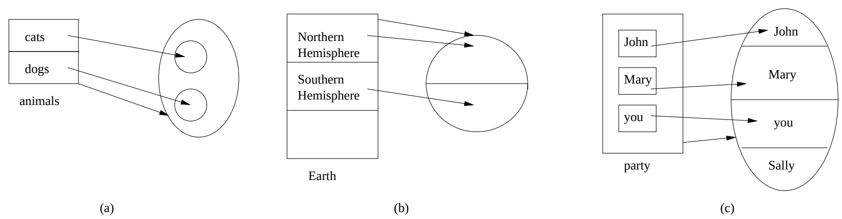

Figure 7: Examples of non-full and non-cumulative partitions.

We humans are often aware of the fact that our partitions are not full or not cumulative, or both. In the former case empty cells are often included into the cell structure in order to make explicit our awareness of these shortcomings. Consider, for example, the periodic table of the elements, which contains empty cells (labeled ununnilium, etc.), for elements not yet synthesized. Moreover the fact that the periodic table in its current form has a structure that can only accommodate a certain number of cells is not taken to imply that in reality there could not be elements whose discovery would require the table to be extended.

8 IDENTITY OF GRANULAR PARTITIONS

As a step towards a definition of identity for partitions Smith and Brogaard (2002b) propose a partial ordering relation between partitions, which they define as follows:

$$A \leq B \equiv \forall z : Z(z, A) \rightarrow Z(z, B)$$

They then define an equivalence relation on partitions as follows:

DE: $A \approx B \equiv A \leq B$ and $B \leq A$

Now, however, we can see that a definition along these lines will work only for partitions which are full. What, then, of those partitions which are equivalent in the sense of DE but not full? What are the relationships between the presence of empty and redundant cells and the question of the identity of partitions? And what is the bearing on the question of identity of the phenomenon of empty space? Can partitions that have empty or redundant cells be identical? Can partitions which are not full be identical?

The question of whether or not partitions that have empty or redundant cells are identical cannot be answered without a theory of labeling. Corresponding empty cells in two distinct partitions, if they are to be considered as identical, need to have at least the same labels. We can only address this question informally here, leaving for another place the task of developing the necessary formal theory.

As already noted, one makes a different sort of error if one thinks that there are dodos from the error which one makes if one thinks that there is an intra-Mercurial planet. The two corresponding empty cells are thus distinct: they are signs of different sorts of failure arising when we project a partition onto reality.

Consider an inventory of the goods you plan to sell in your not-yet-established chain of beer-brewing stores. This is, surely, different from the inventory of the goods I plan to sell in my not-yet-established chain of wine-marketing stores, And this is so even in the event that our respective plans are never realized.

Consider the partition of the people in your building according to *number of days spent behind bars*. You can construct this partition—which is little more than a simple array of numbered boxes—prior to undertaking any actual inquiries as to who, among the people in your building, might be located in its various cells. Thus even before carrying out such inquiries you can know that this is a more refined partition than, for example, the partition of the same group of people according to *number of years spent behind bars*. The two partitions are distinct, and they will remain distinct even if it should turn out that none of the people in your building has spent any time at all in jail. In both cases all the people in your building would then be located in the cell labeled *zero* and all the other cells in both partitions would be empty. Yet the two partitions would be nonetheless distinct, not least because their respective maximal cells would have different labels.

We can now return briefly to our question whether partitions that are neither full nor cumulative can be said to be identical. One approach to providing an answer to this question would be to point out that, even though two partitions are outwardly identical, they might still be such that there are different ways to fill the corresponding empty space. Suppose we have what are outwardly the same biological taxonomies used by scientists in America and in Australia at some given time, both with the same arrays of empty cells. Suppose these partitions are used in different ways on the two continents, so that, in the course of time, their respective empty space gets filled in different ways. Were they still the *same* taxonomy at the start?

9 RELATED WORK

Variant forms of our granular partitions have been advanced elsewhere, in particular in the literature on Spatial Information Science. What makes spatial or geographic granular partitions particularly interesting is the fact they are very well structured in the sense that they not only obey our master conditions MA1–4 and MB1–6 but are also mereologically monotonic (CM). In addition they are often full and cumulative as well as exhaustive with respect to predicates such as land use, political affiliation, and so on.

Examples of geographic granular partitions are categorical coverages (Chrisman, 1982), which are thematic maps depicting the relationship of a property or attribute to a specific geographical area. A prototypical example of a categorical coverage is the land use map, in which a taxonomy of land use classes is determined (e.g., residential, commercial, industrial, transportation) and some specific area (for example a city) is then evaluated along the values of this taxonomy (Volta and Egenhofer, 1993). Another prototypical example is soil maps, which are based on a classification of soil covering (into *clay*, *silt*, *sand*, etc.). The zones of a categorical coverage are a jointly exhaustive and pair-wise disjoint subdivision of the relevant region of space (Beard, 1988). As discussed in (Bittner and Smith, 2001) categorical coverages can be understood in terms of the interplay of two granular partition: one targeting a region of space, the other targeting some attribute domain.

Another important feature of spatial granular partition is that they also recognize topological and geometric properties of the domain they project onto. In order to recognize structure beyond mereology the cell structure of a granular partition must have structural features in addition to the subcell relation between its cells. Examples of granular partitions that take topological structure, i.e. neighborhood relations between adjacent partition cells, into account can be found in Frank and Kuhn (1986); Bittner and Stell (1998); Erwig and Schneider (1997). Applications of granular partitions taking the ordering of the cell structure and the shape of cells into account were discussed in Frank (1992); Freksa (1992); Hernandez (1991).

Recently, spatial partitions were applied also to the representation of spatio-temporal phenomena (Erwig and Schneider, 1999) and temporal phenomena (Bittner, 2002). Independently, the notion of granularity was discussed in Stell (1999, 2000) and Stell, this volume.

10 CONCLUSIONS

This paper is a contribution to a formal theory of granular partitions. We defined master conditions that need to be satisfied by every partition. These master conditions fall into two groups: (A) master conditions characterizing partitions as systems of cells, and (B) master conditions describing partitions in their projective relation to reality.

At the level of theory (A) partitions are systems of cells that are partially ordered by the subcell relation. Such systems of cells are such that they can always be represented as trees, they have a unique maximal element and they do not have cycles in their graph-theoretic representations. But partitions are more than just systems of cells. They are also cognitive devices that are directed towards reality.

Theory (B) takes this latter feature into account by characterizing partitions in terms of the relations of projection and location. Cells in partitions are projected onto objects in reality. Objects are located at cells when projection succeeds. We then say that a partition *recognizes* the objects that are located at its cells. To talk of granular partitions is to draw attention to the fact that partitions are in every case selective; even when they recognize some objects, they will always trace over others.

Partitions do not only recognize objects, they are also capable of reflecting the mereological structure of the objects they recognize through a corresponding mereological structure on the side of their cell array. This does not mean, however, that all partitions actually do reflect the mereological structure of the objects they recognize. For it is an important feature of partitions that they are also capable of tracing over mereological structure. There are, for example, large classes of partitions that simply list objects, without caring at all how these objects hang together mereologically.

Our discussion of granularity showed that partitions have three ways of tracing over mereological structure: (1) by tracing over mereological relations between the objects which they recognize; (2) by tracing over the parts of such objects; (3) by tracing over the wholes which such objects form. The tracing over of parts is (unless mereological atomism is true) a feature manifested by every partition, for partitions are in every case *coarse grained*. The tracing over of wholes reflects the property of granular partitions of foregrounding selected objects of interest within the domain onto which they are projected and of leaving all other objects in the background where they fall in the domain of unconcern.

ACKNOWLEDGMENTS

Our thanks go to Berit Brogaard, Maureen Donnelly, Pierre Grenon, Jonathan Simon, and John Stell for helpful comments. Support from the American Philosophical Society, and from the NSF (Research Grant BCS-9975557: "Geographic Categories: An Ontological Investigation") is gratefully acknowledged. This work was also supported in part by DARPA under the Command Post of the Future Program, by the NSF Research on Learning and Education Program, and by the Wolfgang Paul Program of the Alexander von Humboldt Foundation.

REFERENCES

Beard, K. (1988). *Multiple representations from a detailed database: A scheme for automated generalization*. PhD thesis, University of Wisconsin, Madison.

Bittner, T. (2002). Approximate qualitative temporal reasoning. *Annals of Mathematics and Artificial Intelligence*. To appear.

Bittner, T. and Smith, B. (2001). A taxonomy of granular partitions. In Montello, D. R. (Ed), *Spatial Information Theory: Foundations of Geographic Information Science*, volume 2205 of *Lecture Notes in Computer Science*. Berlin: Springer-Verlag.

Bittner, T. and Stell, J. G. (1998). A boundary-sensitive approach to qualitative location. *Annals of Mathematics and Artificial Intelligence*, 24:93–114.

Casati, R. and Varzi, A. C. (1994). *Holes and other Superfacilities*. Cambridge, MA: MIT Press.

Casati, R. and Varzi, A. C. (1995). The structure of spatial location. *Philosophical Studies*, 82:205–239.

Chrisman, N. (1982). *Models of Spatial Analysis Based on Error in Categorical Maps*. PhD thesis, University of Bristol, England.

Degen, W., Heller, B., Herre, H., and Smith, B. (2001). GOL: A general ontological language. In Welty, C. and Smith, B. (Eds), *Formal Ontology and Information Systems*, pp. 34–46. New York: ACM Press.

Erwig, M. and Schneider, M. (1997). Partition and conquer. In Hirtle, S. C. and Frank, A. U. (Eds), *Spatial Information Theory: A Theoretical Basis for GIS*, volume 1329 of *Lecture Notes in Computer Science*. Berlin: Springer-Verlag.

Erwig, M. and Schneider, M. (1999). The honeycomb model of spatio-temporal partitions. In *International Workshop on Spatio-Temporal Database Management*, volume 1678 of *Lecture Notes in Computer Science*, pp. 39–59. Berlin: Spinger-Verlag.

Frank, A. (1992). Qualitative spatial reasoning about distances and directions in geographic space. *Journal of Visual Languages and Computing*, 3:343–371.

Frank, A. and Kuhn, W. (1986). Cell graphs: A provable correct method for the storage of geometry. In Marble, D. (Ed), *Second International Symposium on Spatial Data Handling*, pp. 411–436, Seattle, WA.

Frank, A., Volta, G., and McGranaghan, M. (1997). Formalization of families of categorical coverages. *International Journal of Geographic Information Science*, 11(3):214–231.

Freksa, C. (1992). Using orientation information for qualitative spatial reasoning. In Frank, A. U., Campari, I., and Formentini, U. (Eds), *Theories and Methods of Spatio-Temporal Reasoning in Geographic Space*, volume 639 of *Lecture Notes in Computer Science*. Berlin: Springer-Verlag.

Galton, A. (1999). The mereotopology of discrete space. In Freksa, C. and Mark, D. M. (Eds), *Spatial Information Theory: Cognitive and Computational Foundations of Geographic Science*, volume 1661 of *Lecture Notes in Computer Science*, pp. 251–266. Berlin: Springer-Verlag.

Hernandez, D. (1991). Relative representation of spatial knowledge: The 2D-case. In Mark, D. M. and Frank, A. U. (Eds), *Cognitive and Linguistic Aspects of Geographic Space*, pp. 373–386. Dordrecht, The Netherlands: Kluwer Academic Publishers.

Lewis, D. (1991). *Parts of Classes*. Oxford: Blackwell.

Mark, D. M. (1978). Topological properties of geographic surfaces: Applications in computer cartography. In *Harvard Papers on Geographic Information Systems*, volume 5. Cambridge, MA: Laboratory for Computer Graphics and Spatial Analysis, Harvard University.

Mislove, M. W., Moss, L., and Oles, F. J. (1990). Partial sets. In Cooper, R., Mukai, K., and Perry, J. (Eds), *Situation Theory and Its Applications I*, volume 22 of *CSLI Lecture Notes*, pp. 117–131. Stanford, CA: Center for the Study of Language and Information.

Schuhmann, K. and Smith, B. (1987). Questions: An essay in Daubertian phenomenology. *Philosophy and Phenomenological Research*, 47:353–384.

Searle, J. R. (1983). *Intentionality. An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.

Smith, B. (1991). Relevance, relatedness and restricted set theory. In Schurz, G. and Dorn, G. J. W. (Eds), *Advances in Scientific Philosophy. Essays in Honour of Paul Weingartner*, pp. 45–56. Amsterdam/Atlanta: Rodopi.

Smith, B. (1998). The basic tools of formal ontology. In Guarino, N. (Ed), *Formal Ontology in Information Systems*, Frontiers in Artificial Intelligence and Applications, pp. 19–28. Amsterdam: IOS Press.

Smith, B. (1999). Truthmaker realism. *Australasian Journal of Philosophy*, 77(3):274–291.

Smith, B. (2001a). Fiat objects. *Topoi*, 20(2):131–148.

Smith, B. (2001b). True grid. In Montello, D. R. (Ed), *Spatial Information Theory: Foundations of Geographic Information Science*, volume 2205 of *Lecture Notes in Computer Science*, pp. 14–27. Berlin: Springer-Verlag.

Smith, B. and Brogaard, B. (2002a). Quantum mereotopology. *Annals of Mathematics and Artificial Intelligence*, 34:153–175.

Smith, B. and Brogaard, B. (2002b). A unified theory of truth and reference. *Logique et Analyse*. To appear.

Stell, J. G. (1999). Granulation for graphs. In Freksa, C. and Mark, D. M. (Eds), *Spatial Information Theory. Cognitive and Computational Foundations of Geographic Information Science*, number 1661 in Lecture Notes in Computer Science, pp. 417–432. Berlin: Springer-Verlag.

Stell, J. G. (2000). The representation of discrete multi-resolution spatial knowledge. In *Proceedings of Seventh International Conference on Principles of Knowledge Representation and Reasoning (KR2000)*, San Francisco. Morgan Kaufmann Publishers.

Volta, G. and Egenhofer, M. (1993). Interaction with GIS attribute data based on categorical coverages. In Frank, A. U. and Campari, I. (Eds), *Spatial Information Theory: A Theoretical Basis for GIS*, volume 716 of *Lecture Notes in Computer Science*. Berlin: Springer-Verlag.

Wittgenstein, L. (1961). *Tractatus Logico-Philosophicus*. London: Routledge and Kegan Paul. Translated by McGuinness, M. F. and Pears, D.