SFB/TR 8: I1-[OntoSpace]

OntoSpace Project Report
University of Bremen
Germany

Spatial Ontology Baseline

Deliverable D2
I1-[OntoSpace]; Workpackage 1

John Bateman and Scott Farrar

Version: 2.0

June 2006

http://www.sfbtr8.uni-bremen.de/i1
e-mail: {bateman,farrar}@uni-bremen.de

Roadmap for baseline deliverables D1-D4

The baseline position on ontology, ontology construction, and ontology use adopted in project I1-[OntoSpace] is set out in a sequence of four deliverables (D1-D4). Each provides an introduction to the respective states of the art and describes the positions within these that I1-[OntoSpace] is adopting for its own work or as proposals for ontology construction within the SFB/TR Ontology Working Group generally.

The baseline is made up of the following components:

- D1 Ontology as such and the principal approaches and methodologies currently available for general ontology construction;
- D2 The ontologies of space: approaches to representing space that have been taken on ontology and qualitative spatial representation and reasoning;
- D3 The ontologies motivated by and for language: approaches to representing the kinds of distinctions that treatments of natural language require—particularly but not exclusively those required for spatial language;
- D4 Inter-Ontology mappings and structuring devices: approaches to constructing ontologies out of submodules and of relating such submodules in systematic ways.

This is the second of these deliverables and introduces the spatial ontology baseline.

Deliverables D1, D2 and D3 are results of Workpackage 1 as described in the I1-OntoSpace project proposal; D4 is a result of Workpackage 3 and cooperation with project I4-[SPIN].

In general, we will describe deliverables either by the long form 'I1-[OntoSpace]:D2' or, when there is no need for disambiguation, the short form 'D2'.

Note: We maintain an extensive and regularly updated webportal for our ontology activities as well as pointers to all kinds of ontologies at the Bremen Ontology Research Group website: `http://www.fb10.uni-bremen.de/ontology`

Abstract

This document discusses the various approaches to representing space that have been taken in ontology and qualitative spatial representation and reasoning. The parameters concerning the ontological modelling of space, both in a general sense and concerning the SFB specifically, are given. The specific projects discussed include: SUMO-space, OpenCyc-space, DOLCE-space, BFO-space, and Geographic Information Systems (GIS). Also addressed are spatial calculi including RCC and other more recent region-based proposals. The discussion is summarized in terms of recommendations for practice and development within the SFB when embodying spatial representations within ontologies.

Acknowledgements

The Collaborative Research Center for Spatial Cognition (Sonderforschungsbereich/Transregio SFB/TR8) of the Universities of Bremen and Freiburg is funded by the Deutsche Forschungsgemeinschaft (DFG), whose support we gratefully acknowledge. We also acknowledge useful comments on earlier drafts from all members of the SFB/TR8 and also particularly from Stefano Borgo and Claudio Masolo (LOA). Mistakes and diverging interpretations remain our own responsibility as ever.

Contents

List of Figures

List of Tables

1 Introduction

This document is the second deliverable (D2) for Workpackage 1 as described in the I1-OntoSpace project proposal. Its function is to provide a detailed view of space, spatial representation, location in space and movement in space from the perspective of ontology. This view draws on the existing states of the art in all of these areas, as well as drawing in particularly relevant work that has attempted to apply ontology to practical concerns—particularly, cognitive modelling, artificial intelligence, geographic information systems (GIS), and the like. The role of an ontology in this area is as it is in all domains and as was introduced in our baseline ontology deliverable D1: that is, (i) to set out a consistent and well-specified general modelling scheme which is free of contradiction and from which follows a set of generic properties that necessarily hold over the entities covered and, (ii) to support problem solving and inference within the domain of concern.

Given the increasing need for accounts of space and a similarly growing awareness of the potential application of ontological methods, it is not surprising that the task pursued here is similar in several respects to some other currently ongoing efforts or proposals. For example, on the one hand there have been proposals that a repository of 'spatial representations and methods' be set up as a module among generic ontological resources—first as part of the Darpa Agent Modelling Language, **DAML**: www.daml.org, and now more recently as **OWL-space** (i.e., an addition for spatial ontology within the Web Ontology Language, OWL)—while, on the other hand, long time proponents of ontology methods for Geographic Information Systems (GIS), such as Frank (2003*b*) and Kuhn (2001), have begun attempting to draw more comprehensive ontologies containing and relating information about distinct kinds of spatial objects and their relationships into the geographical arena. Reasoning and problem-solving in space also has a long tradition in AI and computational cognitive modelling, where some of the most detailed accounts have been developed (cf., e.g., Kuipers 1977, Kuipers 1978, Kuipers 1998).

Across all of these efforts lies a common interest that the representations thus constructed should be useful for problem solving, spatial inference, spatial data retrieval, spatial visualization, consistency checking and maintainance, navigation in space, way-finding in general, explaining human spatial abilities and much more. All of these require highly explicit and spatially appropriate specification: which is what a suitable ontology is required to provide.

Our account here includes a discussion of the various parameters concerning the ontological modelling of space, both in a general sense and, we hope, beginning to move quite specifically to consider the diverse needs of researchers within the SFB. We see the ontological modelling of space as particularly necessary within the SFB for facilitating qualitative spatial reasoning in general, for achieving interoperability over the different spatial calculi used within the participating projects, and for ontologically grounding the spatial expressions found in natural language. Thus, our final aim would be that all of the accounts of space pursued within the SFB could be locatable within the spatial ontology that we begin development of here. Since a necessary (but not sufficient) precondition for achieving this goal is extensive feedback and participation from the other projects of the SFB where the explicit representation of space is an issue, this current deliverable must still in some sense be classified as 'preliminary'. We hope that the directions pursued here will continue to prompt feedback which may well lead us to revise our explorations substantially.

In order to do justice to the diverse issues that are raised with respect to reasoning about, acting in, and communicating about space, we must naturally adopt a rather broad basis for our account—one which does not focus only on 'space as such' or on spatial issues or on particular directions of formalization of spatial issues or on particular kinds of objects in space. This is the only way in which we can approach our own goals of employing ontological methods to achieve integration, reconciliation and mediation between different approaches and traditions. This does mean, however, that we are not able to go into equal depth in all areas, nor are we able to do justice to some of the specialized results and methods that have been achieved in those areas: what we will attempt is to state how these diverse approaches relate to one another—how they may all form part of a general picture of spatial concerns. For this general picture we draw on ontology in its role of providing a structured foundation for more specific accounts.

One consequence of our approach is that we will be concerned with detailed ontological considerations of several domains. A criticism we will make of several proposals that have attempted to provide more ontological content to their treatments of space is that they base themselves on foundations which are too narrow. This can have a variety of undesirable consequences, however. One source of examples of this is in the use sometimes made of 'linguistic evidence' in ontological discussions—we will argue that some of these uses simply move a difficult problem in the ontology of objects and events to an equally difficult problem in linguistic ontology; then, since the foundation for these efforts does not include an equally detailed account of linguistic ontology, the problems are not recognized and the result is presented instead as a 'solution'. We will suggest that the solution can be illusory if the linguistic ontological questions have not also been adequately addressed. We therefore need to combine the best of both areas: distributing the theoretical work to be done across these ontological areas is indeed a potentially beneficial step, but it must be done in a way that respects the ontological requirements of both. Since we devote an entire deliverable to the questions and problems of linguistic ontology (deliverable D3), we will postpone most of this discussion until then, noting in passing potential problems in this deliverable as they arise.

A further source of examples is given by formalizations that pick one particular ontological realm and attempt to build the rest of ontology out of this: i.e., various flavours of reductionism. This might be a physical reductionist viewpoint, where all that is considered ontologically real is quantum physics and all else must be derived from this, or it might equally be a reductionist view in terms of spatial-temporal configurations, where all that really exists are chunks of space-time and the rest of the world has to be constructed from this. We consider such approaches to be flawed both ontologically and cognitively. Ontologically we adopt the so-called 'perspectivalist' approach of Smith and colleagues (e.g., Smith & Grenon 2004) described in our deliverable D1 which considers objects, events and locations such as 'cups', 'robot movements' and 'behind the church' to be as real as the quantum flows revealed by suitable artificially constructed experimental observation. Cognitively, there is little doubt that the kinds of everyday objects, events and places found in commonsense views of the world play an important role in all aspects of cognitive behavior and need to be taken seriously in our models of that behavior. Our description of space and spatial objects and events will therefore mostly draw from attempts to provide models of everyday commonsense reality; this will necessarily, as we shall see, involve aspects of qualitative reasoning and representation as well as traditional ontology.

The focus on spatial issues pursued here therefore requires that we include in our discussion approaches that were not discussed in detail in the general ontology baseline of D1. These involve both theoretical approaches and more practically-directed developments aiming at application. This confluence of input is represented graphically in Figure 1. We see here that our main starting point is given by ontology and ontological methodology—particularly methods for achieving ontological consistency such as Guarino & Welty's (2004) OntoClean method described in D1—supplemented by approaches within qualitative spatial representation and its formalization (cf. Cohn & Hazariki 2001) and concrete applications from GIS and cognitive modelling for the spatial domain. Naturally these concerns, and the approaches, overlap to some degree already; we now attempt to increase this overlap by means of the adopted ontological foundation. In many respect, however, our position is one of 'interested consumers'. We are seeking to define a very broad ontological foundation for working with space and to do this we are well aware that considerable expertise and experience must be 'imported'. Our main contribution is then to examine critically from an ontological perspective how the many diverse approaches can be reconciled and placed in correspondence with the state of the art in ontology design in general.

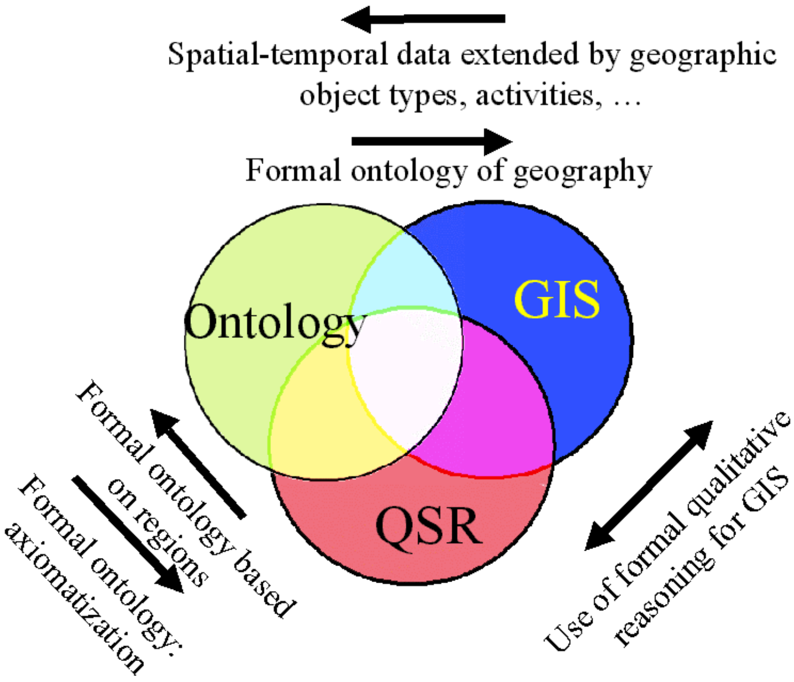

Figure 1: Inputs to this deliverable: Ontology, QSR and GIS

For our discussion in this deliverable, we will assume that the reader has a basic familiarity with the general principles of ontology design independently of their application to space and including basic notions such as granularity, 3D vs. 4D modelling, mereology and so on; these were described in some detail in our deliverable D1 on the Ontology Baseline for the SFB. We refer to several of the ontologies that we introduced generally in D1, picking out their particular statements concerning space and spatial location. For the discussion of space, leading up to our summary and conclusions, we attempt to make the deliverable as self-contained as possible—all of the aspects referred to in the conclusion are introduced in the relevant discussions beforehand.

The structure of the deliverable is then as follows. First we set out broad orientating guidelines for the questions that arise when considering 'space', outlining those approaches to space that are found in the ontological literature. Then we explore in detail some of the explicit approaches to space that can be found in the main ontologies currently available or under development, together with approaches overlapping with, or relevant to, ontology coming from Qualitative Spatial Reasoning and Geographic Information Systems. Finally we summarize and make preliminary recommendations for our own construction of an ontology of space.

2 Space and ontology

In this section, we will focus in our discussion on concrete positions that represent the current state of the art in explorations of space from an ontological perspective, without following these issues to their philosophical roots. The major topics to be discussed include: the question of 'space' as an ontological category, the modelling of concrete spatial objects, (mereo)topology, dimensionality, orientation, shape, spatial relations, paths, and vagueness.

Although the most common approach to computational models of space involves a simple rendition of Euclidean geometry, with space represented as three (or other) dimensional coordinates, there is considerable doubt as to whether this provides a useful style of representation for capturing how humans deal with space. Analogously to concerns raised in the introduction, this doubt applies both cognitively and ontologically.

The kinds of descriptions that we will see running through all of the positions that we describe in this document are accordingly essentially **qualitative**. The basic ontological categories of things and processes are defined by their essential features and interelationships not by physical quantities. Similarly, the relations and constructs involved in spatial reasoning are qualitative not quantitative. In qualitative representation and reasoning with respect to such representations, a situation is characterized by variables which 'can only take a small, predetermined number of values' (de Kleer & Brown 1985, p116) and the inference rules use these values and not numerical quantities approximating them. The nature of these variables and their possible values can then be of ontological import. The accounts of information admitted into our base ontology should be seen as working hypotheses about how humans construct their experience of the spatial world and in which terms.

Our characterization in this section sets out some basic parameters concerning the ontological modelling of space. We begin by considering different perspectives on the nature of space as such that are commonly taken ontologically. We then show that we can consider the minimal formal requirements of an ontology involving space in terms of a sequence of relations. Each relation adds further expressivity to an account and allows a finer, more 'natural' modelling of commonsense views of space and objects. The particular relations adopted, and in which order, is however itself an issue of debate, as each modelling decision brings with it particular consequences. We build on this throughout the rest of this section, surveying the positions taken up and the most central constructs employed.

2.1 Space as an ontological category

The most basic ontological question of traditional philosophical importance is whether space exists independently of any objects that happen to have locations within space or, alternatively, whether space is mainly a matter of inter-relationships between objects. The first view is termed the *Newtonian* or *Galilean* view of space, and the second the *Leibnizian* view (Borgo, Guarino & Masolo 1996*a*, p2); these distinct views have also been termed **substantival**, or **absolutist**, and **relational** definitions of space (cf. Grenon 2003), and **general space** and **local space** by Vieu (1997, p9).

Despite the philosophical-sounding tone of the debate, it in fact has immediate implications for how to explicitly model space in a representation, how space might be used for inference, and how it may be talked about during communication. When building an ontology under the Newtonian approach, for example, space may be modelled directly as a category in that ontology. It then enters into a range of relationships with other entities and should be axiomatized as shown for categories in general in Deliverable D1. In contrast, this need not be the case in a Leibnizian ontology, where space is only present indirectly as relations between objects themselves (see Borgo et al. 1996*a*).

It may also be the case, however, that distinct ontological realms might require differing decisions to be made on the absolutist/relativist position. Aurnague & Vieu (1993), for example, argue that for the purposes of linguistic semantics it is the relativist view that is compelling; although Borgo et al. (1996*a*) favor the absolutist position for "practical applications" involving spatial inference. We can also relate this to a distinction commonly made in the cognitive modelling of space concerning the 'scale' of the spatial framework under investigation. When an agent is interacting within a 'small-scale' environment, such as a room, then it may be plausible to see space in terms of the relationships between the entities within and making up the room; this becomes much less likely however when the agent is considering a large-scale environment, such as a town. This is because for the latter case there is no single point from which the entire environment can be considered; there is a need for a general all-encompasing representation. Particularly for this latter kind of characterization, then, an absolutist view of space may be more useful and it is here that we find most proposals for the adoption of something like a **cognitive map** for integrating information (Kuipers 1978, Kuipers & Levitt 1988, Kuipers 1998).

If space is modelled according to the Newtonian tradition, that is as an ontological category in its own right, then there are still at least two very different ways in which we might proceed. First, there is the traditional geometric view, that considers space as a collection of dimensionless *points*. On the other hand, space can be considered as a system of 'pointless' connected 1-D, 2-D, 3-D, n-D **regions**, as in the Region Connection Calculus (RCC: cf. Section 7 and Randell, Cui & Cohn (1992)) and Borgo, Guarino & Masolo (1997). In this latter case, the definition of regions becomes an important issue in its own right. It has now been shown that geometry can be defined entirely in terms of mereology and topology (see below), and so it is not the case that 'points' are necessarily present as a primitive. This appeals to those who argue that points cannot be primitive as far as a cognitively-adequate representation is concerned because points as such are never perceived. It is therefore points that are the abstraction and the regions out of which perception and our modelling of the world proceeds should by rights be made primitive. As in line with the cognitive (and some would argue, ontological) preference for qualitative representations, it is clear which possibilities will receive the most attention below; in the subsection following we will consider some of the basic tools for this.

A further possibility is then to consider space and its regions as not only a category in the ontology, but as *the* category in that physical objects might be defined as simply whatever is in some particular spatial location. This brings with it certain ontological (and practical problems) that can illustrated well for GIS by, for example, Donnelly & Smith (2003) with the following kind of consideration:

> "How this region-based framework leads to problems becomes clear when we need to formulate a qualitative theory of motion. If we are able even to attempt to characterize movement, something more, for example temporal indexing, must be added to the mereotopology [see below]. But even then, regions-plus-attributes representations of organisms in their habitats must necessarily obscure what is involved when an enduring object is registered at different locations in successive instants of time. For an adequate account of such registration data requires at least two independent sorts of spatial entities: one, the locations, which remain fixed, and the other, the objects, which move relative to them. Since a region-based approach admits only the first type of entity—the location of regions—it must somehow simulate motion, for example via successive assignments of attributes to a fixed frame of reference." (Donnelly & Smith 2003)

This means that, again following Donnelly and Smith's example, instead of talking of particular birds flying around, one must talk instead of mappings of the form:

$$\text{Sparrow}_{152} : \text{time} \rightarrow \text{regular closed subsets of } \mathbf{R}^3$$

One therefore has access only to 'sparrow-shaped regions' which may take on particular attributes and which may vary over time. While it is possible to take such a view, and it may even be *implementationally appropriate* for some particular purposes, it is clear that it finds no particular support cognitively and is, indeed, ontologically problematic in that certain distinctions are necessarily conflated. Thus we can see that, although there are a range of theoretical (and ontological) positions that can be adopted, it is still often possible to evaluate these generally in terms of the consequences they bring for an account as a whole. Equating objects with their location is, then, at least problematic.

2.2 Basic tools: mereology, topology, geometry

Standard ontological modelling allows us to characterize the necessary requirements of a treatment of space very generally. Within formal ontology and several of the approaches to the formalization of qualitative spatial representation, the goal is to adopt a restricted set of primitives which allow axiomatization of the area of concern. In this section we first set out an overview of the kinds of primitives that are adopted in some axiomatizations and the alternatives that are commonly explored. We then show quickly why mereology is not sufficient and what is added when we move to mereotopology.

2.2.1 Alternative sets of primitives

One of the most common selection of primitives for the axiomatization of space is in terms of the sequence:

- **Parthood**
- **Connection**
- **Congruence**

Parthood is the the basic relationship of mereology that we introduced in detail in deliverable D1. Part has long been considered a basic primitive for all kinds of formal ontology that do not adopt set theory. "Mereology is a theory of the binary relation P (for part), originally introduced by Lesniewski ... as an alternative to set theory" (Masolo & Vieu 1999, p4). Connection is the basic spatial relationship between regions allowing the definition of **mereotopology**, a combination of topology and mereology that we will describe in a moment. And congruence allows statements of similarity to be made between regions. It is an extremely powerful relationship allowing geometry-like expressivity without resorting to abstract 'points'. Using these three primitives it is therefore possible to axiomatize formal systems of considerable complexity.

Various further alternatives are possible. In one, it is the *connection* relationship that is adopted as basic instead and this is progressively added to, building parthood and mereology out of connection. Questions need then to be asked concerning the precise notion of connection intended and how this corresponds to 'connection' in the world. In this sequence, some approaches include a *convex hull* operator that also gives stronger ways of talking about shapes—again this allows very complexly shaped regions and relations to be captured. In another alternative, one relies on a spatial interpretation of parthood itself, and defines connectivity in terms of some notion of regular self-connected regions and constraints on their overlapping—i.e., sharing parts. While in one final alternative, the starting relations are instead:

- part of
- **boundary for**

Approach	**Primitives adopted**
Borgo et al. (1996*a*)	Part, Simple-Region, Congruence
Randell, Cui & Cohn (1992)	Connection, Part, Convex hull
Smith & Varzi (1999)	Part, Boundary, Located-at

Table 1: Overview of spatial primitives adopted in qualitative formalizations

- **located at**

This latter, proposed by Smith & Varzi (1999), places a different emphasis on *boundaries* as the way of defining mereotopological concerns. And, in particular, they argue the importance of two distinct kinds of boundaries: boundaries in the physical world, or **bona fide boundaries**, and boundaries created by human convention or cognition, or **fiat boundaries** (Smith & Varzi 2000), that we introduce in more detail in Section 2.3 below.

In all of these alternatives we see a common goal and requirement—essentially, as we shall see in more detail in a moment, an appropriate model of space is to capture not only parthood but also topological concerns and further considerations of shape. And this can be done in various ways. Mereotopology can be created as a single set of axioms, or out of modules which either take parthood as basic and introduce topology or *vice versa*. And once created, it also needs extension in order to approach other important spatial aspects of reality and cognition. In Table 1 we summarize the more common choices of primitives for purposes of comparison.

We will introduce and discuss most of the particular positions taken on these fine levels of formalization as we describe them below. In this section, we begin by explaining why we have to be concerned with at least a mereotopology to say anything useful about space and how this is definitely not all we require. We also see, in the subsections following, that it is necessary to be explicit about the relationship beween space and objects and events, as well as the kinds of boundaries and connections that are assumed, as there are several different positions that can be taken and consequences follow from each.

2.2.2 Mereology and space

If we begin with the most basic formal mechanism adopted in ontology, the **parthood** relationship of mereology, this is not enough for saying very much about space. The parthood relationship is too weak to capture differences that are commonly required when ascribing spatial properties. The basic axioms of mereology state that the part relation is reflexive, antisymmetric and transitive. In addition, various particular axioms are generally selected in order to capture the precise notion of 'parthood' required. There is then a family of mereologies related according to the strength of the additional commitments they make beyond the initial three axioms. This family is described in detail in, for example, Casati & Varzi (1999, p48) and was summarized in Appendix I of Deliverable D1.

We can see that this is not sufficient for dealing with space even in the simplest terms if we consider a graphical rendition of the kinds of relationships that mereology defines; such a

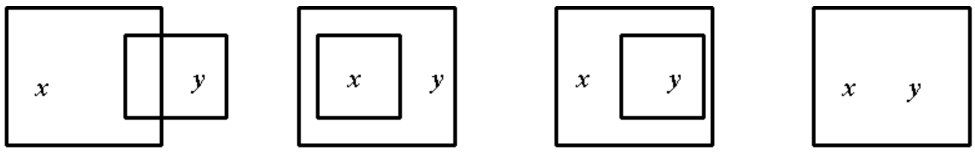

Figure 2: Graphical representation of the parthood relationship from mereology

rendition is given in Figure 2, adapted from Casati & Varzi (1999, Figure 3.1, p37). This shows the various kinds of distinct configurations that can be captured by the parthood relationship alone when defined according to a standard mereology. Moving from left to right across the figure we have the respective cases where x and y share some parts, where all the parts of x are also parts of y, where all the parts of y are also parts of x, and where both x and y share all their parts. According to a pure 'extensional' mereology, this latter case necessarily entails that x and y are then the same entity, because sharing all parts is a sufficient condition for identity.

Many other configurations are commonly encountered in spatial settings, however; this means that these configurations are not sufficient to describe spatial relationships. Central to the description of space are questions concerning an object's or region's connectedness, its parts and wholes, and its overlap with other objects/regions. For this reason, when adopted for spatial representation, mereology is usually supplemented with additional axioms that define topological relationships; this moves us into the realm of mereotopology.

2.2.3 Meretopology and space

The starting point for a whole range of axiomatizations of space is provided by considering the additional notions of **connectedness**. This draws on mereology as a fundamental aspect of any formalization of space (relevant to space regions and the objects located in them) but goes further. Connection allows a surprising complex range of spatially-interpretable configurations to be described. Typically,

> "Theories combining mereological notions and topological ones like those of 'being connected with', 'being an interior part of' or 'being self-connected' have been called mereotopologies"(Masolo & Vieu 1999).

Again, there is a family of such possible mereotopologies related according to the strength of their axioms (Casati & Varzi 1999, p63). In general,

> 'One may view mereotopology as consisting of two independent but mutually related components: a mereological component, concerned with the concept of *parthood* (or *overlap*), and a topological component, concerned with the concept of *wholeness* (or *connection*)" (Casati & Varzi 1999).

There are several possibilities for combining mereology and topology; and there are also several starting points. In Smith & Brogaard (2002) mereotopology is defined as a "theory of

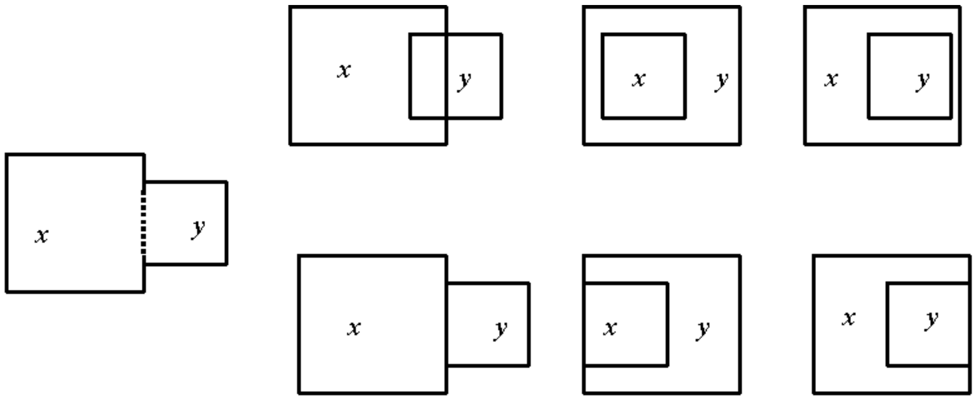

Figure 3: Graphical representation of the inter-entity relationships supported by mereotopology

boundaries, contact, and separation built upon a mereological foundation"—i.e., the 'boundary for' relation mentioned above—whereas in Eschenbach & Heydrich (1995) and several others it is the notion of the *region* that is made basic.

For ease of comparison with the extension beyond mereology, Figure 3 shows a graphical rendition of the kinds of relationships that are then covered by mereotopology; this time adapted from Casati & Varzi (1999, Figure 4.1, p56). From this we can readily see that a number of different 'relative positionings' of the related entities are now naturally covered.

There are several further issues in this area. For example, one common distinction that is drawn is between **open regions** and **closed regions**. Closed regions include their boundaries and open regions do not. Some argue that this distinction is also not cognitively relevant as the difference cannot be perceived: but this is not immediately obvious and Smith & Varzi (2000), for example, argue that the distinction is relatively natural for the case of physical objects within an 'empty space'—such as an orange lying in the open: the orange is then a closed region, it has its boundary (a thinner version of its peel) and the open air it is lying in is bounded by the orange but does not itself contribute a boundary. There are also topological consideration arising from the ways in which regions and contact between regions is formalized. Strict topological contact is seen as sharing at least some part of a boundary but can, as set out in Borgo et al. (1996*a*), be specified more precisely in various ways.

Spatial relations, such as those made possible within a mereotopology, are clearly centrally important for describing how an object is situated with respect to its surrounding space and to other objects. But, as discussed above, there are further relations beyond those which are covered by mereotopology: for example, spatial relations may need to describe an object's location, orientation, distance from other objects, and so on. This may require additional metrical information that is not included within mereotopology. Most of the accounts we will see below have therefore sought to extend the mereotopological view in various respects. There is also the question of what kinds of spatial relationships obtain between entities of *different* sorts, as regions are often conceived of as entities of a 'similar' kind. An early attempt at classification, Pullar & Egenhofer (1988), grouped spatial relations into the following areas:

- direction relations that describe order in space (e.g. north, northeast),
- topological relations that describe neighborhood and incidence (e.g. disjoint),
- comparative or ordinal relations that describe inclusion or preference (e.g. in, at),
- distance relations such as far and near and
- fuzzy relations such as next to and close.

A place for all of these distinct kinds of spatial relations needs to be secured in our overall spatial account.

2.3 Modelling spatial objects and their properties

Whereas the above gives certain starting points for talking about space and inter-relationships between spatial regions as such, we also need to consider the relationship between space and objects very carefully. The possible relationships or non-relationships with space can be used to define distinct ontological entities. Physical entities, for example, are typically distinguished from abstract entities precisely by virtue of their necessary location in time and space; we will see some of the alternatives in our more detailed discussions of particular ontological positions below.

In general, we want to characterize precisely in what ways physical objects (or events) can be said to be *located at* particular locations in space. This raises questions about how the objects concerned are to be identified and how, indeed, the locations are to be identified. We saw in the previous subsection that some notion of regions and relations over regions might provide the latter. But the relationship between such regions and physical entities still needs specification.

In order (a) to obtain the appropriate *identity conditions* and (b) for motion to be added straightforwardly to the account, it should be clear that the particular location of an object is not a necessary ontological feature of that object. An object has to have some location, but that location is, in general, arbitrary. Just where an object can be is nonetheless constrained by its physical constitution. More specifically, following the line of argumentation pursued in Borgo, Guarino & Masolo (1996*b*), we will assume that physical objects have a **physical substrate**—that is, the matter out of which they are formed or made. It is this physical substrate that has a necessary relation with space rather than the object directly. Although, as we noted above, it is also possible to propose space itself as the substrate for physical objects, we will not follow this line of investigation here. An object can only exist by virtue of the existing regions of matter out of which it is made and it is these are situated in space.

These distinct ontological domains lead Borgo et al. (1996*b*) to talk of *stratified ontologies* which decompose the characterizations of physical objects according to several layers, or strata, of description. In particular, they postulate the three ontological domains that we have seen already:

- the topological level, defining the domain of space in terms of simple, connected regions;

- the material of an object;
- the location of an object

plus a further level:

- the morphological level, which extends expressivity beyond mereotopology in order to talk about shapes, holes, edges and the like—here Borgo *et al.* rely on the **congruence** relation that allows generalizations to be drawn concerning shapes and their alignments.

Both boundaries and granularity are then addressed in terms of congruence; granularity is included by reference to an additional parameterization referring to chunks of matter that are *congruent* with granules selected to define the particular granularity intended. We will be able to find some version of these distinctions in most of the particular accounts that we examine below. More extensive considerations of congruence and its possible axiomatization have also been taken up by, for example, Cristani (1999) in terms of the constraint algebra MC-4 and its tractable subclasses.

The stratified view is useful in sorting out a number of otherwise difficult questions. For example, one issue of debate within the philosophical treatment of space is whether more than one object can be at the same location at a time. An ontological formalization of space and spatial objects needs to make sensible predictions—even when the objects considered become complex, as in, for example, the traditional philosophical example of a donut and its hole. Stratification allows us to say that it is not possible for more than one chunk of matter to co-locate, but since that chunk of matter may be a substrate for more than one object, then there may be more than one object at the same place and at the same time. Another traditional example here would be a piece of clay and a statue made out of the clay, or a wedding ring and the gold it is made out of. The statue and the clay, and the ring and the gold, can be distinguished ontologically in a number of ways (we will see this below in more detail); but they are both at the same location—in fact, they are *necessarily* at the same location, which is another task for an appropriate ontological characterization to predict. Another rather different example might be a joint meeting of two committees, which both happen to have the same members. The committees and their members have to be ontologically distinct, but they occupy the same location.

The more morphological information that is admitted to the account, the greater the capability to consider shapes, created spaces (above, below), holes, inherent spatial features such as corners and bumps, and orientation. All of which need to find their place in a full spatial ontology.

Representing these very basic, foundational features of spatial objects is a prerequisite for constructing intelligent reasoning systems. Such system can then operate in terms of situations or settings that are very much more like the kinds of settings that humans take for granted: this is the traditional link that is made to naive physics and modelling situations for intelligent behavior. These foundational properties should be anchored into the representation in rather more fundamental ways than is the case with contingent knowledge that may vary or be effected by events in the world. No matter what occurs, basic ontological relationships between objects, their constituting matter, and the locations of that matter will not change. Moreover, we can expect regularities over shape and change of shape, as well as flexible selections of granularity, to play crucial roles.

Some further traditional areas of philosophical concern relating to the spatial properties of objects include, for example, the spatial relationships that are required when considering *holes* and *boundaries*. Here a particularly useful elucidation has been offered by Smith & Varzi (2000) and Smith (2001) in their distinction between boundaries that correspond to divisions or discontinuities in the physical world, and which therefore participate in a wide range of causal processes independently of human attributions of significance, and boundaries that are 'imposed by fiat', such as the declaration of a particular line of demarcation between two countries as a border. Various aspects of this division will come to light below, here we note that Smith & Varzi (2000) develop out of this division important topological differences: the two kinds of boundaries give rise to different accounts of the notion of connection. This can also be used as a method for determining, or confirming, the kind of boundary one is concerned with.

Essentially, a physical, or **bona fide** boundary, allows the definition of closed regions, or objects, i.e., objects that include their boundaries. Then connection between two entities is defined as the sharing of a common part or boundary, i.e., there is overlap between either the two objects themselves or between one object and a closed version of the other. **External connection**, i.e., connection without overlap, is only possible when one of them is not closed. External connection is then not possible between actual physical objects since these are generally closed, although a physical object, such as a cup, can be in connection with its environment, the air, since the latter can reasonably be seen as open.

This raises problems for 'potential' parts, such as someone's head or hands, in that with bona fide boundaries alone, these also cannot touch but there is no genuine physical discontinuity to mark the boundary. This is particularly the case in examples such as a perfectly homogeneous sphere with respect to which it is still perfectly possible to imagine a dividing line (plane) and talk of the sphere's two hemispheres. One wants the two hemispheres to touch, there is nothing between them after all, but there is no convincing reason why one of the hemispheres should be open, and the other not: no one hemisphere has the clear right to claim the boundary line. This is then different from bona fide boundaries and represents the distinctive property of fiat boundary that they can be shared, or, more precisely, that a fiat boundary in fact calls two boundaries into existence: one for *each* hemisphere, and these are co-located. This is then to combine two previously held views of what boundaries in general are (associated by Smith & Varzi (2000) with Brentano and Bolzano respectively) by allowing them both, but as applying to two ontologically distinct kinds of boundaries: the bona fide and the fiat; we will return for more discussions of fiat entities of various kinds below.

Boundaries are also potentially awkward because of their **ontological dependence** on their hosts. This is more than simply a logical requirement in the sense of a piece of lexical semantics:

> "The dependence of a boundary on its host is a case of genuine ontological dependence, as especially Brentano has emphasized. It is not merely a case of conceptual

> or *de dicto* dependence, as when we say that there cannot exist a husband without a wife. Every husband, i.e., every man who is in fact married, could have been a bachelor (or so we may suppose). But the surface of a table can only exist as a *surface of* a table—perhaps only as a surface of *that* table." (Casati & Varzi 1999, p96)

This cannot be accounted for with the axioms of mereotopology on their own; there is no necessarily dependent notion of boundary already present. An appropriate axiom for extending General Extensional Mereotopology to include a boundary relation B has been suggested by Smith (1996) and, shown in the form used by Casati & Varzi (1999, p96), appears thus:[1]

$$(SCx \wedge \exists yBxy) \rightarrow \exists y(SCy \wedge BPxy \wedge \exists zIPPzy)$$

where SC is 'self-connected', i.e., "x is self-connected if any two parts that make up the whole of x are connected to each other" (Casati & Varzi 1999, p57) thus providing certain regularity conditions for the regions considered, BP is 'boundary part', i.e., a part x of y such that all parts of x are tangential parts of y (Casati & Varzi 1999, p85), and IPP is 'internal proper part'. As Casati and Varzi note, however, there will need to be some reference to necessity somewhere in an account such as this in order to make the statement strong enough.

An alternative they suggest (following Husserl and, for example, Fine 1995) is to have a separate 'module' of ontological dependence as such. This adds a new primitive to the P and C of mereotopology that are already present: i.e., D, where Dxy means that x is ontologically dependent on y. Fine's (1995) axioms for dependence are:

$$Dxy \wedge Dyz \rightarrow Dxz$$

$$Pxy \rightarrow Dyx$$

$$\exists y(Dxy \wedge \forall z(Dxz \rightarrow Pzy))$$

The last two link dependence and mereology; Casati & Varzi (1999, p97) then add:

$$Bxy \rightarrow Dxy$$

in order to link dependence and topology. This is, as they say, most in spirit with their approach to ontology where one begins with mereology and adds components as necessary: here topology and dependence. As we saw in our Deliverable D1, many of the formal upper-level ontologies make an explicit distinction between primary and dependent entities: the axioms given here can be seen as part of the formal underpinnings of such distinctions. They will clearly be essential for providing the necessary glue for holding together several spatial constructs, such as boundaries and their hosts.

Finally here we must also note that spatial objects, in a general sense of entities that stand in a relationship to space, are by no means limited to physical objects but naturally include at least physical processes. These also need to have their relationship to space clarified, although,

[1] Presented with an additional axiom to rule out curious objects such as those with boundary filaments growing out of the object.

as we shall see below, this can either be done directly in terms of 4D (3D+T) treatments (cf. Deliverable D1) or indirectly via the contribution of physical objects that serve as participants in a process. There are also some weaker relationships between certain non-physical objects: e.g., particular social norms of a community might be said in some sense to share the spatial properties of that community: exactly in what sense needs to be clarified considerably.

2.4 Paths

There is some discussion in the literature concerning the notion of a **path**, particularly related to how objects move through space. This notion is also relevant for spatial description, particularly given the role of paths and direction-giving within the SFB. Navigation has been of wide concern in many contexts—including, particularly, robot movement planning in AI; current developments that are deploying car-navigation systems also raise the relevance of generally applicable solutions. Within the SFB there are already several treatments of aspects of the navigation and spatial representation for navigation problems.

We will need here particularly to construct a close connection between the ontological categories involved and the **Route Graph** construct employed for guiding robot actions within the SFB. Providing a characterization of route graphs and the information that nodes in a route graph may include is an important task for the ontology specification. From the perspective of the spatial ontology, we will need to ask just what kind of entity *is* a path? It clearly has rather different relationships to other entities than 'straightforward' locations as such. The first move to an ontological specification for route graphs has now been taken by Krieg-Brückner, Frese, Lüttich, Mandel, Mossakowski & Ross (2005); this is also related to some other ontological domains in our discussion in Bateman & Farrar (2005).

2.5 Vagueness

One further aspect that needs to be mentioned concerns degrees of *vagueness* or inexactness. Spatial representations often need to deal with non-precise information—even when that information is already somewhat vague due to the adoption of qualitative specifications rather than metrical ones. This vagueness can also be considered ontologically, however: here we need to address just what is the status of such non-precise information as part of a complete specification of a foundational ontology. This allows us to be much more precise about just where alleged 'vagueness' lies.

'Vagueness' is a term that occurs frequently and in many different senses. Of these we are concerned solely with those that are taken to have ontological consequences. Vagueness, for example, that is due simply to more or less detailed specifications is not included: this is simply an aspect of possible **under-specification** or generality of description and is more aligned with semantic granularity (cf. Section 6.5 below) than vagueness as we will discuss it here. For the kinds of vagueness that we do consider, there are at least two common positions that have been taken: either ontological entities are themselves vague, or they are precise but simply difficult to pin down when we talk about them. That is, according to the first option, 'mountain' would be a vague entity that has no determinate boundaries in reality: the ontological construct is *inherently vague*; according to the second option, 'mountain' refers to an indeterminate range of precise entities: each particular entity is perfectly precise but we just do not know in general from the use of the word 'mountain' which of these is being picked out.

There is a considerable philosophical tradition in this area and particular logical formalisms have been developed to address these concerns. In terms of space, for example, Casati & Varzi (1999, p95) cite approvingly the position of Lewis, one of the foremost philosophers to have considered this area, thus:

> "The reason why it's vague where the outback begins is not that there's this thing, the outback, with imprecise borders; rather there are many things, with different borders, and nobody has been fool enough to try to enforce a choice of one of them as the official referent of the word 'outback'." (Lewis 1986, p212)

This is broadly compatible with the position proposed by Bennett (2001*b*), where:

> "I regard vagueness as a lack of clearly defined criteria for the application of a concept. Thus, it is a property of language not of the world itself. Typical examples of vague propositions in the geographical domain are: 'All *mountains* are *very high*'; and '*Near* the *marsh* is a *dense thicket*'. The words given in italics are the principal sources of vagueness."

Several of the approaches we will address below concern themselves specifically with the vagueness of certain kinds of spatial and geographical objects (Bittner & Smith 2001*b*, Bennett 2001*b*, Bittner & Stell 2002); this needs then also to be characterized in any broadly usable account.

We note here at the outset, however, two particularly important distinctions: the first is that between **uncertainty** and **vagueness**. As Bennett (2001*b*) explains, uncertainty is an epistemic state and is of less concern to us ontologically. Certain vague statements can still be perfectly 'certain', for example, again from Bennett (2001*b*):

> "...a statement such as 'The chair is in *the corner of* the room' is vague but can often be said with certainty, whereas an exact specification of the location of a chair (or even a range of possible locations) will typically be uncertain."

Bennett takes this as sharply contrasting with **generality**:

> "If I say 'I shall see you again later this month', this is an example of generality, since the claim can be fulfilled in many alternative ways. However, it is not vague, since 'later this month' refers to a precise period of time."

We will have some reason to consider this distinction below and to ask whether it is really as clear-cut as Bennett suggests; but it is in any case an important starting point.

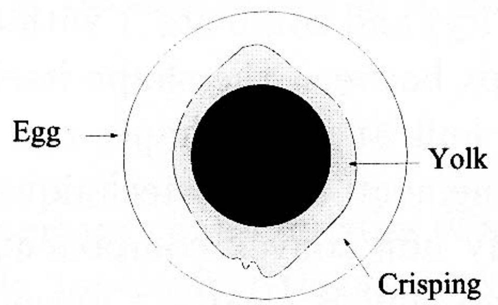

Figure 4: A vague region represented by an egg and its yolk: Lehmann and Cohn (1994)

The second distinction is between vagueness that follows because it is unclear which of several distinct possible definitions might apply—again, geographical terms provide many examples of this problem—and vagueness that follows because it is unclear where a boundary or threshold is to be drawn. Bennett (2001*b*) terms the former **conceptual vagueness** and the latter **sorites vagueness**. Thus, for example, a desert may be defined in a number of distinct ways which may not all come to the same judgement concerning a particular geographical area; while, a 'tall tree' may refer to trees of a variety of different heights depending exactly on which degree of height is currently relevant. These kinds of vagueness are ontologically distinct, although they can combine:

> "For example, to precisely interpret the concept 'tall man' we have first to decide how we are to measure the height of a man: must he remove his shoes and hat? what about hair and prosthetic limbs? what about posture? Once we have resolved these conceptual issues we then still have to deal with sorities vagueness in setting the threshold for tallness." (Bennett 2001*b*)

There have also been attempts to build vagueness directly into a treatment of regions. A number of related directions for extending the region connection calculus (RCC: cf. Section 7 below) are described in Cohn & Hazariki (2001), most prominent of these being the egg-yolk view of regions proposed by Lehmann & Cohn (1994). This sees regions as made up of a lower and upper boundary: the yolk and the egg respectively. The 'true' region then lies somewhere between these boundaries. This is shown graphically in Figure 4 taken from Cohn & Hazariki (2001). There are clearly similarities here with the notion of fuzzy sets and other vague or imprecise representations. We do not pursue this particular line of inquiry further here however, although we may be forced to return to it at a later stage.

2.6 Summary of spatial categories requiring characterization

The most basic parameter for an ontology of space is the way the notion 'space' itself is modelled. Space can be modelled as either an entity in the ontology in its own right or indirectly, dependent on the relations among the objects located in space. Concerning the

	static		dynamic	
	relation	**shape**	**character of change**	**trajectory based**
topological	interior, exterior, boundary, border, touch, contact, separation, ...	connection, hole, crack, fissure, ...	locomotion, separation, perforation, extension, ...	continuity, path, loop, ...
ordering	perspective, direction, horizontal, vertical, orthogonal, ...	cyclic, curve, convex, concave, corner, opening, indentation, ...	rotation, shift, deformation, bodily movement, ...	route, straight, turn, ...
metric	distance, angle, congruence, ...	length, volume, symmetry, proportions, curvature, ...	growth, shrinkage, approach, velocity, acceleration, ...	path length, ...

Table 2: Examples of spatial concepts organized by domain and dynamicity: taken from Habel and Eschenbach (1997)

entities located in space, the way physical objects are modelled is a major ontological concern. This includes the modelling of holes, surfaces, shapes, boundaries, edges and borders, as well as attention to the fundamental issues of dimensionality.

In the modelling of spatial objects, basic structures of reality need to be provided for, including accounts of part-whole relationships (mereology), connectedness (topology), or a combination of the two (mereotopology). A thorough inventory of the possible spatial relations which hold between the objects located in spaces is also necessary, as is the notion of a path, as many navigation tasks will be dependent on it. Finally, the concept of **frames of reference** (cf. Levinson 1996, Levinson 2003) will need to be noted; this will become crucial, particularly when dealing with the linguistic expression of spatial relationships. Although frames of reference bring to bear considerations of a perceiving subject that are often removed from purely spatial accounts within an Ontological perspective, we shall see below (particularly in Section 8) that some proposals for modelling spatial relations are also naturally interpreted in terms of particular frames of reference. We will have more to say on these issues when we return to consider linguistic aspects of ontology and spatial language however.

A useful overview of spatial concepts has been proposed by Habel & Eschenbach (1997); their table of examples is repeated here as Table 2 for ease of reference. All of these constructs must find their place within our general ontology of space, even if only in the form of a placeholder.

In the remainder of this deliverable we will set out the approaches that deal with space within a broadly conceived area of 'ontology'. We will see that there are many, often overlapping, approaches. In some places these amount to genuinely possible alternative perspectives on the phenomena at issue, in others there are modelling decisions that we need to evaluate for their appropriateness within a general account of a spatial ontology. One function of the overview is to bring all of the alternatives to bear equally, going beyond views prevalent within one theoretical tradition but perhaps less often considered explicitly in others. This wide ranging nature of the overview will mean that in some cases we will not be able to argue for (or in some cases even to particularly clarify) the modelling decisions made—there we simply present what has been done. The overview should allow the reader to have a common foundation in all of the areas addressed and provides a, we hope reasonably complete, stocktaking of spatial issues relevant for ontology construction.

The approaches adopted follow and build on the introduction to ontologies given in Deliverable D1 with the extension noted in the introduction to include input from qualitative spatial reasoning and geographic information systems. We begin with the more mainstream ontology approaches, which divide as before into the essentially hierarchically organized definitions (SUMO, OpenCyc) and axiomatized systems (DOLCE, BFO), before turning to the other sources of input new for this deliverable.

3 Spatially related entities in SUMO

The following sections contain a description and evaluation of the spatially related categories of SUMO (Niles & Pease 2001*b*, Niles & Pease 2001*a*), evaluated according to the fundamentals of space in ontology and according to the specific needs of the SFB.

3.1 Basics of SUMO space

The SUMO is a 3D ontology with many categories related to physical objects, a mereotopology, an extensive inventory of spatial relations, and a number of attributes and functions related to the spatial domain. The origins of the spatially related entities in SUMO are: Casati & Varzi's (1994) theory of holes and elements of Smith's and Guarino's respective mereotopologies. Spatial entities in SUMO are part of the physical universe, that is, subsumed by Physical, as opposed to being subsumed by Abstract. The relevant distinctions are shown in Figure 5.

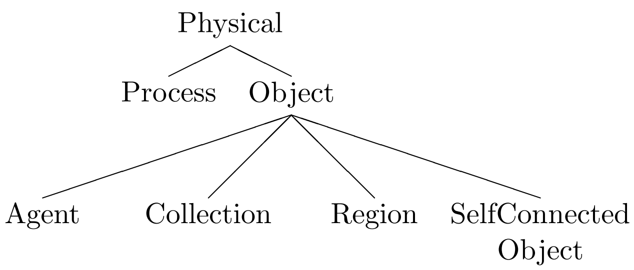

Figure 5: Taxonomy concerning Physical

From the figure, we see that there are concepts of Collection and Agent that are disjoint from the obviously spatial concepts of Region. Collections "have a position in space-time and members can be added and subtracted without thereby changing the identity of the Collection". Instances include flocks of sheep, football teams, and beds of flowers and so contrast with the mathematical notion of a set, which is abstract and has no position in space-time. Agents introduce in turn further considerations not immediately relevant to us here. Both take their spatial properties more from the other two concepts. We will, therefore, focus here primarily on Region and Self-connected object.

3.2 Objects in SUMO

SUMO contains an extensive inventory of physical objects. Instances of Object roughly correspond to:

> "...the class of ordinary objects. Examples include normal physical objects, geographical regions, and locations of Processes, the complement of Objects in the Physical class. In a 4D ontology, an Object is something whose spatiotemporal extent is thought of as dividing into spatial parts roughly parallel to the time-axis" (Niles & Pease 2001*b*, Object).

Physical objects necessarily have a location in time and space; this is captured by the axiom:

```
(<=> (instance ?PHYS Physical)
        (exists (?LOC ?TIME)
                (and
                  (located ?PHYS ?LOC)
                  (time ?PHYS ?TIME))))
```

That is, all instances of physical objects are also subject to a relationship of being located at a location `?LOC` and of 'being at a time' `?TIME`. We see the kinds of entity that may fill the location slot below.

A major topological distinction is made based on the connectedness of an Object: Objects that are of type SelfConnectedObject do not "consist of two or more disconnected parts".

3.2.1 Self-connected objects in SUMO

SUMO's breakdown of SelfConnectedObject is relatively straightforward if not particularly self-evident; it is shown in Figure 6.

The SUO-KIF (cf. Deliverable D1) axiom provided for SelfConnectedObject is:

```
(<=>
    (instance ?OBJ SelfConnectedObject)
    (forall (?PART1 ?PART2)
       (=>  (equal ?OBJ (MereologicalSumFn ?PART1 ?PART2))
            (connected ?PART1 ?PART2))))
```

The axiom may be glossed as: if something is an instance of a SelfConnectedObject then its constituent parts are connected mereologically—i.e., all pairs of parts exhaustively decomposing the object must be connected; this is essentially the definition for self-connection

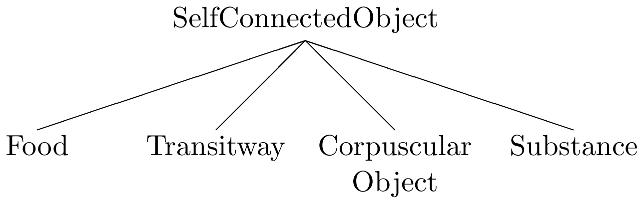

Figure 6: SelfConnectedObject in SUMO

given by Casati & Varzi (1999) as described in Section 2.3 above. Substances are then those objects "in which every part is similar to every other in every relevant respect" (Niles & Pease 2001*b*, Substance). This contrasts with CorpuscularObject "whose parts have properties that are not shared by the whole" (Niles & Pease 2001*b*, CorpuscularObject). The subclasses of CorpuscularObject include Artifact, OrganicObject, and ContentBearingObject.

3.2.2 Regions in SUMO

The next basic type of physical object is Region, i.e., a topographic location and space itself. "Regions encompass surfaces of Objects, imaginary places, and GeographicAreas" (Niles & Pease 2001*b*, Region). A relevant axiom pertaining to Region is:

```
(=>
   (instance ?REGION Region)
   (exists (?PHYS)
      (located ?PHYS ?REGION)))
```

This can be glossed as, 'if something is a region, then there is something located in that region'. The subclasses of Region are shown in Figure 7. AstronomicalBody is included as a Region presumably due to its relative size. As for GeographicArea, it is "a geographic location, generally having definite boundaries", e.g., Asia, the AtlanticOcean, or Texas.

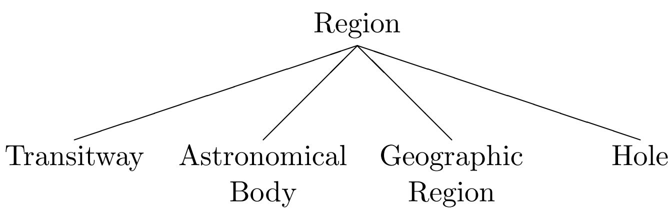

Figure 7: Region Taxonomy

There are certain modelling decisions made here which need to be considered more closely. For example, regions are themselves modelled as subclasses of objects: which means that rather different entities find themselves as siblings—'hole' and 'astronomical body' are very different (assuming we are not talking of ‘black holes’ here but rather the holes in donuts). The unifying feature appears to be that all of these may, by virtue of being regions, function as possible *locations* for some physical object (see below). An object can, therefore, be ‘on a transitway of some kind’, be ‘on an astronomical body’ (on the moon), be ‘in a geographic region’ or be ‘in a hole’. Why, again, just these possible entities are selected rather than others is not clear to us.

Morever, the axiom given above means that regions are necessarily filled by physical objects; which means that there is no such thing as an empty region. This is a rather strong philosophical position to take in its own right and may also be subject to critcism at the very least on grounds of cognitive plausibility.

3.3 Attributes pertaining to space

SUMO treats two aspects of space, shape and position, in terms of attributes. Attributes in SUMO are either internal or relational. Internal attributes inhere in some entity, whereas relational attributes exist by virtue of some relationship between two or more entities. Shape is treated as an internal attribute, as shown in Figure 8.

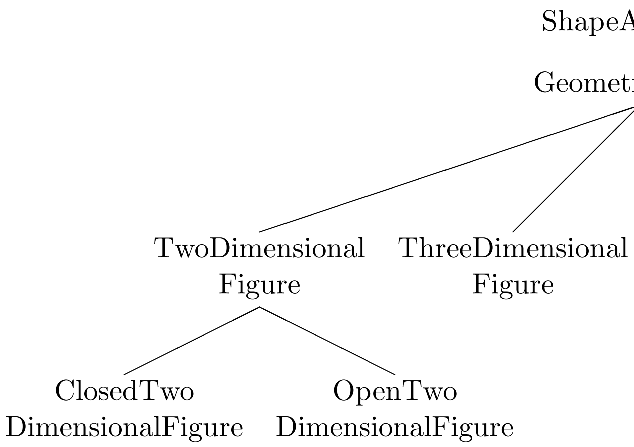

Figure 8: ShapeAttribute taxonomy in SUMO

Position on the other hand is treated as a relational attribute. A complete inventory of its instances (indicated by the i) are listed below:

```
PositionalAttribute
       Vertical(i)
       Horizontal(i)
       Above(i)
       Below(i)
       Adjacent(i)
       Left(i)
```

```
Right(i)
Near(i)
On(i)
```

PositionalAttribute has only one subclass, DirectionalAttribute, whose instances are the cardinal directions: North, South, East, and West. This reflects a more scientific than commmonsense view for at least the languages of Europe, although (cf. our deliverable D3 and Levinson 2003) there are languages where cardinality might be accepted as part of a model of the commonsense world also.

The axioms for instances of DirectionalAttribute and those of PositionalAttribute are similar and only indicate inter-concept contrasts and relations. A typical axiom for an instance of PositionalAttribute represents the opposition between attributes, but otherwise provides little semantics. Compare the axiom for Vertical to that for Left, for example:

```
(<=> (orientation ?OBJ1 ?OBJ2 Vertical)
     (orientation ?OBJ2 ?OBJ1 Vertical))


(<=>   (orientation ?OBJ1 ?OBJ2 Right)
       (orientation ?OBJ2 ?OBJ1 Left))
```

We will clearly need rather more detail.

3.4 Spatial relations in SUMO

The SUMO contains a number of relations which are instances of the very general SpatialRelation class. The documentation for this class is: "The class of relations that are spatial in a wide sense. This class includes mereological relations and topological relations." This means that for SUMO, even the part relation is given a spatial interpretation (thereby conflating mereology and topology). Overlap is defined, for example, soley in terms of part and biconditionally, so that sharing parts is sufficient condition for spatial overlap:

```
(<=>
   (overlapsSpatially ?OBJ1 ?OBJ2)
   (exists (?OBJ3)
     (and
       (part ?OBJ3 ?OBJ1)
       (part ?OBJ3 ?OBJ2))))
```

Two entities are connected if they either overlap or meet.

SpatialRelation is placed as an immediate subclass of the concept Relation, which is itself placed under Abstract. There is no defining axiom for SpatialRelation in general, but a partial taxonomy is given in Figure 9 and Figure 10. There is, however, a relationship drawn between a spatial relationship holding and temporal overlap of the entities related.

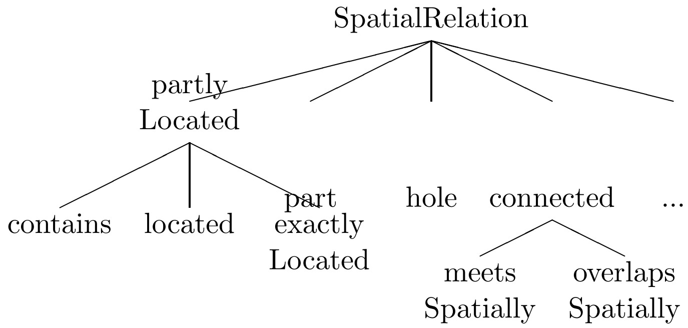

Figure 9: Partial taxonomy of SpatialRelation

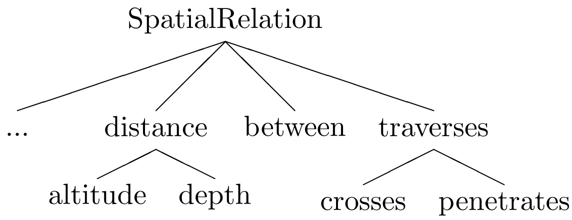

Figure 10: Partial taxonomy of SpatialRelation (continued from Figure 9)

There are some further value restrictions on the entities that may be related via these relations however. For example, located is restricted to be a two-place predicate holding between Physical entities and Objects. This is therefore rather general in that it might suggest that any object can 'function' as a location. It is, however, further restricted by inheritance from partlyLocated, which only allows as second argument entities of type Region (cf. Figure 7).

3.5 Paths

The category path is included as an instance of CaseRole in SUMO. Case roles are largely inherited or derived from the traditional notion of case and semantic roles employed in linguistic accounts, which we discuss at a more fundamental level when we turn to linguistic ontologies in our deliverable D3. It is unclear why the purely spatial notion of 'path' ends up in SUMO alongside such linguistic notions, although one of the motivations quoted for some of the SUMO decisions involve linguistic analyses. For now, we note simply the definition for path as follows:

```
(path ?MOTION ?PATH) means that ?PATH is a
   route along which ?MOTION occurs. For example, Highway 101
   is the path in the following proposition:
   the car drove up Highway 101.
```

Path relates a motion to a region and participates in the following axiom:

```
(=>
  (and
      (path ?PROCESS ?PATH1)
      (origin ?PROCESS ?SOURCE)
      (destination ?PROCESS ?DEST)
      (length ?PATH1 ?MEASURE1)
      (not
         (exists (?PATH2 ?MEASURE2)
            (and
               (path ?PROCESS ?PATH2)
               (origin ?PROCESS ?ORIGIN)
               (destination ?PROCESS ?DEST)
               (length ?PATH2 ?MEASURE2)
               (lessThan ?MEASURE2 ?MEASURE1)))))
   (forall (?OBJ)
      (=>
         (part ?OBJ ?PATH1)
         (between ?SOURCE ?OBJ ?DEST))))
```

There are a number of further spatial relations that may usefully be related to paths and path-following; for example: traverses, penetrates, crosses as well as between from above.

Transitway is another kind of path in SUMO and is subsumed by both Region (see Figure 7) and SelfConnectedObject (see Figure 6). A Transitway is defined as "the broadest class of regions which may be passed through as a path in instances of Translocation". Instances of Transitway include Roadway.

3.6 Summary and Comments

The major notions of SUMO-space have been surveyed. Perhaps due to its multi-theoretical basis, SUMO does not appear to embrace any one theory of space consistently. Therefore, it is not as clear as would be desirable for a general ontology of space within the SFB. SUMO's inventory of objects, in particular those pertaining to the class *Artifact*, may be useful for projects interesting in object recognition in, for example, a typical office environment. Also, SUMO's system of attributes deserves a closer look, as it provides an axiomatically grounded foundation for expansion into domain-specific areas. In particular, this part of SUMO would be useful if some SFB project required an inventory of spatially-related attributes. Furthermore, although the notion of path in SUMO is underdeveloped and not useable in the context of the SFB as it stands, there are clear points of contact both with the approaches taken in other ontologies (e.g., the OpenCyc view of paths that we see below) and with formalizations already present within the SFB such as the **Route Graph**. It should be investigated whether an ontological incorporation of the formalization of the Route Graph can provide a better starting point.

For all of these directions, however, the nature of the axiomatization needs to be considered most critically; we would expect, following our discussion in Deliverable D1, that these extensions would need to be placed carefully against a much tighter web of axiomatization than is the case in SUMO currently. Nevertheless, the modelling undertaken within SUMO stands as a useful point of comparison for any further ontology proposals in the areas that it covers.

4 Spatially related entities in OpenCyc

The following section contains a description and evaluation of the spatially-related categories of OpenCyc, evaluated according to the fundamentals of space in ontology and according to the specific needs of the SFB. This section builds on the basics of OpenCyc presented in D1. Again, the current version at the time of writing, OpenCyc 0.7.0, contains relatively few axioms other than those which relate to the taxonomy and the predicate argument restrictions. The discussion then will contain many listings of taxonomies, the most relevant aspects of which will be described in prose without the support of axiomatization.

4.1 Basics of OpenCyc space

The most general spatial category in OpenCyc is SpatialThing, which is "the collection of all things that have a spatial extent or location relative to some other SpatialThing or in some embedding space" (Cycorp 2004*b*, SpatialThing). This is a very generic class which subsumes both idealized spatial entities (e.g., Arctic Circle), geometric spatial entities (regions, points, lines, etc.), concrete 'tangible' spatial objects and even some events. The taxonomy in Figure 11 shows some of the detail associated with SpatialThing.

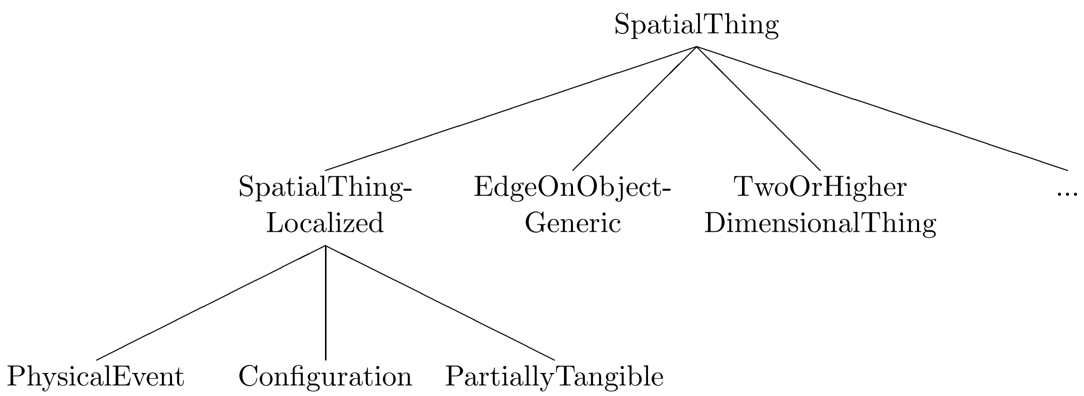

Figure 11: Partial taxonomy of SpatialThing

Everything in OpenCyc belongs to a particular microtheory. The following microtheories are most relevant in terms of spatially related notions: NaiveGeometryMt, Naive-SpatialMt, and NaivePhysicsMt. Most of the classes and relations discussed here, however, will actually come from the NaiveSpatialMt microtheory, documented as follows:

"This microtheory provides concepts and rules to represent the natural way we

> reason about spatial relations. It deals with...the conclusions we can draw from these objects being in static or dynamic situations with regards to only themselves or other objects." (Cycorp 2004*b*, NaiveSpatialMt)

That is, Cyc is primarily concerned with located objects rather than space as such; indeed, as Grenon (2003) argues, regions of space that exist independently of the objects that fill them do not carry much ontological weight in Cyc at all.

4.2 Objects in OpenCyc

Within OpenCyc there is no entity *per se* that corresponds to 'object' or 'physical object' as in other ontologies (cf. Object in SUMO). Rather, there is SpatialThing-Localized, defined as a concrete spatial category that can be empirically observed in the cosmos (e.g., Eiffel Tower, my dog, the earth).

> "The collection of all spatial things, tangible or intangible, that can be meaningfully said to have location or position in the empirically observable universe of the context in question. This includes all PartiallyTangible things, such as pyramids and ships, as well as certain Intangible spatial things, like the Equator. Also included are all Events that can be pinned down to specific places...and thus all PhysicalEvents." (Cycorp 2004*b*, SpatialThing-Localized)

The class SpatialThing-Localized subsumes a variety of naive spatial concepts, as shown below.

```
SpatialThing-Localized
      BiologicalLivingObject
      Border
      CavityInterior-Generic
      CavityOrContainer
      CloudlikeObject
      Configuration
      CustomarySystemOfLinks
      EmptyRegion-Generic
      Event-Localized
      Food
      GeographicalThing
      GeometricThing-Localized
      InformationBearingThing
      PartiallyTangible
      Path-Spatial
      Place
      SomethingToWear
      SpaceRegion-Empirical
      SpatialPathSystem
      Trajectory
```

4.3 SpaceRegion

SpaceRegion is a particularly relevant subclass of SpatialThing. The instances of SpaceRegion are defined as intangible regions of space located in the empirically observable universe, acting as locations for other spatial objects. "Instances of SpatialThing are said to 'occupy' some region of space" (Cycorp 2004*b*, SpaceRegion). A SpaceRegion might or might not be connected and is not limited to any particular dimensionality. Consider the taxonomy related to SpaceRegion in Figure 12.

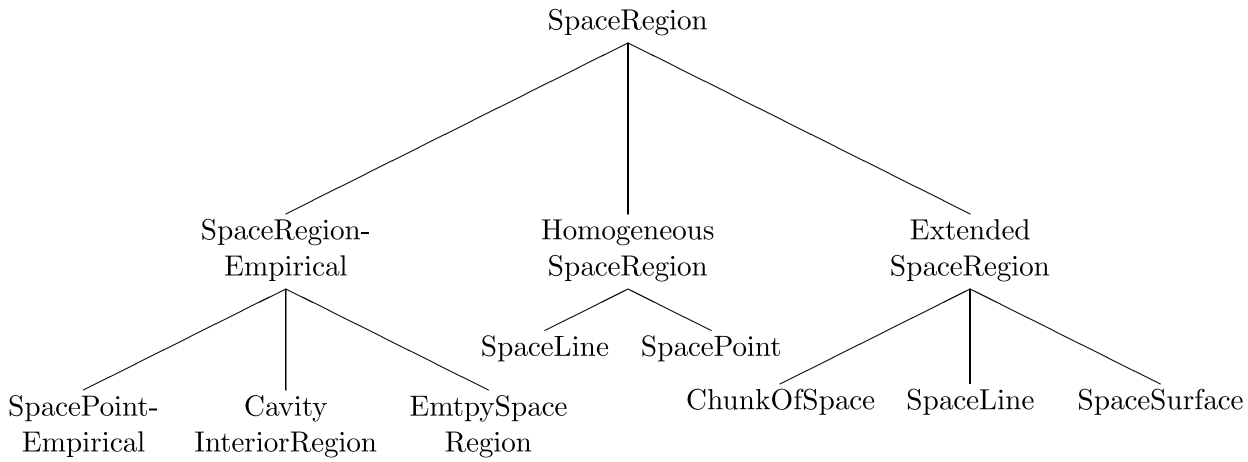

Figure 12: Taxonomy of SpaceRegion

An ExtendedSpaceRegion is the most basic type of SpaceRegion: "Instances of ExtendedSpaceRegion are portions of a three dimensional space that have an extent in at least one direction." (Cycorp 2004*b*, ExtendedSpaceRegion). The subclasses of ExtendedSpaceRegion are either one-, two-, or three-dimensional. In contrast, HomogeneousSpaceRegion "is a portion of a three dimensional space and is of uniform dimensionality". Finally, SpaceRegion-Empirical is both a specialization of SpatialThing-Localized and SpaceRegion.

> "The instances are intangible regions of space located in the empirically observable universe. A space region might or might not be connected (see spatially-Continuous). It might be partially or completely filled with (occupied by) instances of PartiallyTangible, or it might be completely empty (but cf. Empty-SpaceRegion)".(Cycorp 2004*b*).

SpaceRegion-Empirical relates more to a commonsense conception of space as would be reflected, for example, in natural language. OpenCyc freely mixes more scientific descriptions of space with these more naive descriptions, as the OpenCyc knowledge base is intended for natural language processing. However, whether simply mixing more 'technical' and 'commonsense' categories is an effective approach to NLP is certainly not self-evident.

4.4 Spatial features of objects

The shape of an object is represented using OpenCyc's 'collection of collections' mechanism, discussed in D1. In Figure 13 various shape classes are illustrated. Each subclass is instantiated by actual shapes, e.g., CylinderShape, SphereShape, etc.

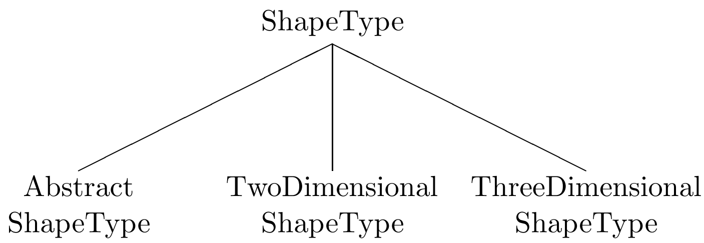

Figure 13: ShapeType and Instances

Instances of SpatialThing-Localized may be related to actual shape via the binary relation shape.

4.5 Spatial parts in OpenCyc

A particular kind of spatially relevant part in OpenCyc is that of surfaces. These decompose into various kinds of entities as such in Figure 14. Again, the particular selection of entities included and not included, and their motivation for being placed as siblings in this particular component of the ontology is not at present clear to us.

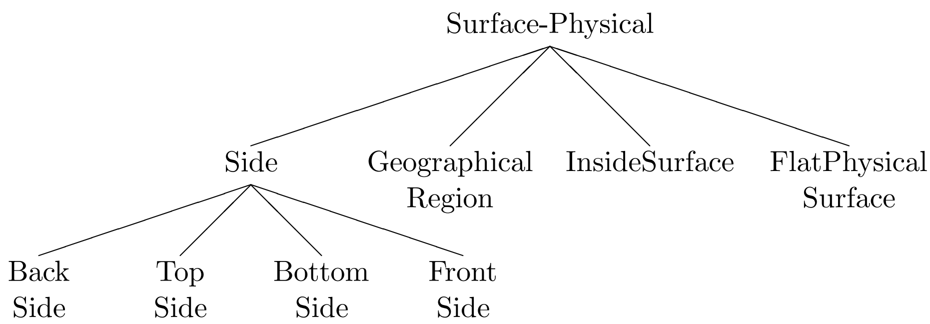

Figure 14: Sides of objects

A slightly different perspective on surfaces is given in the decomposition of Figure 15; this is intended to contribute to a general mereotopology. Note that a Surface-Generic is a specialization of a TwoOrHigherDimensionalThing (see Figure 11 above).

4.6 Spatial relations in OpenCyc

OpenCyc contains a relatively large inventory of spatial relations which may hold between instances of SpatialThing. The following lists the relations, all specializations of the predicate spatiallyRelated, according to the three microtheories most relevant to space.

```
spatiallyRelated [BaseKB]
    coDecompositions
    securedBy-Contributing
    convexHullOf
    fitsIn
    perpendicularObjects
    pointingTowards

spatiallyRelated [NaiveSpatialVocabularyMt]
    aligned
    connectedTo
    hasBeenIn
    parallelObjects
    notFarFrom
    connectedTo
    spatiallyDisjoint

spatiallyRelated [NaivePhysicsMt]
    physicalParts-disjoint
    onSamePlanetSurfaceAs
```

An important notion subsumed under spatiallyRelated is that of "in-ness", which is described with a high degree of granularity in OpenCyc. The first kind concerns fluids.

```
touches[NaivePhysicsMt]
```

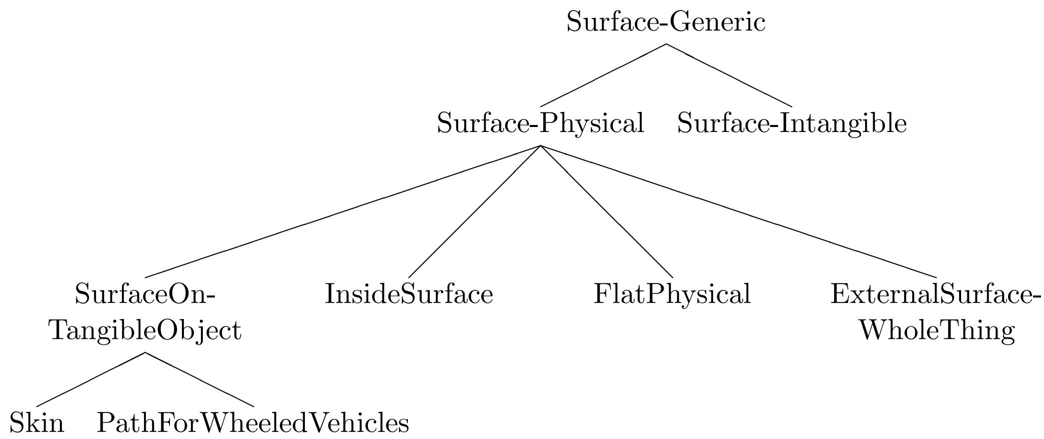

Figure 15: Surface in OpenCyc

```
in-ImmersedGeneric
     in-ImmersedPartly
            in-Floating
     in-ImmersedFully
            suspendedIn
```

The parallel notion of "in-ness" for solids follows from the touching relation covered previously. These relations place constraints on the type of container.

```
touchesDirectly-NotAsPart [NaiveSpatialMt]
       touchesDirectly-Apartonomic
            in-ContFullOf
            in-Embedded
            in-Lodged
                in-Spiked
            in-Permeates
```

The differences here focus on the object's make-up, e.g., in-Spiked can only refer to an object that is sharp at one end. The following focus on the co-locational properties of the two arguments.

```
inRegion [NaiveSpatialMt]
      spatiallyContains
      objectFoundInLocation
               in-Among
               in-ContGeneric
                     in-ContCompletely
                         in-ContClosed
                         in-ContFullOf
                     in-ContOpen
                     in-Held
                     in-Snuggly
```

It is interesting to note that some of these taxonomies may also relate usefully to embodied notions of action, activities and perception (e.g., 'snuggly'); this will only be considered at a later stage in our projects however.

The following taxonomies give the spatial relations which are specific to proximity and location. Proximity, or nearness, is expressed in OpenCyc with the following predicates.

```
near [NaivePhysicsVocabularyMt]
      inFrontOf-Generally
      behind-Generally
      alignedAlong
      hasPortalToRegion
      movesWith
      stuckTo
      spatiallyIntersects
      touches
```

The relation touches deserves special mention due to its granularity—demonstrated by the following taxonomy.

```
touches
      above-Touching
          groundOf
          in-Floating
          on-Physical
      adjacentTo
          bordersOn
          spaceRegionPortals
      alignedAlongSurface
          connectedAlongSurface
          sheetSurfaceTouches
      connectedAtContact
          connectedAlongSurface
          connectedAtContact
          connectedAtEnd
          in-Embedded
          pipesDirectlyConnected
      hangsFrom
      hangsAround
      in-Held
      in-ImmersedGeneric
          in-ImmersedFully
          in-ImmersedPartly
      in-Snugly
      touchesAtEnd
      touchesDirectly
      wearer
      wearsClothing
      wornOn
```

The follow account for how objects are located in space with respect to various frames of reference. The documentation for Cyc (but not OpenCyc) mentions several attributes pertaining to orientation:

- HorizontalOrientation
- VerticalOrientation
- UpsideDown
- RightSideUp

Likewise for Direction, the Cyc documentation mentions these values of TerrestrialDirection:

- Up-Generally
- Up-Directly

- Down-Generally
- Down-Directly
- VerticalDirection
- HorizontalDirection

In OpenCyc, absolute geographic directions are also well supported, as shown in the following taxonomy:

```
VectorInterval [BaseKB]
        UnitVectorInterval
             TerrestrialDirection
                 GeographicalDirection
                      GeographicalDirection-Direct
                      GeographicalDirection-General
                             North-Generally
                             North-Directly
                             South-Directly
                             ...
                             East-Directly
```

There are, then, a considerable number of space-related categories for consideration within OpenCyc–their precise interrelationships will need closer study, however, in order to see how they can best be related with other ontological initiatives.

4.7 Paths and path systems

The path microtheory in OpenCyc contains a particularly rich inventory of knowledge associated with the notion of *path*. Paths in OpenCyc are not modeled directly as mathematical objects such as graphs, transition diagram or maps. Rather, any class instance, e.g., River or Roadway, can act as a path, that is, can be *construed* as path-like or as a path element – point, node, link. OpenCyc's mechanism for path construal consists of a large inventory of role relations which relate arbitrary instances of OpenCyc classes – usually a concrete, physical object – to path constructs. This use of 'roles' in ontology is one which we will return to several time below. Furthermore, instances of OpenCyc classes are structured into *path systems*, which in turn possess well-understood mathematical properties and are organized in a lattice according to those properties.

The most general path construct in the path microtheory is Path-Generic. This class is further subdivided according to the taxonomy in Figure 16. Any instance of a path is either a Path-Simple or a Path-Cyclic. These classes are disjoint, where a Path-Simple:

> "...is a path with two distinct ends that do not overlap each other (in the case of spatial paths, the two ends are spatially disjoint). Since no instance of Path-Simple has ends that join at one point, Path-Simple is disjoint with Path-Cyclic.

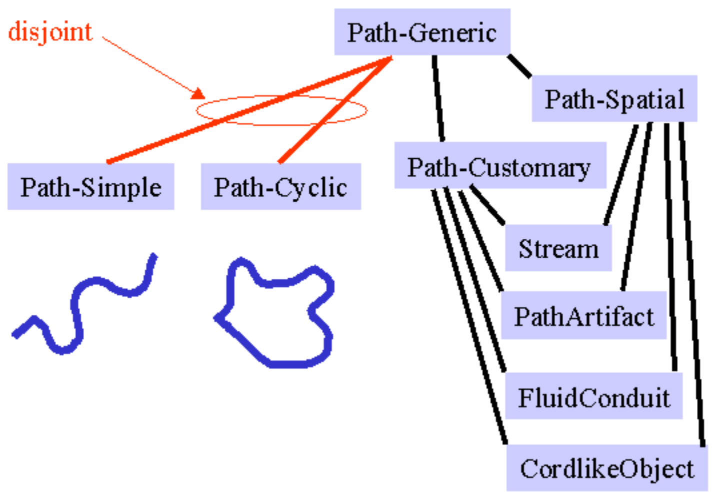

Figure 16: The top taxonomy for paths.

> Although instances of Path-Simple have distinct ends, some instances may have more than two things that are its end-points . For example, a path between Austin and Pittsburgh can also be a path between Texas and Pennsylvania. Notable specializations of Path-Simple include Pipe-GenericConduit, Nerve, and Stream"

Anything that is commonly construed as a path is classed as Path-Customary, e.g., a river, roadway, or hallway. Path-Spatial includes anything that has spatial extent, such as a log across a stream.

Paths themselves are simple entities with limited structure, e.g., sub-paths, nodes, etc. The more complex structures in OpenCyc, of which paths are just one component, are subsumed by the class PathSystem, which:

> "...consists of a nonempty set of points (where each point is an instance of Thing, and is related to the instance of PathSystem via the predicate pointInSystem), a set of nodes (a subset of the set of points — each node is related to the instance of PathSystem via the predicate nodeInSystem) a set of links (where each link is an instance of Path-Simple, and is related to the instance of PathSystem via the predicate linkInSystem), and optionally a set of loops (where each loop is an instance of Path-Cyclic, and is related to the instance of PathSystem via the predicate loopInSystem). In order to specify which link is between which two nodes in the system, which point is on which link in the system, which node is on which loop in the system, and so on..." (Cycorp 2004*b*, PathSystem).

From the OpenCyc documentation, we may conclude that a link joins two nodes while points may be found along the link not necessarily contributing to the structure of the graph.

Whereas nodes are like intersections along a roadway, points are like incidental landmarks such as buildings or landscape features. In a slightly more formal characterization, a PathSystem S is a 4-tuple $\langle P, N, L, (O) \rangle$, such that:

- P is set of points where each point is an instance of OpenCyc Thing and where each point is related to S via the relation pointInSystem.
- N is a set of nodes where N is a subset of P and where each node is related to PathSystem via the nodeInSystem relation.
- L is a set of links, where each link is an instance of Path-Simple and each link is related to PathSystem via the relation linkInSystem.
- optionally, O is a set of loops, where each loop is an instance of Path-Cyclic.

Below are the relations that have the effect of relating real-world concrete objects to path systems (Cycorp 2004*a*):

- pointInSystem – any point in a system, whether it is a node or not.
- nodeInSystem – a node – a point that is a junction, dead-end, designated intermediate point or an isolated point.
- deadEndInSystem – a node with only one link joining it.
- junctionInSystem – a node at which three or more links join, or at which loops and links join.
- isolatedNodeInSystem – an isolated point, not linked to any other, and having no loop.
- linkInSystem – a link of a path system.
- loopInSystem – a loop of a path system (a path from a node back to itself containing no other points) .
- pathInSystem – a path which is within a path system.
- cycleInSystem – a cycle: a loop or else a cycle made of two or more links (and the same number of nodes).
- linkBetweenInSystem – a link links two specified nodes.
- pathBetweenInSystem – a path joins two specified points.
- connectedInSystem – there is some path connecting two specified points

Also, Figure 17 gives a diagrammatic view of the various types of path elements in a PathSystem. The various real-world objects in the context of a path may be related to one another without specifying them in terms of a PathSystem. The following gives an informal listing of these relations:

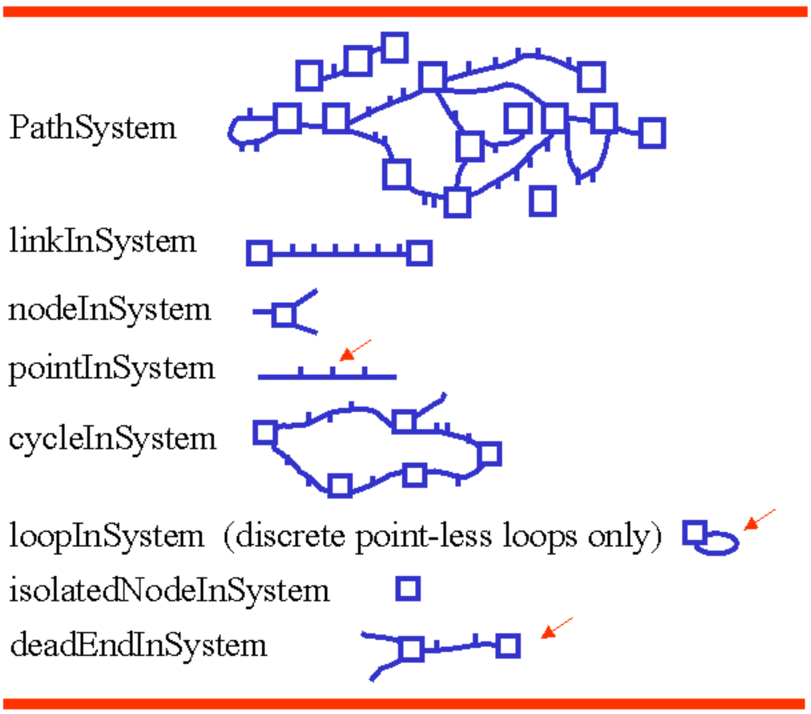

Figure 17: Elements of PathSystem.

- onPath – thing a is on path p
- pointOnPath – thing a is a stationary thing on path p
- betweenOnPath – a is between b and c on path p
- subPaths – p is a sub-path of q
- pathTerminus – point a is an end of path p with no further paths of the same general type extending beyond it
- adjacentPathsAtJunction – paths that join at junction
- endsOfPathSegment – relates path to its endpoints
- branchesInto – main path branches into little paths
- sideBranches – a side path connects along a main path
- pathBetween – p is a path with a and b as ends
- pathConnects – path p connects a and b (and may go on)
- linksOfCustomarySystem – relates a

For example, (subPaths PATH SUBPATH) means that any point on the Path-Simple SUBPATH must also be on the Path-Simple PATH. An example would be a stretch of corridor as included in a larger path through a building.

OpenCyc also provides a class for path traversal, called Traversal (Cycorp 2004*a*). A Traversal is a collection of paths ordered so as to represent movement over some path system. A Traversal is "all the paths traversed in a walk through a path system, in order" (Cycorp 2004*a*). There are a number of functions are relations associated with the class Traversal:

- TraversalFn – A function that returns a Traversal.
- traversalFrom – relating a Traversal and a "starting point" of it.
- traversalTo – relating a Traversal and an "ending point" of it.
- pointOnTraversal – relating a point and a Traversal in the sense that the point is somewhere along the Traversal.
- subTraversals – relating a Traversal and its subtraversals.
- traversalFromToInSystem – a Traversal presented in a PathSystem with its definite "starting point" and "ending point" in the system.

To summarize, 'paths' in OpenCyc are abstractions, or construals of real-world objects as 'nodes', 'links', and 'points' in a path system. The classes Path-Generic and PathSystem are essentially structures for capturing the construals of real-world objects as path-like entities.

The richness of knowledge in this particular domain of OpenCyc lies in the various relations, including the relations from real-world object to path construct and the relations among the real-world objects when they are construed as path elements. Beyond physical objects acting as paths and path elements, the general idea of some entity acting as something else is taken up again in Section 11.1 when the general notion of roles is considered in some detail.

4.8 Summary and comments concerning OpenCyc

OpenCyc's ontology has been described with respect to its spatial relations and categories. With reference to the general parameters of space discussed in Section 2, it is difficult to say what kind of general approach OpenCyc takes to space—is it geometric or region-based, Leibnizian or Newtonian? Perhaps all that could be said is that OpenCyc exhibits properties of all of them. This enormous coverage relates to OpenCyc's main strength, namely that certain spatial sub-domains are intricately detailed. For instance, OpenCyc contains a rich vocabulary of spatial relations as detailed in Section 4.6, although either they are not yet fully axiomatized or their axiomatization is not yet being made available. One area of OpenCyc-space which could be immediately useful for the SFB—in particular to the project A1-[RoboMap]—is the Path microtheory.

Some very important spatial notions are noticeably missing from OpenCyc, however. They include particularly a vocabulary for dealing with **reference systems**—a state of affairs that seems odd considering that reference systems are crucial in the mapping of natural language expressions to non-linguistic spatial concepts (See our deliverable D3 and, e.g., Levinson 2003). Furthermore, spatial entities, as with the rest of OpenCyc, are named and organized based on the lexical semantics of English. On the one hand this would be expected considering that the original Cyc project was intended primarily to aid NLP. But as a result, categories in OpenCyc, especially spatially related categories and relations, reflect the idiosyncrasies of English. Consider, for example, the concept SpatialRelation. As with English verbs, manner of motion or some attribute of the described moving object are conflated with the intended meaning of the spatial relation, e.g., in-Floating. This is not the case for many language families of the world, and so this may complicate the use of a Cyc-like organization. But, given that OpenCyc is so closely related to the lexical semantics of English, some of its components may turn out to be useful in a formulation of linguistic ontology which we explore in depth in D3.

5 Space and spatially related entities in DOLCE

The basics of DOLCE space as presented here are based primarily on Masolo, Borgo, Gangemi, Guarino, Oltramari & Schneider (2002) and Masolo, Borgo, Gangemi, Guarino & Oltramari (2003); the particular forms of axioms that we present are taken from the latter document although their content has not changed substantially between the two. We saw in deliverable D1 that DOLCE must be considered as one of the most rigorous and axiomatically rich ontologies among the ones surveyed for the SFB and its concerns are strictly the 'uppermost' levels of a general ontology. It accordingly has comparatively little to say about the particular detailed spatial relationships that we have seen for SUMO and OpenCyc above. However, as with all of the categories defined in DOLCE, what it does have to say about space is very deeply integrated into the basic axiomatization of its description of physical objects and events as a whole. This section will accordingly be relatively brief, and will focus on those aspects of DOLCE that will become relevant should the ontology be further enriched in the direction of more specific spatial commitments. It should also be noted that the developers of DOLCE have been involved with ontological formalization of space for a considerable period (cf. Masolo & Vieu 1999) but this has not yet found its way into the DOLCE ontology proper; adding this information is in many respects highly compatible with an extension of the spatial area that might be undertaken within the SFB.

5.1 Basics of DOLCE space

As discussed in Deliverable D1, one of the most fundamental divisions made in DOLCE is between entities that unfold in time, called perdurants, and entities which are wholly present in time, called endurants. Both of these are involved with spatial information of various kinds. Indeed, several of the basic assumptions of DOLCE bear on its conception of space and related notions. The upper level defined by DOLCE is primarily concerned with ensuring that the appropriate categories of existents are necessarily associated with spatial information where appropriate and that ontological dependencies (such as spaces related to some object, e.g., 'under the table' or 'the hole in the donut') are correctly captured: this latter is particularly important in order to prevent such parasitic spaces having an existence that diverges from their hosts.

We saw in deliverable D1 that DOLCE takes a particular view on the notion of 'qualities': particular qualities inhering in some object, such as its color, are related to general **quality spaces**. It is the quality space that allows qualities to be compared and contrasted: that is, because the particular color of some rose is located within a general color space, we can compare the color of this rose with others and can also cope with the color of the rose changing: which means more precisely that the particular quality that is the color of the rose is placed differently within the color quality space. To separate out these constructs more transparently, the *value* of the particular color that a rose has (the quality as such) is called a **quale** (cf. Figure 12 from Deliverable D1).

Thus, for some rose `rose#1` and its particular unique color `c#1`, we have the DOLCE description (Masolo et al. 2003, p18):

$$\text{PED}\ (\texttt{rose\#1}) \wedge \text{PQ}\ (\texttt{c\#1}) \wedge \text{qt}\ (\texttt{c\#1}, \texttt{rose\#1})$$

i.e., the particular rose is a physical endurant (PED), the particular color is a physical quality (PQ), and there is a quality relationship holding between them. The particular value of the color quality is then given by the further description:

$$\text{PR}\ (\texttt{color\#1}) \wedge \text{P}\ (\texttt{color\#1}, \texttt{color space}) \wedge \text{ql}\ (\texttt{color\#1}, \texttt{c\#1}, t)$$

That is, a particular element `color#1` is within (part-of: P) the physical region (PR) of color;

that element might be, for example, 'red'. This element is then the quale (value) of the particular color `c#1` that the rose has. This level of indirection is necessary in order to stop actual colors (such as 'red') changing when the rose begins to droop and is no longer red. The possibility of the value of the color changing over time is allowed for with the time index given as parameter t.

This account is important for us here because it is also how DOLCE treats space. Thus, within DOLCE there are also physical regions that correspond to 'space'; this as a whole is termed the space region. A particular location that an object or event may have is then given a value by relating it to this general space region. The relationships are exactly as illustrated for color. Therefore, the rose, in addition to color and by virtue (as we shall see below) of being a physical endurant (PED), has a spatial location:

$$\text{PQ}\ (\texttt{l\#1}) \wedge \text{qt}\ (\texttt{l\#1}, \texttt{rose\#1})$$

That spatial location, `l#1`, then also enters into a quale relationship such as:

$$\text{PR}\ (\texttt{location\#1}) \wedge \text{P}\ (\texttt{location\#1}, \texttt{space region}) \wedge \text{ql}\ (\texttt{location\#1}, \texttt{l\#1}, t)$$

Again, there is a temporal parameter to the quale relation which means that the location of objects can change over time (i.e., they have a different value within the space region) but, crucially, an object cannot become separated from its own nature of 'having a particular location' (the quality `l#1`).

The DOLCE upper taxonomy categories that are relevant for space are shown in Figure 18; this is an extract from the full taxonomy that we saw in Deliverable D1 (Figure 11). Here we see that although endurants and perdurants are crucial for DOLCE, space is actually modelled as neither. Spatial notions are divided among the physical qualities themselves—i.e., objects have real physical spatial locations—and the abstract regions that give those locations values. These regions are 'quality spaces' in the sense of Gärdenfors (2000) and should not be confused with the notion of 'region' given in other ontologies (cf. BFO in Section 6).

Within the axiomatization of DOLCE, that is, the part of the specification which guarantees that the above categories fit together in the way intended, the basic spatial properties are captured as follows. This also gives a further good example of how a dense axiomatization can capture very succinctly a wide range of necessary generalizations; it is interesting to compare these with some of the axiomatizations given above.

$$PED(x) \rightarrow \exists y(qt(SL, y, x))$$

i.e., all physical endurants have a spatial location (this ternary form of qt specifies additionally the type of quality that holds).

$$MSD_S(PQ, PED)$$

This guarantees that there is a *mutual specific spatial dependence* (cf. Deliverable D1) between physical qualities and physical endurants; that is, physical qualities can only exist when there is some specific physical endurant to 'carry' them and specific particular physical objects have

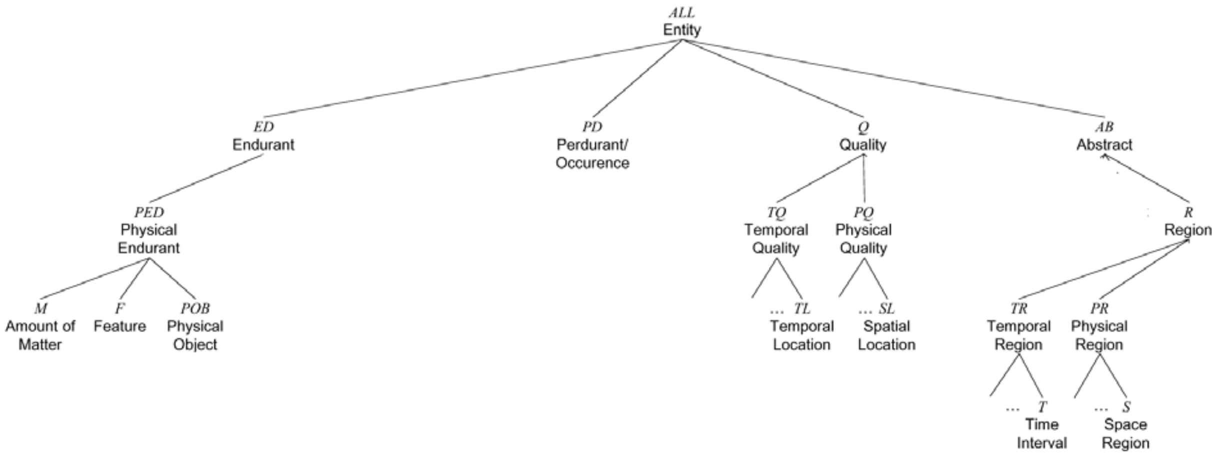

Figure 18: Spatially relevant extract from the DOLCE taxonomy (cf. Masolo *et al.*(2002: 9))

their own particular qualities (not someone else's). Moreover (by the definition of MSD_S), both the physical endurant and the physical quality are present in the same settings (s, s', ...) and at the same times—where being present is defined as:

$$PRE(x, s, t) \equiv PRE(x, t) \wedge \exists s'(ql_S(s', x, t) \wedge P(s, s'))$$

This states, in addition to specifying that being present has a temporal aspect, that if something is present then there is also some setting that is the spatial location of that something (ql_S).

The relation ql_S ranges not only over physical endurants, but also over physical qualities and perdurants, i.e., events and states, although differently for each. The case for physical endurants we have just seen, whereas for both physical qualities and perdurants the spatial aspect of being present is indirect: the spatial quality only comes via their associated physical endurants. For physical qualities, this is straightforward; for perdurants, the relevant physical endurants are defined by the maximal physical participants, which draws on the DOLCE view of how participants are involved in activites, states and events. The details are not relevant at this point, but the intuition is well-known in the modelling of perdurants—i.e., that while physical objects have their own locations, those of activies and events only have locations by virtue of where their participants are.

Despite these finely specified dependencies, the actual approach to space as such is still left completely open within DOLCE itself. As Masolo et al. (2002) explain, treating space in terms of quality spaces "...allows an homogeneous approach that remains neutral about the properties of the geometric/temporal space adopted (for instance, one is free to adopt linear, branching, or even circular time)" (Masolo et al. 2002, p12). The axiomatization of particular views of the space region and its parts is left to particular extensions that may be built on the DOLCE foundation. This is then amenable to other diverse ontological decisions concerning space such as those we see in other sections of this deliverable. The space region may be purely geometric—i.e., defined in terms of space-points—or it might be defined more qualitatively in the ways that we alluded to in the introduction and which we will describe in more detail below (see Section 7). All of these decisions leave the foundational connections and dependencies between physical endurants, perdurants, qualities and locations unaffected.

5.2 Physical objects, features, and spatial dependence

Here, we turn briefly to the kinds of entities that DOLCE sees as entering into spatial relationships in more detail. The most interesting additional cases are given by the subcategories of physical endurants. Physical endurants are categorized as: physical objects, amounts of matter, and features.

A physical object corresponds to the intuitive notion of an object as such, e.g., a hammer, a house, a computer, or a human body, and is the simplest case described by the spatial location axioms given above. Such physical objects enter into further important relationships. For example, a physical object is said to be *constituted* by an amount of matter. Again, axioms of the kind illustrated above are required to guarantee that the combinations of an object and the matter which constitutes that object are spatially and temporally well-behaved; this can be done without requiring any particular view of space. Note that this is necessary because both amounts of matter and physical objects are physical endurants—that is, both can enter into the kinds of spatial quality attributions discussed so far. The task is to make sure that these attributions necessarily remain in step.

This treatment of objects and their matter reflects the basic "multiplicative approach" that DOLCE adopts to entities and space. This implements the stratification of objects introduced in Section 2.3 above, by which two different spatial entities can be "co-located in the same space-time" (Masolo et al. 2003, p13) and still be considered non-identical. The multiplicative approach follows as one of the consequences of the OntoClean methodology (Guarino & Welty 2002) in which entities that have distinct identity properties must be considered to be distinct entities. Masolo et al. (2003, p14) give the following example of a vase and an amount of clay to illustrate the point: "necessarily, the vase does not survive a radical change in shape or topology, while, necessarily, the amount of clay does. Therefore the two things must be different, yet co-located." The vase is therefore said to be *constituted* by an amount of clay, and is not an 'instance' of an amount of clay. That is, identity is differentiated from constitution (what something is made of).

The distinguishing criteria for establishing identity, or the lack thereof, is not based strictly on the spatial property of co-locatedness, but on the *essential* properties that must hold in order for something to be considered an instance of a particular class. DOLCE thus recognizes the possibility of variable perspectives according to specific properties. What might in naive terms be considered the "same" object , e.g., the vase and the clay it is made of, must be considered as two different instances of two different classes. Ontologically, the two perspectives taken on the 'naive object' correspond to different entities which must then be bound together appropriately by axiomatization. This approach can be contrasted with that of, for example, Bennett (2002), who attempts to define everyday objects, such as vases, in terms of the spatio-temporal distribution of variously typed 'chunks of matter'; we will return to this approach in Section 9 below.

The other subcategory of physical endurants, features, are also of interest spatially. A feature refers to those tangible, physical characteristics of an object that are 'parasitic' in that they cannot exist without the existence of their hosts. That is, they are not distinguishable as 'parts' of an object in the sense that they could be isolated (even potentially) from their wholes. This includes traditional problem cases such as 'holes' (cf. Casati & Varzi 1994)—as in the holes of donuts—gulfs, openings, boundaries and so on. In particular DOLCE distinguishes two kinds of features:

- **Relevant parts** of entities, such as a bump or an edge;
- **Places** such as 'underneath the table', 'in front of the house', etc.

The former are rather peculiar parts that cannot be separated from their hosts in any way—which is what gives particular ontological poignancy to Lewis Carrol's smile of the cheshire cat. The latter are clearly not 'parts' of their host entity in any sense, but they would not exist without the host and define certain spatial expectations. Just which spatial expections will be a continuing topic of our work in the project I1-[OntoSpace].

Clearly, instances of PhysicalObject and Feature are **spatially dependent**. Instance x is spatially dependent on y if the following two conditions hold: x cannot be present at time t unless y is also present at time t (i.e., x and y are **co-present**); and, x and y are both located at place p (i.e., they are spatially **co-localized**). E.g., a hole in a donut is spatially dependent on the donut. DOLCE in general posits a generic dependency between features and their hosts rather than a specific dependency: as Masolo *et al.* note, this might be thought strange (how can the 'edge' of a table be the edge of any other table?) but they draw attention to difficult featural cases such as 'whirlpools', which are dependent on some water being present rather than particular amounts of water. It might be that the account will need to be made more differentiating here, although even the edge of a table might be said to continue its existence with a differing physical basis if it is, for example, repaired.

5.3 Summary and Discussion

DOLCE appears to do its job of providing a foundation for further work very well. It makes some fundamental relationships between entities and their locations clear while leaving the actual task of filling in particular models of space and location open. Traditional spatial relations such as above, below, left-of, etc. that we have seen above are properly left to the properties of the Space Regions that are to be defined.

DOLCE's quality inventory may have the potential to be exploited by an object recognition system. That is, if objects are identified by their qualities (using feature detection), then DOLCE provides a tight ontological connection between qualities and the objects in which they inhere. A marble sculpture may be positively identified based on either its constitutive marble quality or its morphological sculpture-like quality. This potential advantage relates also to the interaction of natural language components with spatio-visual components. That is, as will be discussed in Deliverable D3, the flexibility of natural language partially stems from the multiplicity of its wording, phrasing and semanticization potential. That is, the choice of particular linguistically motivated configurations is partially based on which features are salient in the discourse. In the context of the marble sculpture example, either of the following natural language expressions might yield the same result: ‘Go to the marble thing!’ or ‘Go to the sculpture!’.

6 Space and spatially related entities in BFO

The following sections present a description and evaluation of the spatially-relevant categories of the Basic Formal Ontology (BFO) described by Grenon & Smith (2004) and Smith & Grenon (2004). As we discussed in Deliverable D1, the most unique aspect of BFO—relative to the other ontologies evaluated for the SFB at least—is that BFO separates reality into two sorts of ontologies: SNAP and SPAN. A full description of space requires both types of ontologies and the mechanisms for relating entities of the two. Therefore, we explore the spatially-related entities of SNAP and SPAN in turn. In addition we address spatial relations including the trans-ontological relations that cross the SNAP-SPAN divide; this is an aspect of the second center of interest for us of the BFO, and that is its basic organizational assumptions about how ontologies are constructed and motivated. Of particular relevance for us with respect to space here will be the essential contribution of **granularity** within every ontological specification. We pointed out the potential importance of this in Deliverable D1; here we will build on this further with reference to spatial objects.

It is interesting to note before commencing, however, that the treatment of space is one of the areas in which DOLCE and BFO take rather different modelling decisions. We will need to weigh this carefully in our further construction of ontologies for the SFB.

Although the precise category-subcategory structure of both SNAP and SPAN components is still subject to discussion and we find slight variation across papers, it will be most useful for our purposes to follow the diagrams given in Deliverable D1 (Figures 14 and 15) in order to achieve a consistent presentation.

6.1 Basics of BFO space: SNAP

SNAP ontologies are ontologies where time is not present: they contain a snapshot view of reality, whose elements are endurants. SNAP then contains an ontological account of ‘space’ itself and all concrete objects located therein. In distinction to DOLCE, however, SNAP ontologies include space as one of their top-level categories; that is, BFO does not group space with other kinds of qualities or quality spaces but treats it as something ontologically unique in its own right. This top-level category is called spatial region.

6.1.1 The BFO Spatial Region

Within the BFO-SNAP ontology, space itself is considered an identifiable whole, namely “the entire spatial universe (the maximal spatial region)” of which all spatial regions are parts

(Grenon & Smith 2004, p148):

$$SpatialRegion(x) \equiv Part(x, space)$$

This means that all instances of SpatialRegion are parts of space itself. No other sorts of SNAP entities are parts of spatial region, but all other SNAP entities (BFO is still primarily concerned with physical reality) must be located in some spatial region (see Section 6.3). Space regions on all levels of scale act as locations for other SNAP entities. This is the *container* or *absolutist* view of space (Grenon & Smith 2004, p148) that we met in Section 2.1 above. BFO appears already to suggest a characterization of space in terms of a three dimensional space from which extracts of lower dimensionality can be extracted.

6.1.2 The BFO Substantial entity

One of the top-level distinctions drawn by BFO, as well as by many other ontologies, is between **independent** and **dependent** entities. We introduced dependence, particularly in relation to boundaries, in Section 2.3 above as well its role for DOLCE, particularly for spatial features, in the previous section. Here we will focus more on the independent entities of BFO. These are the essential entities which may receive their own location within space. Within a SNAP ontology, ordinary physical objects and parts thereof are subsumed by SubstantialEntity; a number of subcategories here are relevant for space as well as instances of the category itself being candidates for spatial location: in particular, the subcategory Substance is described as: (i) having continued existence, (ii) being bearers of qualities, (iii) being identity preserving over time, (iv) having location in space, and (v) having self-connected wholes with boundaries (Grenon & Smith 2004). Another subconcept of SubstantialEntity is that of the boundaries themselves, which are also relevant for characterizing space.

6.1.3 The BFO Site

In addition to substantial entities, however, BFO posits a further category of independent entities which are purely spatial in intention. These are the particular locations where entities can be and are called sites. A site differs from a substance in that the former acts as a location for the latter. The two are related by the predicate occupies: substances, then, occupy sites. This is considered ontologically distinct from the notion of 'location' which is used with reference to space region; location in this latter sense is purely mereotopological.

Examples of Site include a room in a house or a landing strip or the alimentary tract of a person. They are not instances of SpaceRegion, but are more object-like, in that both substances and sites are *located* in space. Thus, Grenon & Smith (2004, p150) argue that "SNAP can thus do justice to a certain feature of the relational approach to space" (the Leibnizian view discussed Section 2), namely that the properties of space can be described in terms of a three-way relation between Substance, Site, and SpaceRegion. Sites also begin to move towards what we will describe as a 'functional' orientation, in that particular kinds of sites may be associated with particular purposes. Three subcategories of sites, i.e., sites where an object might be, are: 'holes', 'places' and 'niches'.

Of these categories, the notion of the **niche** is the most distinctive; the others have similar counterparts in the other ontologies that we have seen. Niches are structured in terms of "functional properties (of temperature, foliage density, federal jurisdiction, etc.)" and are worth closer attention in their own right.

6.1.4 Niches: Anchoring objects into space

In Smith & Varzi (1999), the niche is proposed as a further fundamental spatial relationship in addition to *part of*, *boundary for* and *located at*. Smith and Varzi draw on mostly ecological and biological examples to clarify their intention with the introduction of niches, arguing that such entities are fundamental to a variety of accounts of reality. Niches build on 'physical-behavioral units' in order to characterize natural, bounded, unitary entities that separate "some organized internal (foregrounded) pattern from a different external (backgrounded) pattern." (Smith & Varzi 1999) Crucially, a niche is

> "not just a location in space; rather it is a location in space that is constrained and marked by certain functional properties (of temperature, foliage density, federal jurisdiction, etc.)." (Smith & Varzi 1999)

Within any niche there is a privileged spatial locus, a hole, into which its occupant, or **tenant**, exactly fits.

In the definition of niches, Smith & Varzi (1999) first axiomatize space as a standard mereotopology using parthood and boundary. Then there is a primitive 'exactly-located-at' relation (L) between, on the one hand, both entities and spatial regions, and on the other, spatial regions. A single entity can only have a single location, called *the location of* x (l):

$$l(x) =_{df} \iota y(L(x, y))$$

Finally, niches are formed from an additional primitive relation N between a niche and its tenant. Tenants exactly fill a special 'hole' in their respective niches and any niche can only have one tenant. (This latter is argued by Smith & Varzi (1999) by suggesting that any other similar entities present are then necessarily part of the niche rather than co-inhabitants; this is, perhaps, not particularly intuitive although certainly at least biologically defensible. If tenants can be topologically disconnected, i.e., groups or colonies of some kind, then this might also be a useful distinction.) The possibility of niches with vague outer boundaries is left open and is related to whether or not an ontology allows vague objects at all. Here Smith generally follows the position of Casati & Varzi (1999, p95) that we have described at several points above; i.e.:

> "On this account, then, vagueness is no issue for mereotopology. It is at most an issue that arises in connection with the drawing of fiat boundaries. As with the case of the equator, there may be some degree of indeterminacy when we speak of such boundaries. But the boundaries are not in and of themselves vague, so we need not fuzzify our mereotopology."

Niches are well behaved topologically and may have holes, cavities, boundaries and so on. More specific spatial constraints are not imposed. Some spatial details of the niche relationship are:

$$N(x, y) \rightarrow \neg O(l(x), l(y))$$

$$N(x, y) \rightarrow IP(l(y), l(x + y))$$

$$N(x, y) \rightarrow C(x, y)$$

i.e., a niche and its tenant are disjoint, their spatial extents (locations) cannot overlap, the location of the tenant is an interior part of the location of the niche and its tenant, and a niche and its tenant are connected.

For our purposes, the connection of niche with function and its bringing together ontologically of 'settings' and occupants for those settings is very interesting. The axiomatization of niches has so far been applied most to biological examples, where, for example, particular bacteria might have as their niche a cavity within a body, which itself is a tenant of a larger niche, such as, for example, a bear cave (if the body is a bear) or a river (if the body were a fish). For such examples it is clear that there is something more involved in their ontological characterization than pure topology or even geometry; the niche may provide a way of beginning to get at this. We return to this in our conclusion.

6.2 Basics of BFO space: SPAN

SPAN ontologies are intrinsically temporal: they always involve 4D views of reality. Time is inherently and explicitly contained within any SPAN ontology and all of its elements. Process, then, has SNAP entities as participants, but also a time component, as instances of Process are said to *happen*. One consequence is that 4D SPAN entities are not located in space, but in **spacetime**. Spacetime is "the totality of spatiotemporal regions", i.e., regions extended in time. This gives rise to a very similar definition of SpacetimeRegion to that of SpatialRegion given earlier:

$$SpacetimeRegion(x) \equiv Part(x, spacetime)$$

SPAN includes many four-dimensional 'analogues' of the three-dimensional entities found within a SNAP ontology. The SPAN equivalent of the category Site in SNAP is Setting. This acts as the 4D location of some Process. For example, "The Hundred Years War forms a setting for the burning of Joan of Arc" (Grenon & Smith 2004, p154). Here we see again aspects of functionality: the designation 'the Hundred Years War' can scarcely be seen as a pure geometric or topological characterization. This introduces a similar three-way decomposition analogous to the case of sites: a process, e.g., the burning, then takes place within a setting, e.g., the war, which is itself located within spacetime.

As we noted briefly in Section 2 above, Donnelly & Smith (2003) also argue that the SPAN approach is better able to deal with change, particularly motion in the spatial world, as compared to other region-based approaches. The argument proceeded as follows: In formulating a qualitative theory of motion within a region-based approach (cf. Sections 7 and 9 below), some kind of temporal indexing is required to supplement the mereotopology. That is, objects must be described in different locations at separate instances of time. This requires two independent sorts of spatial entities: fixed locations and moveable objects. But, "since the region-based approach admits only the first type of entity—the locations or regions—it must somehow simulate motion, for example via successive assignments of attributes to a fixed frame of locations" (Donnelly & Smith 2003, p2). This is similar to the approach within DOLCE, where the space region provides a fixed frame of locations (however that is defined) and one consequence of movement would be that the values of the location of particular objects changes with respect to that space region; this is why the relevant relations in DOLCE are indexed by times as we saw. In the BFO approach, there can be two perspectives on motion: one as a succession of SNAP entities with different time indexes, and one as a single SPAN entity extended in space-time.

Settings accordingly come in two flavors: the first 'with stationary spatial component', the second with 'mobile spatial component'.

6.3 Spatial relations

In this section, we focus in a little more detail on the particularly spatial aspects of the above and their motivations. Particularly the functional aspects of sites and settings are suggestive for treatments of phenomena that will play a role in many SFB projects; this also has interesting relationships (and raises similar problems) to the notion of situation discussed for the General Ontology Language of Heller & Herre (2003) in Deliverable D1. Despite the avowedly non-social orientation of BFO in its present form, many of the examples of sites and settings (as with the Joan of Arc example above) are essentially social in nature. This needs to be clarified.

Particular spatial relations, such as those found in OpenCyc and SUMO, are not instantiated at the level of BFO's formal ontology. They are left to material, domain-specific ontologies. However, at the formal level there are some very general spatial and spatiotemporal relations given: SpatialLocation, Occupies, and SpatiotemporalLocation. SpatialLocation is primitive in SNAP. Every SnapEntity x has a SpatialLocation y in some ontology [ω] (Grenon & Smith 2004, 148):

$$(SnapEntity(x) \wedge Constituent(x, \omega)) \rightarrow \exists y SpatialLocation(x, y, \omega)$$

Occupies is a more specific, derived spatial relation holding between a SubstantialEntity and a Site. "...i) the substantial entity and site which are joined by this relation do not overlap and neither do their respective locations; ii) the substantial entity's location is an internal part of the location of the sum of this entity with the site which it occupies " (Grenon & Smith 2004, p150). The relation SpatiotemporalLocation describes the location of a SPAN entity in spacetime and is primitive in SPAN ontologies, the counterpart of SpatialLocation in SNAP ontologies. Spatial and locational change are described in terms of relations that hold across two different SNAP ontologies with different time indices.

6.4 Layers

Whereas the former notions of spatial relations provide a basis for describing relations between objects of various kinds, Donnelly & Smith (2003) argue that they need to be constrained further in order to prevent certain non-sensical spatial relationships being admitted. For this, they introduce the ontological construct of **layers**.

Layers are invoked in order to avoid the various problems of (temporally-)spatially coinciding entities which belong to separate mereology chains. They can be seen as another approach to naturally constraining the very general application of the parthood relationship to only construct transitive chains that 'make sense'. In short, layers augment mereotopology so as to allow (temporally)-spatially coincident entities not to overlap mereologically. This account is very much aligned with the DOLCE and BFO approaches where objects are distinguished from the spatial regions that they occupy.

Two objects coincide when they occupy overlapping regions of space. Objects coincide with their regions but also potentially with other objects. Overlap is construed mereologically and this makes coincidence a broader relationship in that coincidence is possible between objects that do not share parts.

An example offered by Donnelly & Smith (2003) is the following involving fish, lakes and pollutants. They suggest four coincident three-dimensional layers:

L1. a region layer, consisting of a regular spatial volume including in its interior the spatial region occupied by the lake,

L2. a lake layer, consisting of a certain concave portion of the earths surface together with a body of water,

L3. a fish layer, consisting of a certain aggregate of fish,

L4. a mercury (or chemical contaminant) layer, consisting of tiny deposits of organic mercury scattered through the lake and through the tissue of the fish.

Crucially, entities from distinct layers can never overlap (in its proper mereological sense). A part of the mercury is not 'part of' the fish; nor are the fish 'part of' the water. The transitivity of parthood is therefore modified so as only to operate *within* layers.

This behavior is provided formally by restricting mereological sums so that they are only possible when entities **underlap**—i.e., there is some whole of which they are parts. Then two objects underlap if and only if they belong to the same layer. Donnelly and Smith explore several further variations and introduce a basic **region layer** that allows relationships of coincedence to be defined across entities from various layers.

Both the occupies and located-at relations of BFO proper may then be definable in terms of the relationship of coincidence.

6.5 Granularity and scale

There is a further dimension of 'perspectivalization' that is formally anchored within BFO and which it will now be appropriate to follow in some detail. This is the notion of granularity, which traditionally has been treated as some specification of 'relative size'. Items larger than some reference object are considered visible or significant, items smaller are invisible or insignificant. Explicit granularity parameters are given in the formalizations of, for example, Borgo et al. (1997), while a notion of 'refinement' is provided within a modal formalization by Asher & Vieu (1995). Our proposed use of granularity here will go considerably further in its implications however, drawing extensively on the view of granularity put forward in the work preparatory to BFO. The direction of development goes in a similar direction to that set out in work on Geographic Information Systems, such as Fonseca, Egenhofer, Davis & Câmara (2002), where the following useful distinction is drawn:

> "Some authors consider *granularity* in a spatial database to be the same as resolution, thus implying that granularity is related to the level of distinction between elements of a phenomenon that is represented by the dataset. Hornsby (1999) points out the difference between resolution and granularity. Resolution refers to the amount of detail in a representation, while granularity refers to the cognitive aspects involved in selection of features. This kind of granularity is called *semantic granularity*. ...In the [Ontology-driven GIS (cf. Section 11)] framework there are different levels of ontologies. ...We follow Hornby's approach because we consider that the level of semantic granularity is related to the level of ontology used." (Fonseca, Egenhofer, Davis & Câmara 2002)

We will be concerned here exclusively with the ontological construction of *semantic granularity.*

We begin the trail with the basic definition of BFO's formal apparatus given by Smith & Grenon (2004). This includes the statement that:

> "It is a fundamental tenet of our realist perspectivalism that the same reality can be captured in a plurality of distinct ontologies, all of which are veridical. We can go from a course- to a fine-grained perspective and back again, as occurs for example when some zoological phenomenon prompts us to investigate the features of the underlying DNA, which prompts us in turn to draw new conclusions on the level of zoology. Reality admits in this way of a sort of ontological zooming." (Smith & Grenon 2004)

This notion of zooming and granularity was introduced by Smith & Brogaard (2002) and developed further in the geographical and spatial domains by, for example, Bittner & Smith (2001*a*), Bittner & Smith (2003*a*) and Bittner & Smith (2003*b*). Granularity in the BFO sense provides a striking way of bringing together several important kinds of consideration—particularly, on the one hand and somewhat predictably, approximation and vagueness, and, on the other and not so predictably, function and purpose. Whereas most other approaches to vagueness and granularity remain fixed with the notion of scale, BFO goes further and relates granularity to the purpose for which a decomposition of the world is made. We consider this latter development as providing an extremely powerful way forward out of a labyrinth of ontological and conceptual modelling problems.

Our introduction here will follow that of Smith & Brogaard (2003). Smith and Brogaard first draw attention to the potential use of granularity for avoiding the well-known problems of mereological identity raised by such low-scale phenomenon as quantum effects on physical objects; for example, a stone, although it may lose (and gain) atoms regularly should still be considered, arguably, as ontologically a single continuing physical object. But the fact that it loses (and gains) parts means that it is not identical to itself at different times and hence, when mereology is assumed as a foundation for ontology, it cannot be the same object after all.

The usual response to this is to appeal to some notion of 'significant' part: an object can lose or gain insignificant parts without effecting its identity. The question is then, of course, just when do we (or rather, our ontology) know that something is significant or not. This is where granularity as a question of scale has been stuck for some time; significance can be defined as being 'smaller' than some specified spatial extent; we will see a formally well-developed version of this approach in our discussion of Bennett's account below in Section 9. The problem with this direction is that significance cannot be reduced to a matter of scale in any straightforward topological or metrical sense. This is just not what 'significant' means and this is no accident. Significance can only be defined against a background of purpose and function, and this has then to be accepted at the heart of a satisfactory ontological account regardless of the challenges it may raise for complete formalization.

We find similar considerations whenever notions of varying 'scales' occur. For example, within a completely different domain, Schlieder, Vögele & Werner (2001) describe how the description of motion may depend on the intention of the agent carrying out the motion which in turn may depend on the 'spatio-thematic region' within which the motion takes place and, crucially, such spatio-thematic regions are frequently organized in partonomies, i.e., hierarchies generated by parthood relations. They illustrate this with respect to actions carried out in a museum and the services that are offered within the museum dependent on the level in the spatio-thematic partonomy: at the level of the museum as a whole the relevant services are 'global navigation' and an agent is 'touring the museum'; at the level of a museum wing the services are local navigation and the agent is 'traversing the wing'; in an individual wing, the relevant services are 'general information' and the agent is 'visiting the room'; and finally, at a particular exhibit, specific information is relevant as a service and the agent is 'looking at the exhibit'. Each type of agent behavior may align with different motion patterns, such as moving quickly, standing still and looking at an exhibit, etc. Clearly, the different museum levels within the partonomy can be considered as differing granularities by means of which the other entities may be indexed. Determining which spatio-region applies is simultaneously a question of determining an agent's purpose and intention at that time. Partonomies may thus also play an important role in mediating across different granularity selections.

Returning then to Smith & Brogaard (2003), they argue that whenever a description or judgement is made, the world is partitioned according to some degree of granularity and statements that are made about the world are then true or false with respect to that granularity. The granularity selected is defined by the selection of a context. Then:

> "the work done by contexts in our theory rests on one single feature, namely on the fact that contexts may be *more or less refined*, or in other words that they may determine a greater or lesser *granularity* of ways in which we relate to objects in the world." (Smith & Brogaard 2003, emphasis in original)

The granularity selected then establishes the basic ontological individuals that are going to be of relevance for a particular model. Such individuals will necessarily respect the foundational requirements of such categories as they have been described in deliverable D1. But they will also in an important sense represent distinct ontologies that cannot be combined without more formal work: this is the perspectivalism of BFO.

Respecting appropriate granularities can serve as an additional mechanism for ensuring ontological clarity when constructing ontologies. In many respects, this is a specialization of OntoClean's identity criteria condition (cf. Deliverable D1 and Guarino & Welty 2004). Entities of differing granularities will often have differing identity criteria; whether we find cases where this is not so remains open at this time. We also see further use of granularity in Geographic Information System ontologies below in Section 11.

To make the use of granularity conceptually more clear for the construction of our own ontologies, consider the following example offered by Smith and Brogaard. When someone who is thirsty drinks a glass of water, they might state afterwards that the glass is empty. However, if a government hygiene inspector were to investigate the same glass, they might well come to the conclusion that the glass is far from empty and contains numerous (no doubt contaminated) drops of water. It makes little sense, argue Smith and Brogaard, to argue that one of these points of view is true and the other false—the truth or falsity of the respective descriptions depends on the purposes for which they are made: they are purposeful judgements not merely logical propositions. For the thirsty person, the judgement that the glass is empty is true because the drops found by the health inspector "fall beneath the pertinent granularity" for the thirsty person. Again, this is not an inaccuracy, or a problem due to observational limits; it is rather an indication that "vagueness is not merely a defect of language; it also often facilitates communication without the cumbersome language required to achieve precision" (Bennett 2001*b*). In fact, we would go further—the use of language inherently involves the presupposition of shared functionalities and purposes. 'Empty' does *not* mean 'containing nothing'; it means something more like 'containing nothing very much and in any case nothing useful for the task at hand'. The parameterizing by purposes is an essential component. It has been found in many of the more serious attempts made to probe the meanings that interlocutors really rely on when interacting (cf. Garfinkel 1972) and will form a central pillar of our approach to language, meaning and ontology in the project I1-[OntoSpace] generally.

The mechanism of variable granularity provides access to this kind of flexibility. It is then

> "...the course-grainedness of our partitions which allows us to ignore questions as to the lower-level constituents of the objects foregrounded by our uses of singular terms. This in its turn is what allows such objects to be specified, not precisely,

> but rather in such a way that a range of alternative but nearly identical objects are simultaneously included within their [...] scopes. The unwitting author of the coarse-grained partition does not recognize this 'many' because she is focused, precisely, on those parts and moments of the matters in hand which lie above the pertinent granularity threshold." (Smith & Brogaard 2003)

Granular partitions are defined more precisely and illustrated with some simple examples by Bittner & Smith (2003*a*). Formally Bittner and Smith draw on Rigaux & Scholl's (1995) notion of 'cutting' a tree so that either one takes all of the children of some node or just that node itself. This ensures that the nodes remaining are 'locally' of the same granularity, although more or less detail may be admitted for non-local nodes, i.e., descendents of other sibling nodes in the tree. When the original tree is well-formed in the sense of representing, for example, a partonomy with guaranteedly disjoint nodes, then any cut of the tree is also guaranteed to produce partitions whose members are also pairwise disjoint. This can be seen most easily graphically. So, with respect to the part-decomposition of Europe offered by Bittner & Smith (2003*a*) and reproduced in Figure 19, the following sets are all partitions in the sense they define:

label	partition
g_0	{Europe}
g_1	{Great Britain, Germany}
g_2	{York, London, Scotland, Germany}
g_3	{York, Hyde Park, Soho, Buckingham Palace, Suburbs, Edinburgh, Glasgow, Germany}

Here we can see the interesting property that a particular granular partition does *not* mean that all the (relatively) atomistic elements are of the same 'size'. Pairwise disjointness is taken to be a more basic property and this fits well with the notion of a partition being for some *purpose*: it is typically because of such a purpose that the world is divided into distinguishable entities. Conversely, this is also, in a sense, to define equivalence classes which group entities that are being considered to be indistinguishable for the purpose at hand.

The very fundamental role accorded to partitions is taken up finally in Bittner & Smith's (2001*a*) discussion of partition theory as an alternative to set theory and mereology—this would place it right at the foundations of ontology construction. Each such ontology has a granularity and that granularity pertains to the ways selected for partitioning the entities in reality. The arrangement of such a partition may then be:

> "purely spatial, as in a map, where the relative positions of neighboring cells are determined by the corresponding positions of those portions of geographical reality to which the cells relate. Or it may be determined by a linear ordering, as for example where partitions are determined via quantitative scales reflecting age cohorts or tax brackets or frequency bands. The arrangement may also be determined in more complex (for example hierarchical) ways, as in the case of a partition determined by *kinds* or *concepts* (for example a partition of the animals in your local zoo into *lions, tigers, giraffes, small marsupials,* etc.)." (Smith & Brogaard 2003)

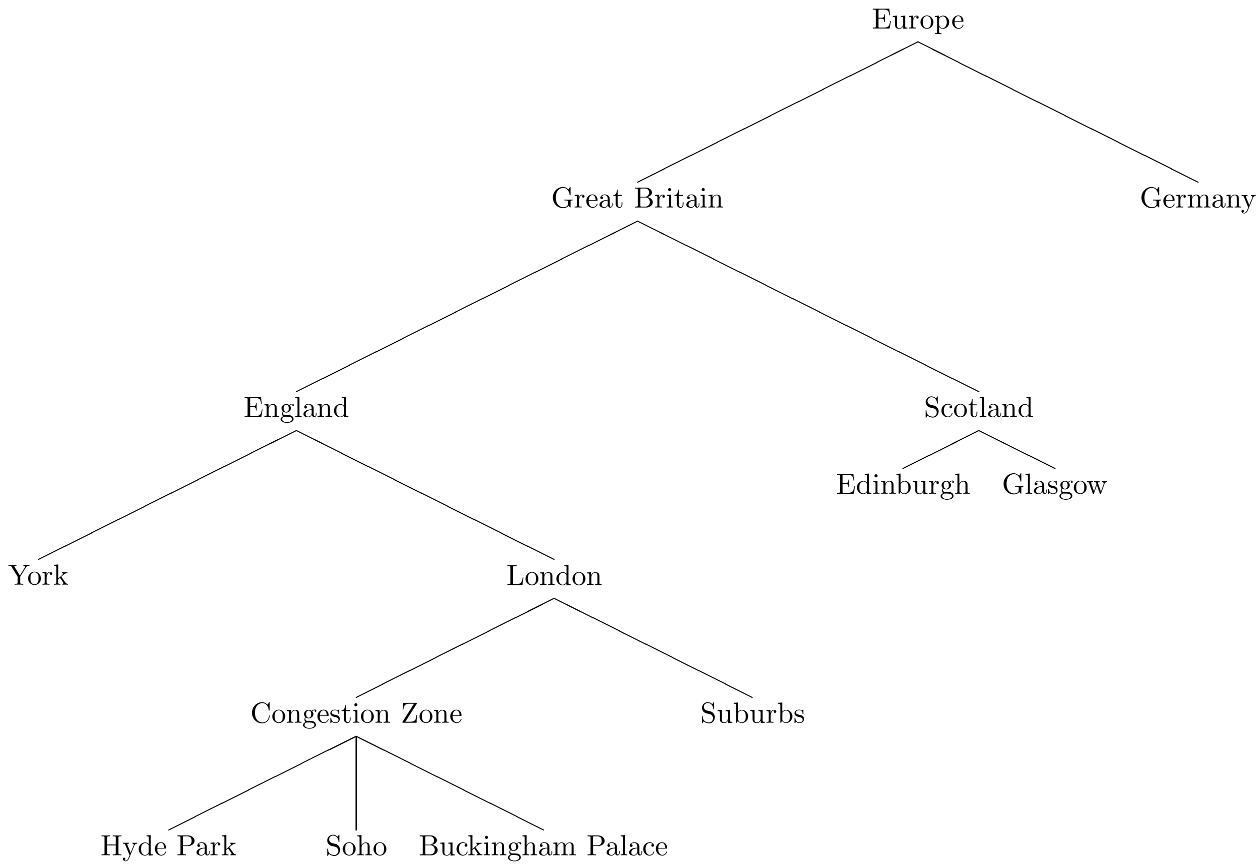

Figure 19: An illustrative partial partonomy of Europe from Bittner and Smith (2003)

We can therefore state that any ontology is itself a partition of reality at a given granularity with a rather complex, generally hierarchical structure. The relations *between* such ontologies then need to be dealt with by formalizing particular trans-ontological relationships in the manner begun by Smith & Grenon (2004).

6.6 Summary and Discussion

Since BFO's aim is to establish the most fundamental upper level structures of reality in general, it does not concern itself with detailed material spatial entities and relations. These are left for domain experts. What BFO does provide is a rather different view of the position of space and spatial entities within a general foundational ontology. In order to proceed, it will be necessary to weigh more exactly the particular position and entities proposed for DOLCE and BFO and to see how they might be extended to include the kind of detail that we see in ontologies such as SUMO and OpenCyc.

Another direction offered is to utilize BFO's general framework for relating separate ontologies, thus taking advantage of the perspectival approach. The separation of SNAP from SPAN entities offers an interesting approach to qualitative reasoning, especially qualitative spatial change as a relation across two different SNAP ontologies. The multi-perspectival approach may also be relevant for the mapping of natural language onto spatial categories, and will be explored in more detail in D3—particularly the notion of granularity dependent on context, for, as Smith & Brogaard (2003) note:

> "The move from an everyday inspection of the glass [drunk empty by the thirsty person] to the more careful inspection involving powerful microscopes amounts to what we can think of as a context switch ... Such switching can be brought about rather easily."

In fact, according to Smith and Brogaard, it can be brought about simply by *talking about* individuals of the granularity required. This flexibility and its relation to linguistic interaction lies at the heart of the empirical investigations being undertaken by I1-[OntoSpace].

7 Qualitative spatial representation with regions

In considering possibilities and requirements for a comprehensive ontological scheme for spatial representation within the SFB, it is also essential to incorporate the very extensive work that has been carried out primarily within the tradition of **qualitative spatial representation**. There are many places of natural contact between the traditions of ontology—general and particular—and specifically spatial representations and formalisms; the goal of conjoining the two is then a natural one. Indeed, the investigation of qualitative spatial representations has been related to the concerns of formal ontology at various points in its history (cf. Egenhofer & Mark 1995, Bennett & Galton 2001). An original motivation for developing representations in this area was also very similar to one of the starting points for the confluence of ontology and artificial intelligence and computational modelling, the 'naive physics' ontology proposed by Hayes (1979): both were seeking an account of our everyday commonsense experience of the world. Qualitative spatial representations under this view are then also candidates for 'realistic' models of space—i.e., views that correspond to, or pick out, properties of spatial configurations of a similar kind to those attended to by humans. And it is in this sense, therefore, that we seek to combine the various strands of spatial representation and reasoning research that have been pursued hitherto.

This section cannot provide anything like an exhaustive overview of the area of qualitative spatial representations (for something approaching this see, for example, Cohn & Hazariki (2001)). Our goal is rather to introduce the basic notions to those who are unfamiliar with them and to present enough of a view of the field to clarify what implications there might be for ontological modelling. Our final goal will be to consider how the two directions can be combined most sensibly.

While it is clear that human spatial processing draws on a variety of qualitative information rather than quantitative information such as precise distances or angles, just which kinds of qualitative information is not yet known. Whereas all such kinds of information relevant for physical objects, including location, relations between objects in space, shape, physical properties such as colors, etc., may be addressed qualitatively and play important roles in the construction of any general ontology, qualitative spatial representations have addressed particularly the spatial relationships holding between objects of various kinds. The basic topological relationships suggested by such accounts are then candidates for 'natural' cognitive representations and are accordingly explored formally, computationally and psychologically.

Incorporating the results of qualitative spatial representation is thus a necessary step to take for any detailed spatial ontology. It can also be seen as offering a significant extension of the computational tools available for dealing with spatial representations. Much qualitative spatial representation work has been carried out paying close attention to formal complexity issues and the requirements for reasoning—for example, within visual processing and cognitive robotics—and so particular techniques have been developed for working with sets of spatial constraints of various kinds; we return to these techniques below.

7.1 The Region Connection Calculus – RCC

As a starting point for discussion we take two independently developed views of spatial relationships: the Region Connection Calculus, RCC, proposed by Randell, Cui & Cohn (1992) and the set of topological constraints proposed by Egenhofer (1991). These involve stating basic spatial relationships that may hold between spatial entities and working out ways of both reasoning with them and applying them to complex spatial configurations. Although Egenhofer describes his relations in purely topological terms and draws on set theory (regions as sets of points) for definitions, while Randell *et al.* draw on a topology of regions with spatial parts and start from the connection relation alone, there are clear similarities between them. Their ontological commitments are, however, somewhat different concerning the particular kinds of spatial objects assumed. The relations proposed are set out in Table 3 with the names employed in both approaches; because there are eight relations, the term **RCC-8**

Egenhofer	RCC
disjoint	disconnected (DC)
meet	external connection (EC)
overlap	partial overlap (PO)
equal	equal (EQ)
coveredBy	tangential proper part (TPP)
inside	non-tangential proper part (NTPP)
covers	tangential proper part inverse (TPP^{-1})
contains	non-tangential proper part inverse ($NTPP^{-1}$)

Table 3: The standard eight 'base relations' of RCC-8 and similar calculi

is commonly employed and we will use this here regardless of the precise approach described.

We saw in Section 2 above that a mereotopology illustrated with two dimensional regular spatial regions and their possible inter-relationships also gives a natural set of eight relations describing whether two regions are touching or overlapping in various ways, or whether they are distinct (cf. Figure 3). It is therefore no coincidence that the relations used in the qualitative spatial reasoning tradition show this similarity. When formalizing them, both approaches make central use of the topological relationship of **connection** and typically begin their accounts with an enumeration of the distinct ways in which spatial entities can be related spatially: thus leading to these standard eight relations.

There are, however, a number of ways of formalizing this view of spatial entities and their relations. While the mereotopological approach common in ontology begins by axiomatizing the relations and their properties in a first-order logical language (in terms of connection and parthood), the computational behavior of such a formalization is not good because of its expressiveness; full first-order formalizations are still often unsuitable for driving reasoning—Table 4 shows an axiomatization provided by Bennett of the account of Randell, Cui & Cohn (1992). The full first-order theory of RCC inherits undecidability and so various, more constrained, adaptations of the full theory have been defined. Within qualitative spatial reasoning, therefore, formalizations of spatial relationships that draw on the formal properties of other mathematical accounts is a very active area. By this means, one attempts to achieve more attractive computational properties that are more amenable to practical reasoning tasks. This also appears motivated cognitively in that certain tasks are readily performable by humans and others less so; such evidence can further constrain the kinds of modelling decisions that may eventually be made in formulating a realistic ontology of space.

The most widely used approach to applying such spatial calculi has been to employ a constraint propagation technique known as **path consistency**. This approach was originally developed for working with the temporal interval representation proposed by Allen (1984); RCC and its relations are seen as the spatial correlates of Allen's approach to time and there are close mathematical affinities between the family of spatial representations and the interval calculus of Allen that support this. Path consistency involves establishing whether some collection of statements involving the relations defined collectively define a possible configuration—if they are, then they can be said to define consistent paths. To compute such

Connection axioms

$$(\forall x \forall y (C(x,y) \rightarrow C(y,x)))$$

$$(\forall x C(x,x))$$

Definitions

$$(\forall x \forall y (DC(x,y) \leftrightarrow \neg C(xy,y)))$$

$$(\forall x \forall y (P(x,y) \leftrightarrow (\forall z (C(z,x) \rightarrow C(z,y)))))$$

$$(\forall x \forall y (PP(x,y) \leftrightarrow (P(x,y) \wedge \neg P(y,x))))$$

$$(\forall x \forall y (x = y \leftrightarrow (P(x,y) \wedge P(y,x))))$$

$$(\forall x \forall y (O(x,y) \leftrightarrow (\exists z (P(z,x) \wedge P(z,y)))))$$

$$(\forall x \forall y (PO(x,y) \leftrightarrow (O(x,y) \wedge \neg P(x,y) \wedge \neg P(y,x))))$$

$$(\forall x \forall y (DR(x,y) \leftrightarrow \neg O(x,y)))$$

$$(\forall x \forall y (EC(x,y) \leftrightarrow (C(x,y) \wedge \neg O(x,y))))$$

$$(\forall x \forall y (TPP(x,y) \leftrightarrow (PP(x,y) \wedge (\exists z (EC(z,x) \wedge EC(z,y))))))$$

$$(\forall x \forall y (NTPP(x,y) \leftrightarrow (PP(x,y) \wedge \neg (\exists z (EC(z,x) \wedge EC(z,y))))))$$

$$(\forall x \forall y (Pi(x,y) \leftrightarrow P(y,x)))$$

$$(\forall x \forall y (PPi(x,y) \leftrightarrow PP(y,x)))$$

$$(\forall x \forall y (TPPi(x,y) \leftrightarrow TPP(y,x)))$$

$$(\forall x \forall y (NTPPi(x,y) \leftrightarrow NTPP(y,x)))$$

Additional axiom

$$(\forall x (\exists y NTPP(y,x)))$$

Table 4: Extract of the axiomatization of RCC given in Randell *et al.* (1992) as provided by Bennett at http://www.comp.leeds.ac.uk/qsr/rcc.html

$\circ_w$		C					
	DR		O				
				PP		PP˘	
	DC	EC	PO	TPP	NTPP	TPP˘	NTPP˘
DC	1	DR,PO,PP	DR,PO,PP	DR,PO,PP	DR,PO,PP	DC	DC
EC	DR,PO,PP˘	1',DR,PO, TPP TPP˘	DR,PO,PP	EC,PO,PP	PO,PP	DR	DC
PO	DR,PO,PP˘	DR,PO,PP˘	1	PO,PP	PO,PP	DR,PO,PP˘	DR,PO,PP˘
TPP	DC	DR	DR,PO,PP	PP	NTPP	1',DR,PO, TPP,TPP˘	DR,PO, PP˘
NTPP	DC	DC	DR,PO,PP	NTPP	NTPP	DR,PO,PP	1
TPP˘	DR,PO,PP˘	EC,PO,PP˘	PO,PP˘	1',PO, TPP,TPP˘	PO,PP	PP˘	NTPP˘
NTPP˘	DR,PO,PP˘	PO,PP˘	PO,PP˘	PO,PP˘	$O \cup 1'$	NTPP˘	NTPP˘

Table 5: A composition table for RCC-8 (taken from Düntsch *et al.,* 1998)

paths it is necessary to define the *composition* of relations, since a path from x to y to z is equivalent to the composition of the relations between x and y and between y and z.

Therefore, employing the path consistency algorithm is made considerably easier by forming **composition tables** that show the results of composing spatial relations (Freksa 1992*a*): here again it is particularly important that the relations have received a clear semantics; constructing composition tables is equivalent to solving a set of theorems in the calculi and can be difficult. Composition table construction is therefore also sometimes carried out using automatic theorem provers (cf. Randell, Cohn & Cui 1992, Eschenbach 2001).

Composition tables can be provided for sets of relations that are jointly exclusive and pairwise disjoint (commonly abbreviated as **JEPD**). Such sets are generally termed **base relations**; the eight relations of RCC-8 form such a set. Table 5 shows the corresponding composition table for these eight base relations (taken from Düntsch, Wang & McCloskey (1998) and Randell, Cui & Cohn (1992)). Each entry in the table shows the results of composing two assertions concerning spatial regions. For example, the result of the following two assertions:

$$x\ NTPP\ y \wedge y\ EC\ z$$

is given by the table under the entry in the NTPP-row and the EC-column; i.e.:

$$x\ DC\ z$$

This is shown graphically in Figure 20.

In general, the results of any composition may require more than one possible relation; these disjunctions are simply listed in the corresponding table entry. It can be a considerable task to prove all of the compositions present in such a table, but once they are present they simplify subsequent reasoning substantially as the results of these proofs have in effect been 'cached'.

There has been considerable work on the computational complexity of both the temporal calculus and the various brands of specialized spatial representations proposed (cf., particularly, Nebel 1995, Renz & Nebel 1998, Renz & Nebel 1999, Scivos & Nebel 2001). The

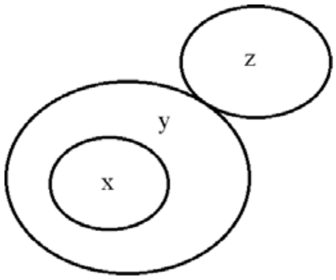

Figure 20: Example composition of NTTP (non-tangential proper part) with EC (external connection)

results of this continuing and very active area of research allow a precise characterization of spatial problems and representations according to applicable algorithms. Maintaining such complexity results may well constitute important 'meta'-information concerning components of a full spatial ontology.

The primary consideration in such complexity analyses can usually be reduced to the question of **satisfiability**—that is, when there is a description of a spatial configuration drawn from the relation set, is it possible to find a real spatial configuration that conforms to that description. If satisfiability can be checked, then descriptions can be used for reasoning. So far it is known that reasoning in both RCC-8 and RCC-5 (see below) is NP-hard. Moreover, satisfiability for RCC-8 is decidable but NP-complete (cf. Bennett, Wolter & Zakharyaschev 2002, Renz & Nebel 1999).

There do exist substantial but more restricted fragments of these calculi that have better computational properties—for example, Renz & Nebel (1999) present a 'maximal fragment' of RCC-8 that includes all the base relations but which is still tractible—i.e., satisfiable in polynomial time. Such fragments are formed by considering subsets of the full set of possible relations that may hold between regions according to the calculus. As the composition table above showed, the result of composing relations may require disjunctions of relations—i.e., a more complex relation is formed out of the base relations. The full set of possible relations formable by combination of the RCC-8 base relations consists then of 256 relations. Restricting the relations so that only subsets of these 256 are available can substantially improve the computational properties of the resulting calculus.

The practical application of these fragments is still being studied: it has now been shown (Renz & Nebel 1998), for example, that they can be used to speed up backtracking during constraint satisfaction. This works as follows: some possible relation posited to hold between regions may need to be refined further during the search for some solution; the calculus fragment adopted provides the possibilities that are available for such a decomposition—analogously to decomposing a number into a product of prime numbers. The maximal decomposition is naturally provided by the eight base relations of RCC-8: this then provides the largest set of alternative paths to be checked for solutions. The smaller fragments can improve this situation by providing combinations of relations that similarly serve to decompose the complex relation being decomposed, but drawing on fewer relations than the eight base relations. This then results in fewer possible paths to be followed, which consequently brings a (sometimes con-

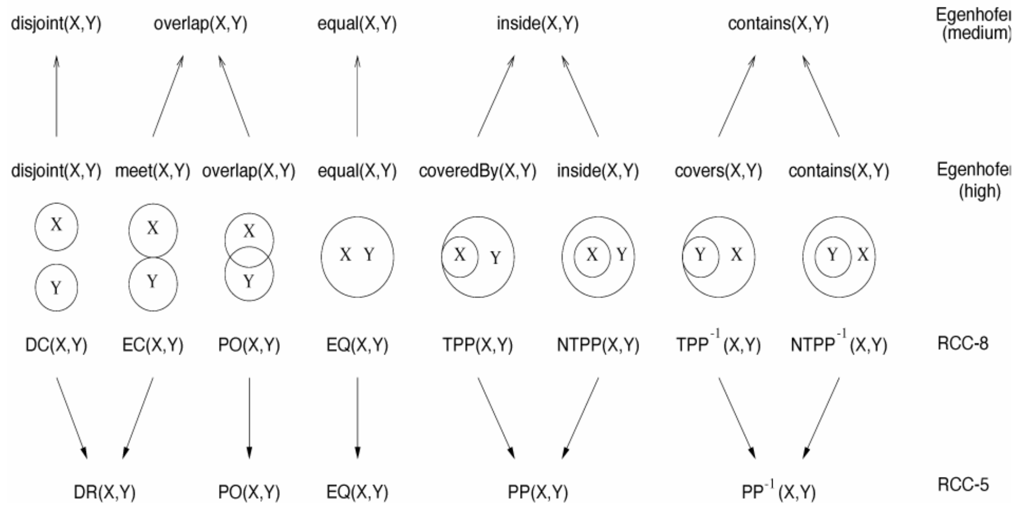

Figure 21: The basic eight spatial relations of RCC-8 and Egenhofer, plus their suggested simplifications as RCC-5 and the medium resolution set

siderable) improvement in reasoning performance. Obviously, since not all the base relations are being adopted there will always be some complex relations that are not within the range covered: for this reason the resulting fragments will necessarily include fewer than the full 256 relations. Renz & Nebel's (1999) fragment ($\widehat{\mathcal{H}_8}$) includes 148 relations; two further fragments have been found by Renz (1999) containing 158 and 160 relations. These three are the only maximal tractable fragments for RCC-8. The restriction that a problem must be expressed using only the relations defined in the fragment is still a considerable one; explorations of these fragments and their benefits continue—already clear, however, is that many kinds of 'difficult' cases are more amenable to automated reasoning than was previously suspected.

One can also attempt to find a useful smaller set of topological relations that has more attractive computational properties. This latter course has been taken for both the RCC and Egenhofer-style of approaches: the former is illustrated in RCC-5 (Bennett 1994), a region connection calculus with 5 base relations, the latter in the 'medium resolution' set of Grigni, Papadias & Papadimitriou (1995). These two simpler cases are set out in Figure 21 in a form taken from Knauff, Rauh & Renz (1997); this representation usefully shows the precise ways in which the simpler sets of relations are related to the full set of eight. Interestingly, Knauff et al. (1997) report that they have found no psychological motivation for the simpler sets of relations, although a slight preference may be for the grouping exhibited for RCC-5. Moreover, it has been seriously questioned whether the simpler set of relations is sufficient for realistic problems. Accordingly, there have been studies of other selections of base relations for region connection calculi; Düntsch et al. (1998), for example, also present composition tables for RCC-7 and RCC-10.

Finally, combining the purely topological descriptions with other sources of constraint, such as metrical information such as distance and size, can also change the computational properties of the descriptions used substantially. For example, adding appropriate size constraints

(a) RCC-5 (b) RCC-8

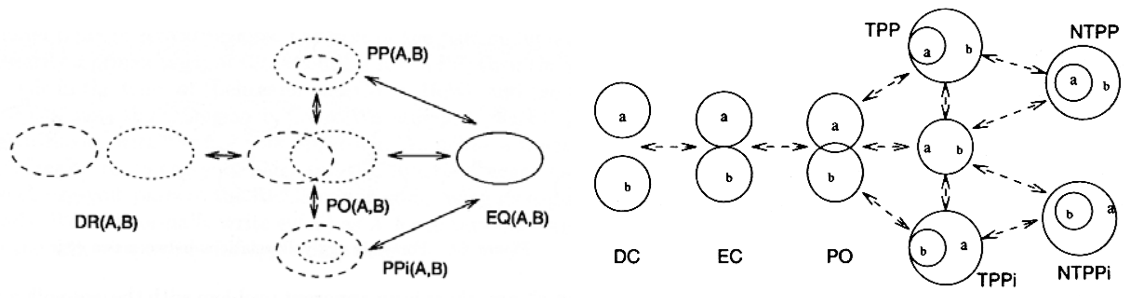

Figure 22: Conceptual neighbors shown as possible continuous transitions between RCC relations for RCC-5 and RCC-8 (taken from Cohn *et al.* 1997; Figure 21, p127 and Figure 10, p112)

between regions (e.g., keeping their size constant) renders all reasoning in RCC-8 polynomial in complexity (Gerevini & Renz 1998).

7.2 Movement

A further consideration that has been addressed within RCC is that of changing spatial relationships: that is, when regions move. We will reserve particular attention to this aspect of our ontologies for a later date, but note here in passing that this will also obviously require attention for various research topics within the SFB and its engagement with space. Within RCC formalizations have been proposed that emphasize the notion of **continuous** change. This can be represented well in terms of **conceptual neighborhood** as originally proposed for temporal relations by Freksa (1991). This is best represented diagrammatically as shown in Figure 22.

An interesting result also dicussed by Cohn, Bennett, Gooday & Gotts (1997) relates composition tables and conceptual neighborhoods: the disjunctions arising from the composition of two relations necessarily lie within a conceptual neighborhood; that is, all the members of a resulting disjunction will be connected according to neighborhoods such as those defined by the diagrams in Figure 22. This can be used as an aid for deriving composition table entries since whenever two relations have been found to belong to the disjunction, one knows that the relations lying between these within a neighborhood are also necessary.

We return to a further view of movement that is currently under active development as a topic in its own right in Section 10 below.

7.3 Dimensionalities

The region connection calculus and the straightforward cases of mereotopology addressed above only relate objects that are alike: in the case of RCC, regions. There have also been approaches that concern themselves with relations among different kinds of objects. These are drawn mostly from work on Geographic Information Systems where a restriction to objects that all similar dimensionally has been considered less plausible.

Egenhofer & Herring (1991), for example, define a complex set of topological interrelationships that bring together regions, lines and points in a single classification system. This has later been refined with metrical information of various kinds in order to provide a fine-grained characterization of spatial configurations aimed at providing information retrieval within Geographic Information Systems. Rashid, Sharif, Egenhofer & Mark (1998) take this further and apply it as a basis for defining the semantics of natural language terms involving spatial relationships—we will return to this aspect of the classification in our deliverables concerning spatial language and ontology. Here we indicate the kind of spatial relationships captured by repeating the so-called **9-intersection** diagrams used by Egenhofer & Herring (1991): this is shown in Figure 23 and gives a graphic illustration of the very different kinds of configurations captured compared with those we have seen so far. Within this framework the distinct topological possibilities are given in terms of the possible intersections involving interiors, boundaries and complements of the regions (sets of points) that are being related.

A further approach combining distinct kinds of spatial entities is that of Isli, Museros, Cabedo, Barkowsky & Moratz (2000). Here there are distinct kinds of connection relations for point-like, linear and areal features. This builds on the approach of Egenhofer.

In terms of an ontological representation, one could place these kinds of configurations as basic distinctions within the spatial ontology module along with the supporting concepts of boundaries and interiors. The appropriateness of such an organization both for spatial representation alone and for representations of space for communication in natural languages will need to be addressed.

Here a particularly relevant study is that of Egenhofer & Rodríguez (1999), who propose a representation of rooms and some everyday objects in terms of a relation algebra. The spatial entities considered are mixed ontologically in that they contain both surfaces and containers—i.e., entities of fundamentally different dimensionalities. Here, rather than beginning with geometric or even spatial considerations, Egenhofer and Rodríguez instead adopt the Gibsonian notion of **affordances** (Gibson 1977) and formalize what *can be done* with surfaces and containers. Thus the formalization first specifies that entities can be placed on surfaces and can later be removed from them, that entities can be placed in containers and moved around in them, etc. Significantly, Egenhofer and Rodríguez claim that the purely spatial relationships then fall out of their formalization automatically without separate formalization. They also relate this to the use of natural language expressions for describing the spatial situations encountered. As we shall see in our deliverables concerning spatial language and the requirements that it raises for ontological modelling, there is much in this approach that echoes our own direction of development.[2]

[2]See also Bateman (to appear) for some preliminary remarks.

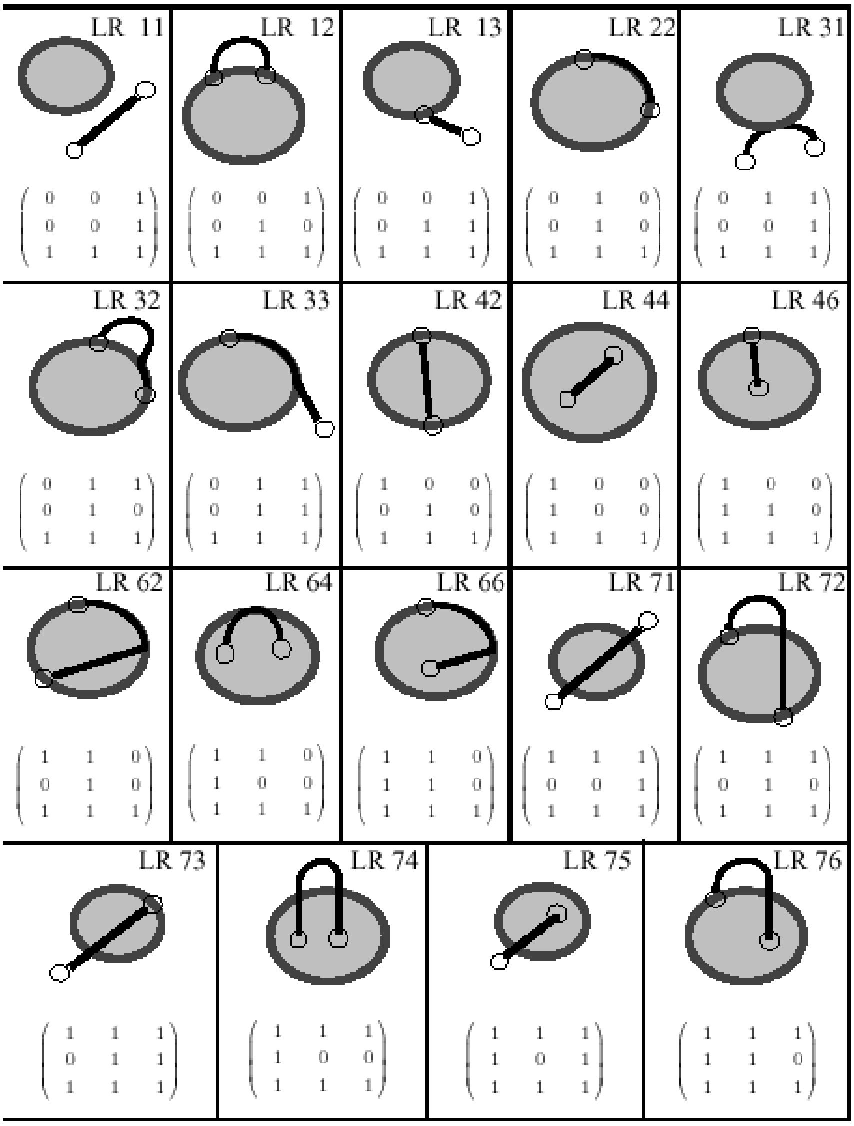

Figure 23: Geometric interpretations of the 19 line-relation relations drawn by the 9-intersection model of Egenhofer and Herring (1991)

7.4 Conclusions

One primary motivation for exploring the kinds of relation sets seen here in close detail is that their mathematical properties can be used to characterize their computational properties—i.e., to answer questions concerning whether descriptions employing them can be used effectively (or at all!) for reasoning. In many respects, however, these descriptions (and as we shall see below, even much simpler specifications) are already computationally complex—although there are particular processing mechanisms that can be employed in order to employ them for reasoning. The full characterization of RCC as a first-order theory with one primitive of 'connectedness' is already too expressive computationally; it is known to be undecidable. So we have seen that there is an active research direction to find more computationally attractive representations.

Although such issues of computational techniques for working with spatial representations have not traditionally played an important role in ontological discussions, we will, however, need to consider them far more closely if parameterized applications of the spatial ontology developed within the SFB are to be employed. Although our purpose in the present document is restricted to providing a basis for considering these diverse traditions within a single overarching ontological perspective, the application of the concepts developed will need to make contact with actual processing techniques as developed in the individual research areas involved. The extent to which such specifications can be used by theorem provers will be a natural area for further investigation when such an approach is considered for our ontological specifications in general.

The axiomatizations available for all versions of RCC make it relatively clear how this kind of specification could relate to the other kinds of ontology presented above—particularly those organized around subsumption lattices. The axioms for each node in this lattice fragment may be drawn from the standard RCC axiomatizations. Although here again there are various possibilities; Bennett et al. (2002) shows how RCC can be expressed in terms of a modal logic, which is a substantially different (but equivalent) kind of representation. This is potentially useful in the context of ontology specifications because of the correspondence between modal logic and description logic mentioned in Deliverable D1; because of this a modal logic representation makes it possible to capture RCC-like representations in a form amenable to reasoning with knowledge representation systems such as Racer and FaCT; some of these issues are the concern of the SFB project R4-[LogoSpace]. There have also been attempts to represent fragments of the region calculi directly in description logics (cf. Schulz, Hahn & Romacker 2000, Schulz & Hahn 2001); this is not a full representation, although the authors report that a 'large fragment' of RCC-5 is covered. Such issues of correspondence across related formalizations will need to be addressed more fully in the formal specification of the ontologies for the SFB: cf. Project I4-[SPIN].

As the simplification of RCC-8 to RCC-5 suggests, we can also relate the various sets of relations according to degree of generality/specificity—which in turn may also bring implications for other areas of a spatial ontology, for example, whether 'boundaries' are present in the account or not. Clearly both of the simplified relation sets in Figure 21 make it impossible, for example, to state whether a related object is situated 'at a boundary' or not. Figure 24 shows the base relations organized in a subsumption lattice. Similarly, extensions to RCC-8

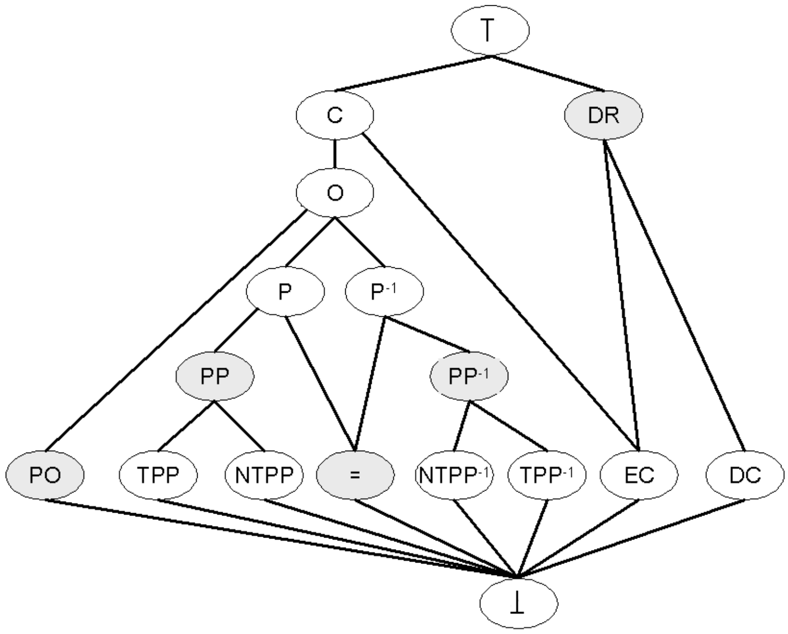

Figure 24: Subsumption lattice for the basic eight spatial relations of RCC-8 and the five of RCC-5 (shown in grey)

are commonly formed by decomposing the base relations further: the computational properties of variants involving 10, 15 and 23 relations, i.e., RCC-10, RCC-15 and RCC-23 have all received attention. The existence of distinguished subalgebra—such as the maximal tractable subsets for RCC-8—will also need a place as these have been shown by Renz & Nebel (1999) to bring substantial benefits.

We can see in the lists of spatial relations addressed in this section similar relations to those that have been adopted in several of the ontologies that we have discussed so far. However, what is less often carried over to the ontology versions is the precise mathematical characterization of this set of relations along with their intrinsic modularity. We will need to ensure that these valuable properties are maintained, even when they are placed alongside the other information maintained in a more general ontology—such as general patterns of subsumption relationships (as in SUMO) or of part-whole relationships (mereological). It is also worth citing the conclusions of Cohn & Hazariki (2001) at this point:

> "...it is unlikely that a single universal spatial representation language will emerge—rather, the best we can hope for is that the field will develop a library of representational and reasoning devices and some criteria for their successful application. Moreover, ...qualitative and quantitative reasoning are complementary techniques and research is needed to ensure they can be integrated—for example by developing reliable ways of translating between the two kinds of formalisms."

Moving between different kinds of specifications of a problem area must also be one of the tasks that an adequate spatial ontology takes up. However, in constructing relationships between, for example, RCC and traditional ontologies, both formal and AI-oriented, such as Cyc and SUMO, a number of questions arise concerning the degree of flexibility required or allowed.

The RCC-approach is, deliberately, extremely general—relying in its pure form only on the Connection relation. However, when placed in a larger perspective, this relation can be seen to cover a range of not always interchangeable perspectives. Borgo et al. (1996*a*) point out, for example, that in their framework the connection relation of RCC can be interpreted "as denoting strong connection, line-connection, point-connection or a combination of these"; their own formalization allows a more precise characterization of its adopted primitive. They conclude:

> "This freedom in the interpretation could be an advantage of the RCC approach, in the sense that the theory is apt to capture a very general notion of connection, which may be useful for various purposes. However, the theory appears too weak for a formal characterization of space in its present state." (Borgo et al. 1996*a*)

They are then particularly concerned with an additional element of *cognitive plausibility* as motivation and restriction for the properties for the primitives adopted. This is what leads them away from approaches that adopt "debatable" distinctions such as open and closed regions, to place a greater weight on the ontological level of morphology for physical objects—which they capture via their congruence relation.

Similar concerns arise in Grenon's (2003) attempt to embed RCC within Cyc.

> "For Cyc, objects in space (and physical events) are primitives. If RCC is generalizable to SpatialThing and if regions (SpaceRegion) are but a subtype of spatial things, it would be redundant and sub-optimal to develop the theory on regions only. Yet, RCC takes regions as primitives. Pushed to the extreme, this latter position leads to defining spatial objects in terms of regions. There would be only one substance, space. An object would be no more than a qualitative singularity in space. Such an eliminatist view (doing away with the primitivity of entities in space) is not conceivable for Cyc's upper-level ontology. The compromise still appears to be straightforward generalization of RCC's notions at the level of SpatialThing in Cyc." (Grenon 2003)

Grenon then investigates what occurs if the connection relation of RCC be interpreted not as originally given in the RCC literature but as a connection relation holding over objects—i.e., the status of the primitive entities assumed is changed from the neutral characterization of RCC to the object-centered perspective embodied in Cyc. As we saw in Section 4, Cyc offers a number of finely discriminating 'touching' relationships that are distinguished not only by degree of touching but also by the kinds of objects that they are to hold over. The connection relation of RCC might then be assimilated to the higher parts of this hierarchy of touching relations in Cyc. Grenon argues, however, that such a reinterpretation is problematic and what is necessary is "an interpretation according to which RCC was a theory of regions of substantival space." To accomodate this, he then has to introduce new components to Cyc dealing with pure spatial regions within which objects may be located. Presumably a similar step would need to be taken if working with SUMO also as there, as we have seen, there is also a rather unclear distinction between objects, spatial relations and spatial regions. Grenon also concludes by stating that other interpretations of RCC may be possible, but would require an ontology to provide a more general mereological foundation than that available in Cyc.

Grenon raises similar concerns with the notion of region as such:

> "However, what is arguably not trivial is precisely what regions are, or, in other words, what their ontological status is: are they dependent or independent ontologically, that is, for their existence, on other entities which are not themselves spatial regions? In fact, depending on whether the term 'region' refers to a part of a substantival space or to a portion of a relational space, the answer and the ensuing interpretation of the theory would be significantly different ontologically speaking. In other words, the alleged reality that the theory intends to capture is not that evident." (Grenon 2003)

Thus, the notion of spatial region in RCC is actually ontologically unclear and is another place where RCC does *not* commit itself.

The limitations of RCC alone are well recognized and have been addressed within the RCC-tradition also. One direction for such an extension entirely within the spirit of RCC is Bennett's region-based work, to which we turn in Section 9 below. Another is the extension by a **convex hull** operation. Davis, Gotts & Cohn (1999) show that this taken together with RCC-8 provides a very powerful system that can define relations (such as *just inside* or *just outside*) which are capable of distinguishing "any two regions not related by an affine transformation." There are then very many spatial relations that may then be defined with very few primitives: just the two of connection and convex hull. However, as Cohn & Hazariki (2001) then ask:

> "The question arises: when to stop? In Cohn, Randell & Cui (1995) we propose some criteria based on computational and predictive properties of the representation, but ultimately it must be a domain specific question: certain distinctions will only be useful for certain domains, but for these domains they may be crucial. ... the significance of qualitative distinctions depends largely on their relevance to the behaviour being modeled." (Cohn & Hazariki 2001)

It is precisely to narrow down possible answers to this question that relating the RCC approach to a broader notion of ontology is potentially useful.

This discussion shows well some of the issues that arise when we begin to bring together the qualitative spatial approach and ontology proper. For the mathematical or formal definition of RCC, what exactly a region is is of little import. But if we pursue an ontological foundation with the kind of rich 'horizontal' linking of a DOLCE-style of axiomatization, we need to know more.

8 Directedness and Orientation

A further important aspect of spatial representations that has not been dealt with in the qualitative accounts that primarily relate regions is **directedness**: that is, not only can a spatial representation be concerned with static regions or lines that stand in some kind of relation, but these regions or lines may additionally possess some kind of directionality. Since one area of investigation within many projects of the SFB involves navigation, issues of directionality cannot be ignored. Navigation problems can be modelled in terms of an object moving in some direction, followed by turns by some qualitatively described amounts in qualitatively specified directions. Problem solving might indicate whether the moving object has gone in a circle, i.e., arrived at its starting point, etc. The aim here is then to find plausible and sufficient notions of directionality that correspond to human concepts of reasoning about navigation and paths. Directionality, both of moving objects and as intrinsic to particular kinds of objects—i.e., certain object located in space also have 'fronts', 'backs', 'tops', 'bottoms', etc.—will therefore need to captured and represented appropriately in our general spatial ontology.

Directions of various kinds were mentioned in the SUMO and OpenCyc ontologies above, but there have also been qualitative spatial calculi proposed for dealing with this kind of information. One way in which this has been approached is by adding **orientation** to the spatial regions represented; we present here some of the most well-known examples of this—again, for references to further examples, see Cohn & Hazariki (2001)—and then summarize their import for our ontological considerations as a whole.

8.1 Cardinal directions

Perhaps most obviously from the view of navigation are considerations of reasoning in terms of the cardinal directions: north, south, east and west, or finer classifications including north-east, north-west, etc. Such direction and descriptions of routes and locations are commonly found in so-called 'large-scale spaces', spaces that are too large to see all at one go and for which a composite map needs to be constructed.

Reasoning with formal descriptions involving cardinal directions has been approached in a similar way to the connection calculi described above: i.e., in terms of defining relations between directions that can be composed into longer path descriptions. Given such a formalization, standard reasoning techniques can be employed in order to find models corresponding to the description set or to demonstrate inconsistency. The properties of such descriptions have been addressed by, for example, Frank (1991), Frank (1992), Frank (1996), Lizogat (1998) and others.

Taking Frank (1996) as illustrative of the approach, cardinal directions can be captured in terms of an algebra with two operations applicable to the cardinal directions adopted: the first operation simply reverses the direction of travel, the second combines two segments of a path. The directions themselves can be defined in a number of ways. For example, one way is simply to divide up the two-dimensional space proportionally to the number of cardinal points adopted; another is based on projections. Frank (1996) shows that the latter semantics

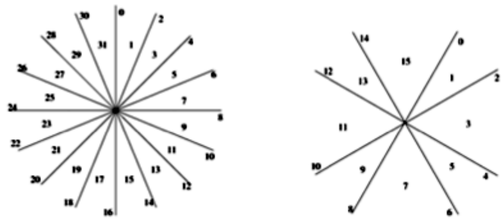

(a) $\mathcal{STAR}_8(0)$ (b) $\mathcal{STAR}_4[30,60,120,150](30)$

Figure 25: Two STAR calculi from Renz and Mitra (2004)

provides more exact and definite results for path descriptions and localization.

Reasoning with such systems is, nevertheless, somewhat limited and it is clear that to be effective further information, such as metrical information, often needs to be provided (cf. Clementini, Felici & Hernandez 1997).

More recently, several generalizations of these cardinal direction calculi have been proposed. Renz & Mitra (2004), for example, develop a $\mathcal{STAR}$-calculus in which it is possible to specify arbitrary angles between the adopted 'cardinal' directions. Two examples from Renz and Mitra are shown in Figure 25. Here the lefthand configuration describes a calculus in which the plane is divided by 8 axes, forming equal angles, and aligned with respect to a global orientaton in the plane. The righthand configuration in contrast to this, is made up of 4 lines, forming angles of 30, 60, 120 and 150 with respect to the reference direction. These configurations define regions which are numbered in a specific way: the numbering is formed with respect to the first identified half-axis.

An interesting application of such a generalized view of directionality is that the regions defined can be made to conform more to qualitative regions identified cognitively. For example, the notions of 'front', 'back', 'left' and 'right' do not seem to divide the plane into equal areas: it is often claimed that the 'back' region is rather more extensive than the 'front' region. Allowing arbitrary angles in the calculus provides a mechanism for this, although many of the precise implications of this for reasoning remain very much current research tasks.

8.2 Double-cross calculus

A further scheme not defined in terms of cardinal directions but instead in terms of relative changes in direction over a path is the **double-cross** calculus introduced by Freksa (1992*b*). The primitives of this description consist of triples of points a, b, c classified according to where c lies with respect to the directed segment a to b. This can naturally be conceived as three points along a navigation path—the moving agent goes from a to b and then makes a turn in the direction of c.

The qualitative relations defined for the direction of c then involve two components: (i) whether the movement is forwards of a line drawn perpendicular to ab and through b, on that line, between that line and a line drawn perpendicularly to ab and through a, or behind that line; and (ii) whether the movement is in a line with ab (forwards or backwards) or to the left

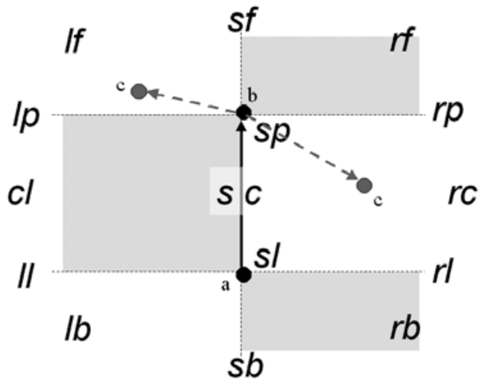

Figure 26: The basic relations of the double-cross calculus of Freksa (1992)

or right of *ab*. This set of relations is summarized graphically in Figure 26.

The resulting calculus can also be used in ways similar to the connection calculus, although an extra complication is provided by the fact that the defined relations now relate three points rather than the two of RCC. Further, more refined, qualitative decompositions can easily be imagined; the question is whether there is any particular evidence from any source that some decomposition is 'better' for humans than others.

It is also known that reasoning in the double-cross calculus is NP-hard and in PSPACE (Scivos & Nebel 2001); which means that, despite the relative simplicity of the descriptions allowed, reasoning within the framework still presents a considerable challenge.

8.3 Dipole calculus

Moratz, Renz & Wolter (2000) define a further calculus centered around directed line segments, or **dipoles**. They are intended to be used for objects with an inherent orientation; taking such objects simply as regions does not do justice to how such objects are generally perceived and employed by humans interacting with them. Moratz et al. (2000) describe a set of relations that can relate dipoles and which are jointly exhaustive and pairwise disjoint in exactly the same way as done for RCC-8 and the other comparable descriptions. This then allows the same techniques for reasoning to be employed. The 24 base relations defined for the dipole calculus are set out in Figure 27. Here we can see that the basic spatial entities related in this account are *pairs* of arrows, or directed line segments. The result is the dipole relation algebra $\mathcal{DRA}$.

Naturally the aim was to keep the number of relations admitted to the base set sufficiently small so as to support theorem proving while still covering the spatial configurations that can occur with sufficient granularity as to be useful. Moratz et al. (2000) show, however, that reasoning with their dipole calculus is NP-hard.

It may be possible to find variations on the dipole calculus that are more supportive of reasoning; Moratz *et al.* note that Isli & Cohn (1998) define a tractable fragment involving

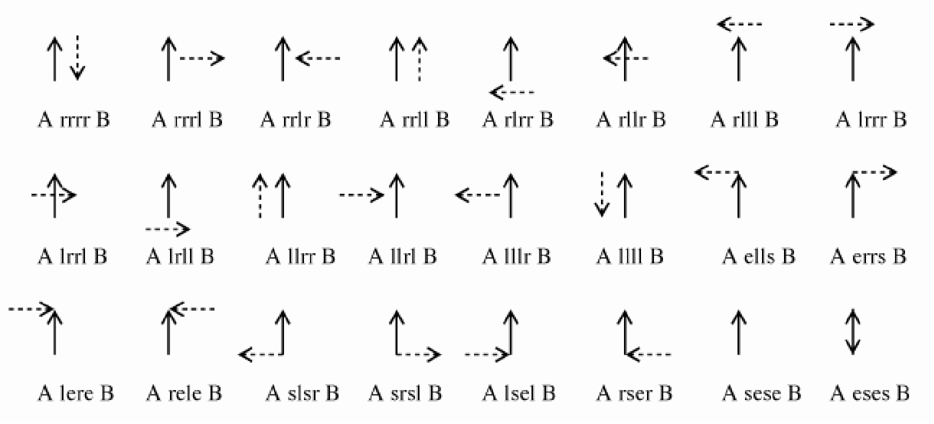

Figure 27: The 24 basic relations of the dipole calculus of Moratz *et al.* (2000)

orientation, although this also needs to relate three points (as does the double-cross calculus) instead of their two. Needing to relate three points instead of two presents certain difficulties for the construction of composition tables for reasoning.

8.4 The Oriented Point Relation Algebra

With the increasing number of spatial calculi on offer, there are accordingly attempts to provide generic frameworks that can mediate, or relate, between the individual calculi. One such direction is that proposed by Moratz, Dylla & Frommberger (2005) with the **Oriented Point Relation Algebra** $\mathcal{OPRA}_{\mathtt{m}}$. This calculus generalizes previous approaches to oriented lines, such as the dipole calculus, and positions with respect to orienting lines, such as the double-cross calculus, to consider oriented points: the relations specified then hold between points on the 2D-plane plus a direction. The calculus also, as with the $\mathcal{STAR}$-calculus, supports a parameterized granularity that imposes a certain specified set of possible qualitative directions. The simplest variant, $\mathcal{OPRA}_1$, accordingly divides the plane into 4 equal regions, which due the orientation of the point considered, which functions as origin, can also be labelled as front, back, left and right. The relations that make up the calculus are then relations between such oriented points, each of which brings its own oriented reference system into play. $\mathcal{OPRA}_2$ then divides the 2D-plane into 8 regions, $\mathcal{OPRA}_4$ into 16 regions, and so on, following the divisions also made in the $\mathcal{STAR}$-calculus. It is then relatively straightforward to transform descriptions between granularities, at least when one granularity is a multiple of the other.

An example of a possible configuration specified in $\mathcal{OPRA}_2$ is shown in Figure 28. Here we see two oriented points, A and B, and their positions relative to the other in terms of the regions within which they are located.

This relationship shown in the figure is expressed with the notion:

$$A\,{}_2\angle_7^1 B$$

which means that in $\mathcal{OPRA}_2$ A stands in the region 1 with respect to B and B stands in the

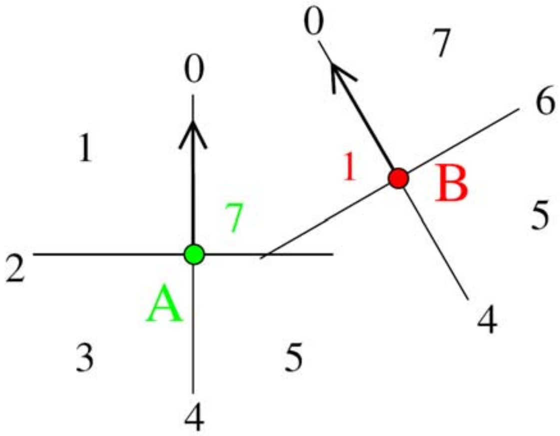

Figure 28: Two oriented points standing in a relation of the $\mathcal{OPRA}_2$ calculus from Moratz *et al.* (2005)

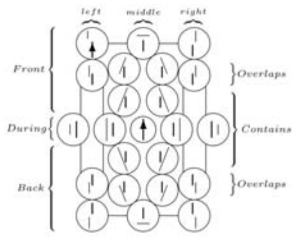

Figure 29: The basic relations of the BA calculus of Gottfried (2004)

region 7 with respect to A. A discussion and some results concerning the composition of such relations is given in Moratz et al. (2005).

Constructions of a variety of calculi, including double-cross and dipole, have been now been provided in various granularities of $\mathcal{OPRA}$ (Dylla & Wallgrün n.d.) and reasoning is provided by constraint satisfaction. Results concerning the formal complexity of the standard problems of satisfiability, etc. for the calculus are not yet very detailed.

8.5 Bipartite arrangements

Gottfried (2004) presents a related characterization of a qualitative calculus operating over relations between intervals on the plane that has been applied in rather different areas to the spatial characterizations seen so far. The original application was for supervising traffic scenarios and, since then, there have also been characterizations of **qualitative shapes** (Gottfried 2005). As with the other spatial calculi seen, Gottfried's proposal is for a relation algebra over the domain of interest. The relations of this algebra, called **Bipartite Arrangements** ($\mathcal{BA}$) is shown in Figure 29, taken from Gottfried (2005).

The paths formed by connected sequences of lines related by these relations has been proposed as a succinct and effective way of describing the shape of plane polygons. The mechanism here is that a relation can be derived for each vertice of the polygon and the collected set of relations then characterizes the overall shape at hand.

8.6 Ordering as an additional level of description

Finally, we briefly present a rather different kind of approach developed by Schlieder (1995). Here orientation is added as a generalization into two-dimensions and above of the temporal ordering relations developed for time by Allen (1984). Schlieder's approach gives a way of formally describing orientation that is more powerful than simple mereotopology but less powerful than complete metrical descriptions. This clarifies descriptions further and allows us to specify exactly what is being committed to by particular descriptive schemes. Schlieder describes how approaches such as Frank & Kuhn's (1986) **cell complex representation** for geographic databases and information systems in fact combines topological and ordering information of the type he defines. Further discussion of possible formalizations of this level of description are given in, for example: Schlieder (1996*b*) and Schlieder (1996*a*).

8.7 Conclusion

While it cannot be doubted that directionality and orientation must appear somewhere in our account, the question of precisely where they should appear is also somewhat open. There are certain indications that orientation cannot be made to fit into a spatial ontology in the same way non-orientational accounts because the locations identified depend on further factors in addition to simply the object being located. This suggests the question as to whether orientation is 'real' in the sense of belonging to an observer-independent reality, whether it is something induced by a cognitive agent's interaction with its environment, or whether it is a way of talking about the enviroment in certain contexts.

Significant for all of the possible accounts is the increasing role of an 'observer' who is situated within the world that is being modelled; we no longer have an independent, observer-free characterization of a spatial world independent of those who are situated within it and the objects that must be located. For the inclusion of directionality and orientation, both kinds of information entail situated observers who can take up points of view within the modelled spatial world. The extent to which this may influence the spatial ontology is still to be investigated; it is certain, however, to play a crucial role in our consideration of spatial language—there the role of the observer, particularly in terms of adopted **frames of reference**, is well-known as an inescapable component of any adequate model. Within the spatial ontology realm, however, it is far from clear whether orientation is to be included simply alongside other spatial properties.

Formally there is certainly no problem with admitting directed line segments in addition to lines; this can be generalized to solid objects in various ways. But this is less than clear cognitively. When we consider the cases where orientation plays a role, it appears natural also to consider notions of *movement*. That is, a path is only directed because an entity has to move along it: that movement occurs in time and time has an inherent directionality. This is sufficient to induce an ordering of points or regions met along the path. This direction of formalization has been pursued by, for example, Eschenbach & Kulik (1997) and Habel & Eschenbach (1997). Thus, both space and time might commit to an ordered geometry of some kind—which makes the relation of 'between' a basic definable relation, but it is only time that takes that further by adding direction. Habel & Eschenbach (1997) use this to motivate some asymetries between prepositional usage in the spatial and temporal domains, relating these to "whether the domain in question has an inherently distinguished direction".

A further indirect suggestion that it may be appropriate to attend more to the movement of an entity when incorporating directed within the ontology comes from Smith & Brogaard's (2003) and Mulligan's (1999) agreement that "the way your perceptions relate to external reality depends upon your trajectories of possible action." This, with its echoes of Gibson and ecological psychology mentioned elsewhere in these deliverables, certainly argues in favor of placing action more at the center of at least the directed spatial entities and relations.

This may also present a finer consideration of spatial dimensions. For there are, of course, spatial dimensions which do in some sense have an inherent direction and which relate interestingly to trajectories of possible action: in particular, the up-down dimension is very different to those perpendicular to the vertical. Less so rooted in physical reality, but still significant, certainly cognitively, is the dimension of **front** *vs.* **behind**; this does not itself have a directionality though. The relationship between such bodily dimensions and spatial description, particularly of **paths**, has been taken up in detail in Eschenbach, Tschander, Habel & Kulik (2000).

The precise ramifications of these considerations will need to be returned to in greater detail; for the present we suggest the utility of considering directionality as less of a normal spatial relation as SUMO or OpenCyc would have it and more of a derived property of potential motion and interaction. As a line of pure exploration, perhaps directionality forms some kind of virtual (i.e., fiat) niche within which particular kinds of spatial relations might be defined: including the gravitationally-induced dimension of up-down, the cognitive-perceptive dimensions of front-back and perhaps the purely fiat dimension of left-right. Only the first of these is in any way a candidate for inherent direction; the second becomes directed during action.

9 Bennett's spatio-temporally founded ontology

A case where the link between purely spatial calculi and ontology has already been made considerably stronger is the explicitly ontology-oriented work growing out of RCC pursued by Bennett and colleagues (e.g., Bennett 2001*a*, Bennett & Galton 2001). This work has, very broadly, taken the development of a generalization and extension of RCC to include a much more powerful treatment of space called the **Region-Based Geometry** (RBG) of Bennett, Cohn, Torrini & Hazariki (2000): this uses basic mereological parthood and a 'morphological' predicate which, for purposes of axiomatization, can either be the **Sphere** predicate or the **Congruence** predicate. Bennett *et al.* show how their axiomatization allows the distinct types of connection relation described by Borgo et al. (1996*a*) to be defined, and includes the generic connection-relation of RCC (Section 7). RBG is a very expressive theory, which naturally brings both positive and negative properties: for example, it is has very poor computational properties and there is little expectation that accounts written within the full formal theory will be directly usable for reasoning, whereas, on the other hand, the account can represent arbitrary geometric configurations and provides a fully qualitative version of geometry.

RBG is then used as a foundation for Bennett's ontological work so that distributions of 'matter' in space can be described. Particular chunks of matter are then identifiable as individuals of various sorts, which brings us to the general concerns of ontology that we have seen above and in Deliverable D1. Thus, as Bennett describes his research programme, the aim is one of

> "establishing an ontology bottom up (starting from just the spatio-temporal distribution of matter types)." (Bennett 2002)

Bennett acknowledges that higher order objects such as those with "biological or artifactual significance" are going to present a variety of difficulties, but starts with "generic physical objects", which he terms "mesoscopic rigid physical objects", in order to get his logical treatment underway. Of particular concern to him is an account of physical objects and their identity over time: i.e., how they can be modelled when breaking, being worn or damaged, colliding and moving, and so on. The connection of identity to the history of an individual, particularly with respect to the history of the chunks of matter making up that individual, is of particular note: there is little doubt that accounts of identity will need to include such considerations. Elsewhere, however, there are modelling decisions that we will need to view more critically: in particular, Bennett appears to be combining the ontology of space and time, that of physical objects, formal ontology and linguistic semantics:

> "We want to encompass the essential insights of all these different approaches within a single semantic framework. Our paper may be regarded as an attempt to produce a detailed *ontology* (as advocated by e.g.,(Guarino 1998)) of time and events accounting for their various linguistic manifestations." (Bennett & Galton 2001)

This means that the ontological distinctions proposed should also have linguistic consequences. They then contrast their approach to "much recent work on ontology for AI systems" that has been carried out at "the level of axioms", stating that their approach is at "the semantic level". We thus have an explicit combination of ontology/semantics, at the one extreme subject immediately to logico-philosophical constraints and at the other end superficial linguistic details in grammatical categories such as nouns (mass and count), verbs, tenses and the like. We will critique this combination of very different ontological domains and its consequences in our deliverable D3.

The main components of the ontology development undertaken center around the definition of two logics: $\mathcal{D}$, whose denotational semantics is intended to be an explicit classical (i.e., not relativistic) model of physical reality, and $\mathcal{VEL}$, the **Versatile Event Logic**, for covering the description of events. The motivation of the approach pursued with the logic $\mathcal{D}$ is that it may be contrasted with other axiomatic approaches to formal ontology which

> "specify a conceptual vocabulary by large numbers of axioms stating logical dependencies holding among concepts, so the concepts do not have any explicit denotational semantics apart from the standard, very general semantics of arbitrary first-order theories." (Bennett 2001*a*, p105)

$\mathcal{D}$ is then to allow the definition of a comprehensive vocabulary using only a few primitives which have been completely axiomatized according to the given semantics. The precise extent to which this substantively differs from formal axiomatically specified ontologies such as DOLCE is yet to be characterized.

The vocabulary of $\mathcal{D}$ relies on five basic semantic types:

- spatial regions,
- times,
- individuals,
- mass nouns,
- count nouns.

We will in our discussion here ignore the superficial linguistic interpretation of the latter two types and treat them as technical terms within the definition of the logic. Time is organized in terms of **histories**, allowing for branching time and possible worlds. **Indexes** are defined for describing states of the world that combine a history and a time. **Individuals** provide for temporally persistent objects, i.e., endurants, as a category ontologically distinct from the spatially extended regions that the matter of those objects occupies. Formally, individuals are modelled as functions from given histories and given times to spatial regions. These spatial regions represent the extensions of the individuals concerned. At each index, i.e., given history and given time, a **count noun** picks out a set of individuals; this is based on Gupta's (1980) logical exploration of a particular class of linguistic expressions and provides a "double indexicality" in order to cope with the linguistic phenomenon that referring expressions can be used to pick out individuals for discussion at one time in terms of properties that they may hold at another time. Bennett cites the example:

Some girl will become president of the USA.

in which an intended reading is that some individual who is a girl now will later, i.e., when she is no longer a girl, become president. For Bennett, then, the 'count noun' *girl* picks out an individual now, i.e., at the current history-time index, and that individual is itself a function mapping history-time indexes to spatial extensions. This means that we can then find the girl when she is no longer a girl and has, in fact, become president.

Bennett defines within his logic notions of continuity which enable changes in matter distribution to be restricted to continuous changes (apart from when there are creations, destructions, breakages and other catastrophic events). The extensions of individuals are then defined relative to these evolving (i.e., changing with history and time) distributions of matter. In particular, identity criteria for physical objects are given in terms of **chunks** of matter. A chunk is defined as: 'a maximally self-connected region of some matter type' (cf. Bennett 2001*a*, p113). These chunks do not, however, "directly correspond to the objects referred to in everyday communication" (Bennett 2002) and so, crucially:

> "When we talk of a cup or a brick we are referring to an object that continues to exist even when it is chipped or scratched. Hence we need to characterise types of object which persist through the loss of certain parts. To construct a rigorous ontology of commonsense physical objects one needs a theory which, although it must somehow relate them to the matter from which they are formed, does not simply reduce them to idealised chunks." (Bennett 2002)

This is a complex task; especially when the foundation adopted is that of spatially-distributed collections of matter as in Bennett's case. Well-known problems that Bennett discusses include those of:

- the mismatch between logically accurate descriptions and the way people talk about things: e.g., a block may be described as "cubical" even though "when we look at it closely we find that its surface is ridged and uneven" (Bennett 2002). This Bennett terms an *idealizing approximation*.
- the problematic notion of identity of objects that may gain or loose parts over time.

We will make some comments on each of these in turn.

First, stating that descriptions are approximations places an ontological emphasis on a 'real' that is then approximated to. Although this may be necessary given Bennett's position, it is also reductionist. There is a physical reality, and people's language does not measure up and instead has to approximate; this is accepted by Bennett and formalized in terms of a **supervaluationistic** logic (Fine 1975) that allows 'vagueness' of the kinds we saw in Section 2.5 above. This has the positive feature of allowing certain precise formalizations to be entertained independently of underlying vagueness in the precise extensions of terms. The acceptance of approximation is, however, also a necessary step given the methodological position of providing *positive* definitions of terms: what can 'cubical' mean if not that an object has certain ideal three-dimensional geometric properties, properties that in the real world, unfortunately, will never actually be satisfied. There are alternatives perhaps: for example, with the 'negative' (paradigmatic) definitional style of de Saussure's linguistics (de Saussure 1959/1915), one would look instead at 'cubical' as one term in a set of oppositions: what 'cubical' means is that we divide out possible shapes into several categories, including 'spherical', 'flattened', etc. and the object we are describing is to be asserted to belong to the 'cubical' term rather than the others. In this sense, the use of 'cubical' is exact and not an approximation: there are interesting connections to be drawn here with the function of granular partitions that we introduced above (Section 6.5). According to this view, it is only when we enter the very distinct discourse of ideal geometry that there is a problem of inexactitude. Ideal geometry may not be the only discourse where a described object falls short of being sufficiently ‘cubical’: but in all cases the degree to which the term fits or not depends on the discourse involved, which is another way of saying that it depends on the function and purpose of attempting the description in the first place. This may be compatible with Bennett’s position in as far as the inexactitude is absorbed into the various ‘precisifications’ that a supervaluationistic approach demands for its vague terms, but this would need further investigation.

There is also a relationship here with the second problem area above: i.e., dealing with objects gaining and losing parts. Here, since identity is defined in terms of the amount of matter and that matter’s spatial extension, Bennett needs to introduce a range of further theoretical apparatus. The basic fact of this kind of identity is well known and much debated as we saw above: an object remains the same object when it only looses or gains *insignificant* parts. The question is how ‘significance’ is to be admitted (or excluded) from the account. Bennett (2002) sets out three ways in which a part may be significant:

1. “A part may be significant purely in virtue of its physical extension. That is it may be too small to consider important or even too small to observe.”
2. “A part of an object may be insignificant relative to the geometry of the extension of that object—differnce between the geometry and topology of the object with and without the part is in some sense ‘negligable’.”
3. “A part of an object may be insignificant in relation to the purpose or functionality of that object.”

In both Bennett & Galton (2001) and Bennett (2002), the problem is essentially attacked from the perspective of the spatial-temporal masses adopted as foundation: an insignificant part is one which is too small to be significant. Just what that size may be depends on the particular objects involved. This ‘maximally insignificant sphere’ can then be built as a parameter into the account, but actually there are few answers concerning just how big or small this might be. And, as we argued above, there can in general be no answer to this question without considering the third of the possible kinds of insignificance listed here: that of functionality and purpose.

There are in Bennett’s account the beginnings of moves to include function. One of these is to allow the designation of “essential parts” or **identity sustaining pieces** (ISP) that have to remain for identity to continue. Such pieces can be assumed to have definitions in terms of functionality without requiring that those definitions themselves be provided—which is a useful way of avoiding discussing all the possible ways in which some part may be deemed significant or not. Then:

> “In respect of its identification criteria, an artifact of a given kind generally has some function which it must fulfill or be capable of fulfilling in order to count as falling under that kind. For instance a CUP might be defined as ‘a rigid body that

has a capacity to hold liquid and (in its usual sense) can be held easily in one hand'." (Bennett 2002)

Bennett is then hopeful that such criteria can find a formalization in terms of identity sustaining pieces and that his logic $\mathcal{D}$ thus provides a useful account. But he also admits that *identity criteria* (as opposed to identification criteria) for artifacts and biological objects might not be so readily reduced to the rigid distributions of matter of his physical objects. This leaves for us a rather large theoretical question mark over the applicability and utility of a substantial component of Bennett's formalization.

The main properties of the other component of Bennett's work, the Versatile Event Logic, moves in a similar direction to the logic $\mathcal{D}$ but focus more on time and events. It allows statements that freely combine explicit reference to time points with an ordering relation, temporal intervals and their interrelationships, and tense-like operators holding over propositions. It adopts a similar ontology to that of $\mathcal{D}$ with respect to time, histories and objects, but adds in a discussion of events (described by 'verbs') and the participations of objects in events. Within $\mathcal{VEL}$, particular segments of the history structure are called **episodes**; event types are then defined as sets of episodes and event tokens are regarded as "an episode seen from a certain perspective." (Bennett & Galton 2001) Whereas space and matter appear to play the major role in $\mathcal{D}$, for the versatile event logic a similar role appears to be granted to time in the form of a structured history. The approach focuses on formalization issues and it is unclear to us as yet to what extent this will be compatible or not with other ontological formulations of the areas of concern. Nor are we entirely convinced that the ontological considerations, particularly in the area of significance, are those most suitable for the domains addressed. This also requires further investigation.

10 Qualitative Movement

As we have mentioned only in passing above, views of space without time, particularly of *movement* in space, are somewhat restricted in their utility. In a series of quite recent papers, there have accordingly been proposals for the direct formalization of accounts of motion in the spirit of qualitative spatial calculi for other aspects of space. We focus for the purposes of this overview on the **Qualitative Trajectory Calculus** (QTC) and its variants under development by Van de Weghe, Cohn, Tré & Maeyer (2005) and Van de Weghe, Cohn, Maeyer & Witlox (2005).

The essential idea of the QTC is that it is possible and useful to capture the relative motion of entities in a qualitative fashion. That is, a calculus is developed which abstracts away from the details of precise movement in order to capture qualitatively distinct categories of relative motion. As a starting point Van de Weghe *et al.* develop a calculus involving relative movement in one dimension; this they term QTC 'basic' in one dimension (QTC_{B1D}). The calculus expresses relations between two point-like entities that are disconnected. In one dimension, it is possible to enumerate an exhaustive and mutually disjoint set of relations holding between these entities on the basis of whether they are moving towards each other, moving away from each other, moving in the same direction, moving towards or away from the other which is stationary, or are stationary with respect to each other.

A succinct representation of the possibilities is offered by Van de Weghe *et al.* as follows. The relation between two entities k and l, possibly in motion, in one dimension may be captured qualitatively by a triple in which the three elements are determined as follows:

1. if k is moving away from l: +
   if k is moving towards l: -
   otherwise: 0
2. if l is moving away from k: +
   if l is moving towards k: -
   otherwise: 0
3. if k is moving faster than l: +
   if k is moving slower than l: -
   otherwise: 0

Sequence of such characterizations of the relations between entities characterize possible changes in relative motion.

Van de Weghe, Cohn, Tré & Maeyer (2005) provide a useful example of this in which they characterize the changes in relative motion in a situation where a lion sees a zebra that is at rest, starts stalking the zebra, the zebra sets off, outruns the lion for a while, the lion gains, then the lion gets tired, and they both, somewhat later, come to a stop. This is expressed in a two dimensional variant of the QTC (QTC_{B2D}) which is constructed quite simply by focusing on the changing *distance* between the entities related. Thus it is suggested that, at least in the case of the lion and the zebra, it does not really matter precisely how the two entities move in space, the qualitalitively relevant aspect of their motion is whether or not they cease to be disconnected! This reduces the two dimensions to a single dimension which can be described as above.

The resulting **conceptual animation** suggested by Van de Weghe *et al.* is then as follows:

$$\{(000) \leadsto (\text{-}0\text{+}) \leadsto (\text{-++}) \leadsto (\text{-+}0) \leadsto (\text{-+-}) \leadsto (\text{-+}0) \leadsto \\ (\text{-++}) \leadsto (\text{-+}0) \leadsto (\text{-+-}) \leadsto (0\text{+-}) \leadsto (000)\ \}_{B2D}$$

Van de Weghe *et al.* provide a loose natural language gloss of the events being described here but this does not pick out the precise contribution being made by the qualitative details of the specification. In Table 6, we set this out in a way that relates the QTC relations more directly to the motion issues involved. Particularly interesting is their suggestion that this formalism may provide a means of capturing the semantics of natural language terms of motion and, indeed, their characterization does appear to be picking out certain such elements, for example, to 'pursue', to 'escape', etc. are intrinsically concerned with the relative relations of motion and direction of the two entities invoked.

A rather more complex version of the QTC is also suggested to operate in two dimensions and bringing in some of the orientation and directionality properties of the Double-cross

	relation	*gloss*
	(000)	lion (k) is not moving relative to zebra (l), zebra (l) is not moving relative to lion (k), there is no difference in speed.
$\leadsto$	(-0+)	the lion moves towards the zebra (-), the zebra is still stationary (0), the speed of the lion is greater than that of the zebra (+).
$\leadsto$	(-++)	the lion is still moving towards the zebra (-), the zebra is now moving away from the lion (+), the lion is still faster than the zebra.
$\leadsto$	(-+0)	the lion is still moving towards the zebra (-), the zebra is still moving away from the lion (+), the zebra matches the speed of the lion (0).
$\leadsto$	(-+-)	the zebra moves faster than the lion.
$\leadsto$	(-+0)	the zebra slows down and the speeds are matched.
$\leadsto$	(-++)	the lion gains speed and is faster again.
$\leadsto$	(-+0)	the lion slows and speeds are matched again.
$\leadsto$	(-+-)	the lion slows more and the zebra is now faster than the lion.
$\leadsto$	(0+-)	the lion stops, the zebra is still moving away from the lion, the zebra is faster than the lion.
$\leadsto$	(000)	the zebra stops too.

Table 6: Conceptual animation of a chase with the qualitative trajectory calculus

calculus (see Section 8.2 above). This variant, termed QTC_C is described by Van de Weghe, Cohn, Maeyer & Witlox (2005) and involves considerably more relations. These now describe not only the one-dimensional relative motion but also the orientation along the lines made possible by the Double-Cross configurations. Thus one of the entities may move in one of several qualitative *directions* away or towards the other entity. This is shown well graphically in Figure 30, taken from Van de Weghe, Cohn, Maeyer & Witlox (2005). Here each entry in the table is to be interpreted as indicating the direction of movement available: the identified entity can either move in parallel to the other, or can move to anywhere on the arc indicated.

Van de Weghe *et al.* also suggest a naming scheme in terms of a 4-tuple of values formed similarly to the triple for QTC_{B1D} and QTC_{B2D}; this is shown in the figure.

Again as a practical example of the potential utility of this kind of calculus, Van de Weghe, Cohn, Maeyer & Witlox (2005) present an illustration of the formalization of an 'overtaking event' as might occur in a traffic scenario. As with the previous example, it is possible to construct a conceptual animation showing the qualitatively necessary spatial and motion relationships of two cars involved in such a manoeuver. Principles of continuity, which state that a value cannot change from one value to another without going through the available intermediate values, then allow a segmentation of the events involved. The example also represents the precise difference between overtaking in the UK and overtaking in Continental Europe, as can be seen in the following contrasting conceptual animations:

CE $\{$(-+00) $\leadsto$ (-+-+) $\leadsto$ (00-+) $\leadsto$ (+−+) $\leadsto$ (+-00) $\}_{QTC_C}$

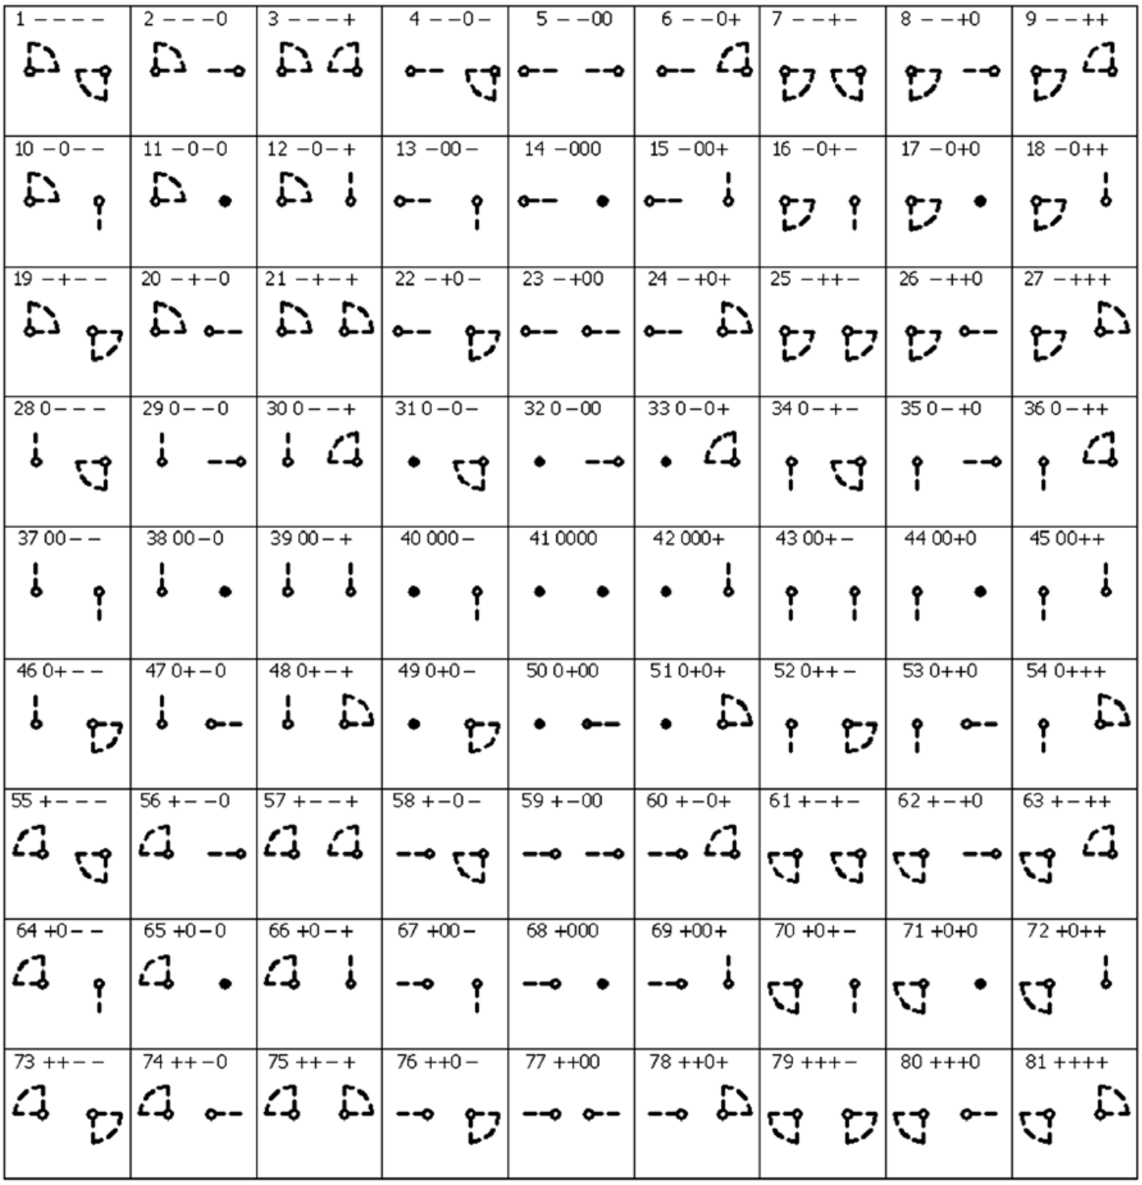

Figure 30: The basic relations of the qualitative trajectory calculus (QTC) combined with orientation (Van de Weghe *et al.*)

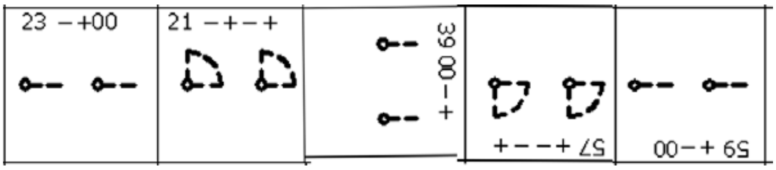

(a) Continental Europe Overtaking Event

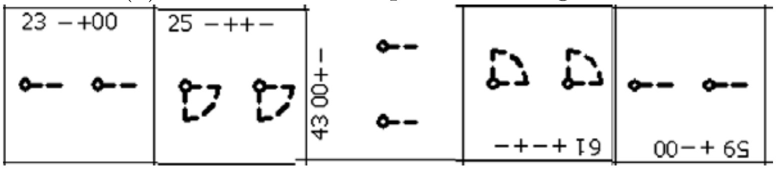

(b) UK Overtaking Event

Figure 31: The contrasting overtaking manoevers of the UK and Continental Europe represented in the qualitative trajectory calculus QTC_C (Van de Weghe *et al.*)

UK $\{$(-+00) $\rightsquigarrow$ (-+-+) $\rightsquigarrow$ (00+-) $\rightsquigarrow$ (+-+-) $\rightsquigarrow$ (+-00) $\}_{QTC_C}$

This representation brings out nicely the similarity in the manoevers involved. To make this clearer, Figure 31 shows the same conceptual animations but folding in the graphical representations for the individual relations, rotating individual graphical elements as necessary to maintain an iconic sense of orientation (which does not change the nature of the qualitative relation holding of course). As Van de Weghe *et al.* propose, this kind of semantic representation may well provide formal mechanisms for distinguishing particular categories of movement and their related subevents: such as, for example, recognizing the particular changing lanes or passing events that together make up an instance of overtaking.

Composition for the relations of the QTC are discussed in Van de Weghe, Kuijpers, Bogaert & Maeyer (2005).

11 Geographical Information Systems and Geographical ontology

As with several of the other areas discussed in this deliverable, it cannot be our purpose here to provide an introduction or overview of Geographic Information Systems (GIS) in general (for such introductions and further pointers, see, for example: Fonseca, Egenhofer, Agouris & Câmara 2002, Longley, Goodchild, Maguire & Rhind 2001). Our concern is restricted to the growing interaction between GIS and ontology design and application since there is now a recognized need for ontological specifications in many areas relevant to GIS. Whereas the basic standards that have been developed in this field, such as the GML (Geographic Modelling Language), are still essentially geometric and define points, lines and areas and some relationships between these, the richer account of the meaning of geographic objects and relationships that is possible within an ontology offers a stronger basis for intelligent systems—particularly when moving into the areas of administrative and institutional geography and change over time. As a consequence, there are now extensive research efforts within

GIS concerning far richer geographically-relevant ontologies. These concern themselves with issues of spatial representation and are also increasingly overlaping with traditional areas of computational ontology which seek general modelling schemes for ‘commonsense’ objects, particularly those relevant for geography in its widest senses.

The development of ontologically-motivated accounts of geographic information also brings with it several useful sets of consideration of their own. One of the primary motivations for pursuing such ontological modelling, for example, is the promise of integration and interoperability across systems that draw on very different kinds of data (e.g., differing kinds of remote sensors operating at different geographical scales) and involve different classification schemes according to the kind of data considered. As Fonseca, Egenhofer, Agouris & Câmara (2002) write: what is sought is

> “a GIS architecture that can enable geographic information integration in a seamless and flexible way based on its semantic value and regardless of its representation.”

Ontologies are then seen as a way of achieving this aim (cf., e.g., Visser, Stuckenschmidt & Vögele 2001, Fonseca, Egenhofer, Agouris & Câmara 2002, Frank 2003*b*).

GIS also provides a good testing ground for many aspects important to ontology and ontological engineering. As Frank writes:

> “The design of Geographical Information Systems, which cover information about objects and properties in the world with respect to their location ..., involve ontologies. Indeed, such systems are ontologically more demanding than administrative information systems. They span a much larger diversity of kinds of things: from the description of the elevation of the surface of the earth to the description of the natural land cover (woods, fields, etc.) and morphology (mountains, valley, etc.). They also include man-made features like roads and buildings as well as artificial boundaries between a range of different sorts of political and administrative units.” (Frank 2001, p667-8)

We take the fact that there are actual, very real and concrete tasks for managing and using geographic information that need to be resolved—with or without ontologies—as a very positive force for driving ontological research and development. This pushes ontology proposals within GIS beyond simpler views of ontology—such as, for example, those prevalent in semantic web discussions— and leads naturally to the more powerful notions of ontology that we are mostly concerned with in these baseline deliverables.

We focus on three of the most prominent lines of interaction between GIS information modelling and ontology: the ontologies and ontological mechanisms proposed by Egenhofer and colleagues, by Frank and by Kuhn. Some of these proposals overlap to a certain degree—both in their developmental histories and issues raised—but provide useful complementary perspectives on the functionalities necessary within GIS. We do not address issues internal to GIS and its development (cf. Longley et al. 2001); nor do we consider existing standards such as ISO 19107 (Spatial schema), ISO 19112 (Spatial referencing by geographic identifiers), OpenGIS or particular formats for geographic information such as GDF (Geographic Data Files) or GML (Geographic Modelling Language). While essential for a proper consideration of geographic infomation in the large, these standards are, in their current form, less relevant for the particularly ontological concerns of this deliverable.

Our discussion also focuses solely on questions of ontology design and is essentially programmatic in that we are developing a basis for further ontological developments—it is to be hoped that these developments will then feedback into application, including ontologies for GIS. The aim is that this will both improve modelling decisions and provide more of the more advanced functionalities required of ontologies in their application within GIS.

11.1 Ontology-driven geographic information systems

The first line of ontology development for GIS that we consider is that described by Fonseca, Egenhofer, Agouris & Câmara (2002). Fonseca *et al.* start from a characterization of levels of abstraction from computer science and add to this a **cognitive level** that is to capture "what people perceive about the physical universe". This is made up of objects, relations and processes and fits most naturally within the views of ontology that we have seen above. Such a representation is to be provided with a formal specification and a computational implementation in order to support practical use within GIS.

One particular aim is to derive the computational implementations from the ontology specifications themselves, mapping, for example, ontological organizations to classes and methods within a Java implementation. This relates directly to one of Fonseca *et al.*'s major concerns, that of using ontologies as an interface between geographic information systems and their users. Under this view ontologies are structures that are to be **traversed**: the user is presented with a graphical rendition of the ontological organization and may select various paths to navigate around both the ontology and the data that that ontology organizes. Various methods associated with the classes in the ontology provide support for navigation as well as for merging or transforming information in various ways. Many of the mechanisms proposed for ontology implementation and the form of the discussion of these mechanisms then draw on Object-Oriented programming techniques of inheritance over structured objects with associated 'methods' for operating on those objects. In short: "Ontologies . . . are seen as dynamic, object-oriented structures that can be navigated." (Fonseca, Egenhofer, Agouris & Câmara 2002)

There are two particularly important principles involved in the ontology-driven GIS (ODGIS) approach that we will consider more closely here.

First, as with most GIS-ontologies, it is regarded as essential that one can maintain information at a variety of levels of detail. Information at a fine-scale from sensors of various kinds needs to be related through various levels of detail to make-up larger-scale views of an entire environment—in the limiting case, the entire planet. Shifting between scales is naturally a traditional concern of geography and geographic mapping. Relating levels of detail to ontology navigation, Fonseca, Egenhofer, Agouris & Câmara (2002) refer to **vertical navigation**, where a user browsing an ontology can choose to consider either more specific or more general nodes in the ontological hierarchy. As we mentioned above, in Section 6.5, this concern of ODGIS treats granularity as semantic granularity, i.e., related not to size but to conceptual aspects selected to be included in the representation.

Second, Fonseca, Egenhofer, Agouris & Câmara (2002) draw attention more than most other authors to the crucial relationship between an ontology and a community of users for that ontology. Taking this relationship seriously is an important prerequisite for achieving real interoperability—a prime consideration for all GIS efforts and for ontology in general. Fonseca *et al.* argue that attempts to build interoperability soley on the definition of shared **standards** are misguided:

> "Since widespread heterogeneity arises naturally from a free market of ideas and products there is no way for standards to banish heterogeneity by decree. The use of semantic translators in dynamic approaches is a more powerful solution for interoperability than the current approaches, which promote standards." (Fonseca, Egenhofer, Agouris & Câmara 2002, [further references ommitted])

Different communities will develop their own agreed ontologies and it is then the task of ontology and ontology engineering to consider ways and means of relating these. Some of the methods for such integration rely on considerable formal apparatus and we will return to these in our deliverable D4; for present purposes, we consider the implications of heterogeneity as a component of foundational ontologies.

This has several consequences for the design decisions taken within the ODGIS-approach. On the one hand it relates to granular partitions as introduced above:

> "We use also hierarchies of groups [geospatial information communities: GIC] to generate ontologies of different levels of detail. For instance, in a city, the mayor and his/her immediate staff view the city at a higher level. The department of transportation, the section in charge of the subway system will have an even more detailed view of the city." (Fonseca, Egenhofer, Agouris & Câmara 2002)

Thus, we can see ontologies at various levels of detail in terms of selected granularities (in the sense of semantic granularity). Moreover, on the other hand mechanisms need to be provided that support the posited equal rights of different communities to define ontological categories.

This task is addressed by Fonseca *et al.* in a novel way. Considering standard examples such as a particular entity being both a 'building' and a 'school' or a 'factory' or a 'shopping development', they note that one approach to modelling this state of affairs that is *not* ontologically sufficient is that of **multiple inheritance**. Using multiple inheritance one could write that a particular entity is both a building and a factory—thereby inheriting properties from both. While there are problems raised here within object-oriented programming concerning how exactly the properties and methods inherited from the combined class' parents are to be merged, there is also the much more fundamental difficulty raised by attempting to violate the basic ontological criterion of *identity* discussed at length in Deliverable D1.

Identity is one of the criteria specified in Guarino & Welty's (2004) OntoClean methodology for ontology construction. A physical object such as a building and an institution such as a school cannot be placed in a single inheritance hierarchy without violating basic ontological meta-properties: if we use a building as a school, we cannot simply state that the physical entity is then an institution because then a physical entity would be a social construct. Neglecting this kind of issue has been one of the main causes of unnecessarily tangled hierarchies in ontology and knowledge representation.

Fonseca *et al.* accept this fundamental problem and avoid it by extending the notion of **roles**. Roles are potentially useful for ontology design, and as seen in Deliverable D1 are accordingly increasingly often appealed to, precisely because they avoid violations of otherwise rigid properties. A person might be a student at one time and a parent at another and a chairperson of a company at another. Each of these have different identity criteria and certainly do not represent mutually rigid properties. Just what the appropriate treatment of entities and the roles that they can enter into is, however, still a matter of active debate and research (for a useful literature review, see: Loebe 2003). For Fonseca *et al.*, roles offer a way of accepting differing views of objects without raising ontological problems; a given entity can adopt various roles throughout its lifetime. Roles are incorporated as an additional 'slot' in the information maintained about some class. Fonseca *et al.* also relate this to varying communities of users. Thus, while a lake may be a geographical region for some community, for the water department it might be a source of pure water, for the environmental scientist a wildlife habitat, for a tourism department a recreation point, and for the transportation department an obstacle. This is captured by saying that the ontological entity lake may play the role of a habitat or an obstacle, etc. Roles are thus used as a way of incorporating perspectives or contexts.

This is then taken one step further in order to accept the equal status of alternative perspectives or contexts. Roles are not restricted to a special class of dependent ontological entities as is elsewhere most usually the case—*any* concept can be used as a role. This, Fonseca *et al.* argue, can then be used as a general mechanism for relating different ontologies. For example, in the ontology of the community of the tourism department, there may be categories that involve recreation points of various kinds and these have their own necessary and rigid properties and identity criteria. One can then relate the entirely distinct ontology of physical objects and geographical regions that contains a category lake to that of the tourism department by saying that the lake plays the role of a recreation point. Equally, however, a recreation point may be considered to play the role of a lake:

> "The application developer can combine classes from diverse ontologies and create new classes that represent user needs. This way, a class that represents lake in a Parks and Recreation department ontology can be built from geographic region ... At the same time, lake can be seen as a port for loading cargo, or it can be seen as a link in a transportation network." (Fonseca, Egenhofer, Agouris & Câmara 2002)

This usage contrasts with those ontologies in which there are some given hierarchy of roles maintained in addition to the entities that may play those roles. It also relates interestingly to an ontology such as that of SmartKom (see Deliverable D1), in which almost all entities useful for domain modelling are placed under the concept role in order to move modelling decisions in precisely the same direction as attempted here by Fonseca *et al.*.

This does, however, involve a considerable extension of the notion of role:

> "The main objective of using roles in this work is to employ them as a tool to connect different ontologies. Therefore we use a more unrestrained definition of roles than other authors ... who argue that roles should have their own hierarchy and can only subsume or be subsumed by another role. ... Each community has a right to its own point of view and information must be integrated on that basis, hence the use of a flexible specification of role." (Fonseca, Egenhofer, Davis & Câmara 2002)

Roles in this sense are incorporated into the **navigation metaphor** for interacting with ontologies as examples of **horizontal traversal**, where different perspectives of entities are brought into the foreground. A new entity, constructed by taking the definition of the role-concept as starting point, can be formed in relation to an existing entity; this is illustrated graphically in Figure 32 taken from Fonseca, Egenhofer, Davis & Câmara (2002).

In the Java-like class representation favored by Fonseca *et al.* horizontal navigation is achieved by an `extract`-method defined for all objects. If, then, a lake were to be seen as playing the role of a link in a transportation system, then a new 'transportation link' object would be created which 'inherits' (horizontally) just those of the properties and attributes of the original object that are relevant and appropriate:

```
new object link = lake.extract (link)
```

The problems of moving across perspectives are then placed within the definitions of the extract method for the objects concerned.

While clearly attacking a problem that is of central import for ontology usage, it is a little less clear that the notion that Fonseca *et al.* define as 'role' is really a role in any traditional sense. As indicated above, at places in their discussion the equal rights of different communities makes it appropriate not only to say that a lake (as geographic region) plays a role as a habitat (as biologically-relevant entity), but that equally a habitat (as a biologically-relevant entity) has a role as a lake: it depends from which community one starts from. Role-playing is then, at least potentially, symmetric. At other points in the discussion, however, Fonseca *et al.* do talk about one of these perspectives being 'real'—generally the geographical region view of a lake as a body of water—and the others then are 'only' roles. In this usage role resembles more its traditional usage: as they write "It is a basic assumption of this paper that a consensus can be reached about which are the basic properties of a lake." There is also hope that the basic assumptions from differing communities will converge, when they are basic enough. Figure 32 also shows this division in that there is an assymetry with respect to which entity is indicated as bearing a role and which entity is indicated as being a role; but presumably the diagram could have been constructed from the perspective of the other 'starting' community.

To what extent this is really compatible with a commitment to heterogeneity is not an issue that we need to address at this point. Significant is only the particular mechanism of horizontal navigation and the treatment of perspective that it employs. Both when we consider

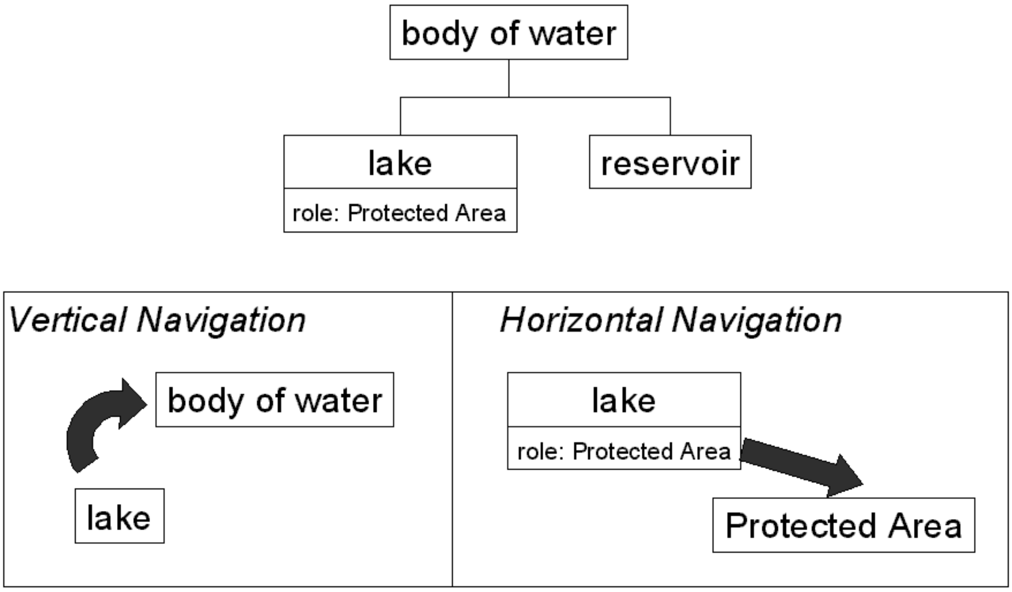

Figure 32: Types of ontology navigation taken from Fonseca *et al.* (2002)

linguistic ontologies and deal with so-called **thematic roles** and when we consider the general participation of entities in activities and states, we will need to return to the 'restrained' notion of role that corresponds more to the traditional assymetric uses of the term. But there is no doubt that we will also need to have some mechanism corresponding to the perspectival shift or inter-partition mapping that Fonseca *et al.* subsume under their use of 'role': its ontological nature will need to be clarified however.

The other components of the ODGIS-framework that Fonseca, Egenhofer, Davis & Câmara (2002) describe are more traditional. Each construct may have parts, functions and attributes; these do not appear to be further characterized however and appear to contain themselves rather heterogeneous entries (e.g., as parts a lake has both water and a beach, and its attributes include both 'acidity' and location—all very different ontologically). The definitions also combine formal and informal glosses. This level of granularity may be sufficient or necessary for the GIS-context, but should be seen as a particular partition on the stronger organizations discussed in earlier sections.

11.2 Frank's proposals for a multi-tiered ontology for GIS

The second direction of ontology development we consider here shares Fonseca, Egenhofer, Agouris & Câmara's (2002) concern with practical use and computational implementation but goes further in the fine characterization of the kinds of information to be maintained. Databases containing spatio-temporal information are to be placed on a sound ontological basis by distinguishing carefully the rather diferent kinds of information that need to be managed. This is in line with a series of articles in which Frank has been proposing an increasingly broader deployment of ontologies to solve problems within the geographic do-

Tier	Domain
Tier 0	human-independent reality
Tier 1	observation of physical world
Tier 2	objects with properties
Tier 3	social reality
Tier 4	subjective knowledge

Table 7: Frank's ontology tiers: taken from Frank (2001:668)

main (Frank 1997, Frank & Kuhn 1999, Frank 2001, Frank 2003*a*). This has now moved away from purely spatial and geographic considerations to include much of what is regularly seen as part of ontology more generally. It is therefore particularly relevant to consider this work within the framework for spatial and general ontologies that we have been developing.

To bring the rather different domains of interest together in a formally and computationally sustainable manner, Frank's suggestion is to construct a **tiered ontology**, in which various ontological perspectives may be organized and related. The basic tiers proposed are shown in Table 7; these tiers are intended to provide a solid foundation for geographic work that can help organize GIS and spatio-temporal databases in general, support information retrieval and geographic information management, and provide more natural interfaces for geographic information systems.

As we can see from the table, there are substantial parallels with other proposals for ontology organization; including general approaches to **stratified ontologies** such as those suggested in Bateman (1995), Borgo et al. (1996*b*) and Poli (1998)—as well as being reminiscent of Donnelly & Smith's (2003) introduction of layers that we saw above (cf. Section 6.4). All of these accounts use the notion of 'stratification', albeit in somewhat different ways and drawing on different historical antecedents (for a broad historical view, see: Poli 2001). The essential notion connecting them remains however: this is that there are substantially different ontological domains which need to be maintained separately. Frank here adds further several useful considerations and motivation for distinguishing tiers which are all rooted in the needs of GIS but can be adopted more generally; these are:

- that differing rules and mechanisms are necessary for consistency checking and maintance within each tier, and
- that "the proposed tiers are ordered from data for which data collections from multiple sources are more likely to agree to data for which disagreement is more likely" (Frank 2003*b*)

Thus, whereas it is assumed that observations of physical reality (tier 0) can be made with a high chance of agreement for any observers (tier 1), institutional interpretations of those observations (tier 3) necessarily depend on a supporting social context and so can vary widely across communities. Moreover individual states of knowledge about those interpretations (tier 4) may be even more subject to disagreement. Note here that Frank deliberate chooses to include accounts that might not be found in traditional ontology: for example, tier 4, with its explicitly epistemological orientation—i.e., concern with states of knowledge—is included as a necessary component of a complete account. Thus Frank is explicitly concerned with constructing a sound basis for a complete information system rather than worrying overly about boundaries of philosophical discussion; here we will focus primarily on Tiers 0–3.

Frank's basic starting point within this ontological framework is a 4-dimensional realist ontology. This, as he makes clear, is crucial for many geographical models, which need to be able to include time, events and changes as well as static representations. Thus, Frank's basic level, 0-tier, is the four-dimensional space-time continuum itself, which can be described "with differential equations". On top of this come further tiers which move progressively towards commonsense human conceptualizations. Particularly interesting is the ordering of Tier 3, social reality, ahead of Tier 4, cognitive agents—a position more similar to the stratification proposed in Bateman (1995) than that more often found in cognitive science, where knowledge (including knowledge of a social world) is made prior. Tier 1, taking in observations of the physical world, then is the most direct mediation between the 4-dimensional space-time continuum and human concerns: such observations are assumed to yield values of various kinds of quantities holding at particular times and locations and within inherently limited degrees of accuracy. As is natural with GIS's concern with the acquisition of actual physical data, this tier is developed more than is the case with other ontologies; this is certainly appropriate for geographic concerns but may well also be an important addition that will play a role in the SFB—for example, via the observation and measurement of the environment by robot sensors.

For present purposes, however, i.e., that of relating the various approaches to space and spatial entities (both objects and states/events) taken in ontology overall, it is Frank's Tiers 2 and 3 that are of most concern. Tier 2 corresponds most closely to much of the above discussions and is intended to capture the world of physical objects and their interrelationships. The objects of this tier are founded on the observations of tier 1 by relaying on "regions of uniform values": that is, when one has a 4D region within which specified groups of observations are 'uniform', we are dealing with a potential entity. Individual point-like observations are assumed primitive (Tier 1); observations of change are constructed out of differences between observations. For everyday human interaction with the world, this clearly occurs below the level of explicit conceptualization as in many domains there is no access to such observational behavior, even if the brain were functioning in this way.[3] The construction of this tier of the ontology is therefore essentially spatially (and temporally) centered.

Frank considers this way of defining objects in terms of uniform regions as a means of avoiding certain philosophical problems—particularly those related to disagreements concerning just what the basic objects of an ontology should be. The properties out of which objects may be defined are essentially Gibsonian: i.e.,

> "The properties must be uniform for an object are related to the possible ways of interaction with an object. Depending on the property, which is uniform, very different types of objects are formed and these objects then follow different

[3]Although there is evidence for perceptions that are not founded in point-observations, for example considering 'visual flows'; language perception also has little to do with points. How these issues are reconciled with the approach described is as yet unclear to us.

> ontological rules. ... The properties, which are fixed to determine uniformity, can be used to define a topological, morphological or functional unity.” (Frank 2003*b*)

This is another way of binding object creation to purpose, just as we saw above (cf. Section 6.5) in Bittner and Smith’s definition and use of granular partitions. Where Frank’s account differs, however, is in the stronger relation assumed between partitioning into objects and observable uniform properties. It is not entirely clear whether this stronger relation is supportable in general: the problem of ‘wholes’ does not appear to be addressed sufficiently; for example, there is no uniform property of ‘being sittable on’ that holds of all the space-time points that one measures within the space-time region of a chair—it is the chair as such that one can sit on. Thus, the property could only be a property of the entire space-time region that the chair occupies: A chair does not have any point-observable properties that serve to define it uniquely and unambiguously apart from the circular ‘region’ that is the chair itself. It may be that this could then be included within Frank’s account by a flexible notion of observation point and its granularity, but this is not immediately obvious as a way of proceeding and so we will not consider this particular (at least for us) unresolved issue further here.

Slightly more formally, the description of **material entities** within Tier 2 is then given in Frank (2001) as follows:

> “For present purposes, we posit a set of tokens e.g., names of material entities, which map to spatio-temporal regions, which we interpret as material objects. . . . Material entities are conceptual and correspond to 3d-t regions, which have approximately the corresponding properties; not all 3d-t regions, which have the right properties are adorned with a name. . . . it is possible that at one and the same location, more than one entity exist.”

As Frank points out, this approach also shows certain similarities with Bennett’s (2001*a*) account, in which matter and its distribution play a fundamental role and regions are taken to describe the distribution of matter at any particular time. Then:

> “Assume a set M of tokens m, which map to spatio-temporal (3d-t) regions, with properties X.”

Which brings us back to the view that material entities are defined by their properties—which, as mentioned above, requires a somewhat extended notion of ‘property’ to be tenable for everyday objects. In summary, material entities

- occupy spatio-temporal regions (they are not these)
- consist of material, the material is not part of them,
- are in various ways involved in events (as agent, as object, as location), they are not part of these,
- relate topologically and mereologically to other entities.

An interesting addition to other notions of objects that we have seen in the ontologies above is a closer consideration within Tier 2 of *change*. Objects are seen as enduring in time but not, necessarily, forever: certain events can bring objects into being or destroy them, or break them into various parts, and so on. Here Frank introduces the notion of **lifestyles**: different types of objects partake in different lifestyles. The possible combinations of lifestyles applicable to some entity can also be used to define the object type of that entity. Thus, we can distinguish object types such as:

- solid objects, which, after being created, may maintain their identity until they become broken and even after being broken may regain their identity by having their parts glued together again;
- fluid objects, which, when merged with other fluid objects fuse so that the contributing objects cannot be recovered; or
- weak aggregates, which easily gain or loose parts without losing their identity.

Identity-criteria over time are clearly an area of ontology that is going to need further discussion and relates to the developments pursued in Bennett's (2002) framework of region-based geometry and time.

The concept events are treated within the ontology of Tier 2 similarly to material entities in terms of properties holding over regions but adds the notion of *change*:

> "The difference between entities and events—endurants and perdurants ... or continuants and occurants—is not found in the space-time region and its properties, but in the more fundamental observation of a property respective to the observation of a change in property."

Frank adopts a straightforward allocation of entities to kinds. Events are classified as distinct types of 'processes'. Processes are assumed to be organized into a type signature as usual. The processes are considered as traditional semantic configurations with typed participants. Here mention is made of grammatical case and clause schemata as well as classes drawn from Wierzbicka's (1996) 'Natural Semantic Metalanguage': a primitive lexically-motivated linguistic ontology that we will describe briefly in Deliverable D3.

Mixing motivations for the categories of an ontology can lead to inconsistencies that are difficult to resolve; we criticize this methodological flaw more extensively in our Deliverable D3 where we deal with linguistic ontologies proper. Despite problems, such mixing is commonly found, particularly in projects where some degree of natural language competence has been required.[4] But there are good reasons for being wary of a simple 'mix' of linguistically motivated categories and categories that are intended to be of 'ontological' relevance. The acceptance of the methodology rests on a working assumption, primarily rooted within loosely 'cognitive' linguistics, that language is more or less directly indicative of conceptual organization. This

[4]Such as in, for example, the ontologies of the computational projects Lilog, Verbmobil and SmartKom: see Deliverable D3 for discussion.

is similar to an established tradition in ontological design where various properties of natural language are treated as if they are directly relevant for the entities and relations that are to appear in an ontology. In Deliverable D3 we will attempt to show that the relation between linguistic forms and possible ontological motivations is further apart than the simple juxtaposition of languistic-conceptual-ontological distinctions would lead one to expect. Simplifying this relationship leads to a view of language as essentially 'naming' conceptual distinctions: this does not support the kind of flexibility commonly seen in language use however.

Frank's Tier 3 then considers the reality that is made up of socially or institutionally constituted entities rather than those grounded in physical reality. Here there is a certain tension in the descriptions that may, ultimately, go back to the restricted view of language taken. We illustrate this with respect to the following extended fragment:

> "Social reality includes all the objects and relations which are created by social interaction. Human beings are social animals and social interaction is extremely important. The reason to separate physical reality [Tier 0], object reality [Tier 2] and socially constructed reality [Tier 3] is the potential for differences in observations: within errors of observations, the results of observations of the same point in time and space should be the same. The construction of objects can be based on the uniformity of various properties, and thus objects may be formed differently—for example, the definition of forest can be based on various criteria and thus leads to different extensions of a "forest" (indeed one should speak of different kinds of forest: legal forest, land-use forest, forest as physical presence of trees, etc.); differences for object formation can be tracked back to different methods in classification if enough care is applied to the domain-specific interests and procedures." (Frank 2003*b*)

What, then, and to re-ask Bennett's question, is a 'forest'? In its existence as a physical object (even ignoring problems of the space-time points at which observations are made: cf. Bennett (2001*b*)), a forest according to Frank depends on recognition of a space-time region of uniform properties. Then, because the legal department, and the land-use department and the commonsense view each may employ different collections of properties, we have the different 'extensions' that are possibly associated with 'forest' as such. Each set of methods defines its own objects. But each of these sets of methods is only applicable within a community of users: they are therefore also social objects and as such belong within Tier 3. There is a problem here which relates, as Frank points out, to Smith's (2001) notion of bona fide vs. fiat objects. The means by which one can go about ascertaining whether a description is real, or 'true', depend crucially on its status as being either bona fide, i.e., supported by physical reality, or fiat, i.e., created by human institutions. In the former case one can go out and take measurements; in the latter one must examine how the social decisions involved were made.

This means that Frank's tiers of ontology allow Tier 3 issues to intervene at Tier 2 too readily. Whereas physical objects should be definable in the terms of physical reality, this physical reality should not already be pre-stuctured in terms of the interests of different social groups. Taking the Gibsonian view of perception, and allowing this kind of input to Tier 2, adds in certain interaction-related properties of physical reality but, arguably, should not open up the door fully to socially pre-structured perception. There is a certain equivocation here:

> "Objects seem to be 'real', but one must always remember that they depend on the classification used for their formation and therefore alternative ways of 'carving' up the world in objects are possible. Some classifications are extremely closely related to fundamental operations of the human body and are therefore likely 'universals' (i.e., the same for all human cultures); others are not." (Frank 2003*b*)

But can physical objects really be carved up in alternate ways? In the world of commonsense objects, we would suggest not. It is not a matter of interest-related choice whether the head of the nail is a part of the nail, or whether the nail is part of the wall, or whether the river runs into this sea of that ocean, or whether this region is a non-tangential proper part of that region or not: different communities can choose to impose differing partitions over physical reality but, as Smith & Brogaard (2003) state, these partitions do not 'add anything' to physical reality. This goes against Frank's position, which maintains:

> "This is a fundamental problem for any object ontology: the division of the world into objects is not unique and depends on the observer and his intentions. A special case is given if a classification is finer than another." (Frank 2003*b*, p55 in online version)

But a finer classification does not change reality: choosing to see something at more or less detail is not an ontological variable in the sense of obtaining new objects. We would suggest, therefore, that a division into tiers as Frank suggests is definitely necessary, but that Tier 2 needs to be rather more restricted than it appears currently to be. An ontological stratum at which real physical objects are located, defined by their physical properties and the continuities and discontinuities of those properties, and with respect to which partitions of various semantic granularities can be adopted, appears to us not only to be a more defensible position but also one which will allow of cleaner formalization.

Finally, we note in passing one further interesting aspect of Frank's approach and that is its intended computational instantiation. In order to support real application-reasoning combined with the sophisticated ontological view that he proposes, Frank argues for an algebraic specification rather than the, in formal ontology more traditional, specification in terms of first-order (at least) logic. This allows direct use of modern functional programming tools such as Haskell. There is here, therefore, an interesting point of commonality with the direction that we will be pursuing in Deliverable D4 since, as we indicated in Deliverable D1, an algebraic specification will also be pursued for our own ontology specifications.

11.3 Cognitively-motivated semantic reference systems

From the ontology-related work of Kuhn we pick out two significant strands: one, a concern with the relation between ontology construction and language, the other, the development of more powerful 'ontological' representations for supporting GIS. The latter is currently being pursued primarily within the Musil project.[5]

[5] URL: http://musil.uni-muenster.de.

Kuhn (2001) proposes that an effective way of constructing domain ontologies is to analyze texts from the domain to be covered. This is an established way of proceeding within artificial intelligence, particularly within natural language processing, where a domain model needs to be constructed for some particular range of behavior. It is therefore particularly interesting to see it attempted within the geographic domain. The particularly linguistic aspects of the resulting ontology will be a topic of deliverable D3.

Of more central importance here, Kuhn (2003) and Kuhn & Raubal (2003) have begun promoting **semantic reference systems** as an extension beyond the mechanisms typically found (or, more often, discussed) for ontologies. Such reference systems may be interpreted as upper-level or foundational ontologies, but with improved processing features particularly relevant for GIS, including, for example, **projection** for relating ontologies across differing granularities; in particular:

> "The spatial reference system analogy suggests something more powerful than today's ontology languages can offer. Producers and users of geographic information need tools for *transformations* among semantic spaces and *projections* to sub-spaces. A transformation may occur within or between information communities and involves a change to the reference system (for example, adding a new axiom to the ontology). A project occurs typically within a community and reduces the complexity of a semantic space (for example, by generalizing two entities to a super-class)." (Kuhn 2003)

These kinds of advanced functionalities are reminiscent (particularly in its suggestions for adopting **category theory**) of discussions within ontology standardization concerning so-called 'meta-ontologies' (such as, most prominently, the Information Flow Framework of Kent (2000)); we return to say more about these inter-ontology capabilities in our Deliverable D4. For present purposes, it is useful to consider both the ontological component of these semantic reference systems and their proposed implementation. We can import from both of these aspects considerations that are beneficical for our general spatial ontology as we are pursuing it here.

Kuhn develops the idea of a semantic reference system by analogy with existing **spatial representation systems**. Enough is known about spatial representation systems employed in GIS, for example, coordinates of various kinds, to enable transformation across differing reference systems. The proposal for semantic reference systems is that similarly effective generalizations be made for semantic content, so that transformations across different reference systems can be automatically supported. This is, in any case, one of the primary goals of inter-operability using ontologies but goes further than most—particularly in the area of spatially-relevant represention s—in considering its concrete realization. This realization is again in a dialect of Haskell, similar to the implementation directions suggested by Frank above.

Kuhn & Raubal (2003) provide a simple but detailed example of the intended use of semantic reference systems built around a domain of cars, roads and car ferries. They show how a domain ontology can be built up in terms of classes and relations and subsequently related to other domains, both by transformation and projection as they define it. They start with a definition of the basic classes presupposed—this they term a **semantic reference frame**. A semantic reference frame is closest in their account to an upper-level ontology as used in ontology design and defined and illustrated at length in our Deliverable D1. Just as with upper-level ontologies, such a frame is intended to be **generic** in that it can extend to cover a variety of domains. In their example, the semantic reference frame contains basic concepts of naming (objects have names and two objects have the same name iff they are the same object), location (objects are 'locatedAt' locations), links, paths and surfaces.

While 'link' describes a basic notion of connectivity, 'paths' and 'surfaces' draw for their definitions on the functional notions of what actions or kinds of interaction the concerned objects support; Kuhn and Raubal relate this both to Gibsonian affordances (cf., also, Jordan, Raubal, Gartrell & Egenhofer 1998) and more directly to Johnson's (1987) **image schemata**. Thus, a path is defined as something providing a 'move to' method, while a surface is something upon which other things can be put, or be taken off from, or which supports things. This part of the definition corresponds well with the functional definition for spatial entities and relations that we (cf. Bateman to appear) and others (e.g., Vandeloise 1985, Aurnague & Vieu 1993) have argued for on linguistic grounds elsewhere; it is particularly interesting to see this perspective being taken here also in a non-linguistic setting. This marks a distinct direction of development for spatial definitions so far not seen in the mainstream spatial ontologies that we have discussed above, although clearly gaining considerable ground in the context of GIS.

The next step of Kuhn and Raubal's example moves into the domain ontology. Here an extension of the semantic frame (upper-level ontology) is given by specializing the generic classes and methods. This domain is intended to illustrate a possible navigation scenario. 'Nodes' are defined as having names and 'locatedAt' is extended to apply to particular objects called 'cars' and 'carFerries'. Moreover, 'edges' are introduced as a further specialization of 'links'. Thus, in the generic ontology, or semantic frame, we have generic classes and relations of objects and links, whereas in the domain ontology, or **semantic reference system**, these resurface as, on the one hand, cars and car ferries and, on the other hand, as edges. Then, for navigation and movement, 'RoadElements' and 'FerryConnections' are defined as a kind of link and, particularly, a kind of path for cars (i.e., cars can use RoadElements to move) and for ferries respectively (i.e., ferries can use FerryConnections to move). Particularly the axiom provided for the latter entity, the FerryConnection-as-path definition, is complicated somewhat by including information that also indicates that cars travel with their respective ferries. This is no in all likelihood a product of making the current example self-sufficient since in general one would want this information to be derivable from the usual properties of moving an object that is itself a container for other objects: this would then be provided by a foundational ontology with broader coverage. However, the fact that a car can be on a ferry, and can be put on the ferry as well as be taken off, is modelled by defining a CarFerry also as an extension of 'surface' from the semantic reference frame. This is sufficient to provide some operational semantics for the entities of this domain.

Both ontological projection and transformation are then illustrated with respect to this model and in strict analogy to their uses in exclusively spatial areas. Projection is **dimensionsional reduction** and the corresponding operation for spatial reference systems is a simplification of some kind. Kuhn and Raubal suggest a simplification in which road elements and ferry connections are no longer distinguished: all that is considered relevant is the fact of connection. This might be useful for planning how to get from A to B: it is generally not relevant whether or not the car moves itself or is carried across a portion of the journey by, for example, a ferry. This 'projection' is then handled in the formalization by appealing only to the definitions for 'nodes' and 'edges' rather than their respective specializations. Transformation is naturally a more general operation and it is not possible in advance to define just what kind of changes will be involved; Kuhn and Raubal suggest appeal to more powerful mathematical mechanisms such as category theory, although their adoption of a higher-order functional programming language already allows them to achieve some of the functionality required without providing the mathematical foundation explicitly. A transformation is then defined by adding additional axioms that express the correspondence: in their example, an axiom is provided that states that cars can move across edges, not restricting further just what those edges are. This is then a transformed system in which no distinction is made between a car moving along a road element or across a ferry connection (in a ferry).

The example as a whole is intended as a proof of concept for the approach rather than a detailed ontology in its own right. The issues discussed relate in interesting ways to other questions that we have raised above, particularly but not only to questions of granularity. However, it is clear that in order to be used on a broader scale, it will be necessary to provide more structuring tools than simple collections of axioms. The particular transformations, projections, reference frames and reference systems will each need its own specified place in a complete framework. Providing such a formal architecture will need to become a proper job of ontological engineering, if not of ontology itself. Here the notion of **formal ontological relations** as proposed by Smith for holding his component SNAP and SPAN ontologies together (see Deliverable D1 and above) may usefully be drawn on: the relations are the logical operations that are necessary for an ontology to work but which themselves do not constitute additions to reality or to the ontology. It is correct, as Kuhn and Raubal state, that this area has been neglected in traditional ontology and ontological engineering: but this situation cannot now continue as the demands on ontology increase.

11.4 Conclusions

One reoccuring theme for all of the approaches seen above is an increased orientation to the Gibsonian notion of affordance. Both Frank (2003*b*) and Kuhn & Raubal (2003) note similarities between Johnson's (1987) **image schemata**, a view of possible manipulations of objects, and the semantic primitives of Wierzbicka (1996) upon which certain aspects of their ontologies are to be constructed; Egenhofer & Rodríguez (1999) define a relation algebra over a set of relations such as moving into, moving out of, etc., also derived from image schemata; and Kuhn (2001) explicitly draws on affordances as such. The view that ontology needs to draw on accounts of human, or more generally, agent, action is thus now widespread.

A further theme is the central role of granularity and, more importantly, being able to move between granularities. We have discussed this extensively above and it is now clear that our formalization of ontologies and mechanisms for interacting with ontologies will need to provide this crucial functionality. Therefore, whereas Kuhn uses this development to distinguish semantic reference systems from ontologies, we will prefer here to extend ontologies and the tools for using them to incorporate more of the functionality for shifting granularities that are clearly required. The term 'semantic reference frame', with its notion of semantics—which, for us, is more suggestive of a linguistic level of description or ontology—will not generally be used in an ontological sense.

Another interesting feature of the more prominent of the recent approaches to employing ontologies within GIS is the role accorded to language. There has been interest for some time in relating GIS-terms to their possible linguistic realizations—this is seen as a contribution to achieving more widely usable person-machine interfaces for GIS (cf., e.g., Freksa 1982, Rashid et al. 1998). But there is now a movement whereby language is seen as a potential source of insight for the ontological distinctions proposed. This appears to be a new move within geographic ontology as both Kuhn (2001) and Frank (2001), Frank (2003*a*) stress the innovative nature of this approach. Also, as we saw above, we have proposals from Frank that his ontology be 'linguistically' justified, from Kuhn for deriving ontological categories from the analysis of texts of a specific domain, as well as Fonseca *et al.* adopting without comment a mixture of domain-specific ontological categories and selections from the lexical database WordNet. In each case it would be important to draw the connection between these approaches and traditional areas of concern in natural language processing and linguistically motivated ontology; this is not, however, done. A consequence is that the use made of linguistic evidence is not always consistent and is certainly under-developed in the sense of only appealing to rather superficial linguistic phenomena. We take this as evidence that the genuine advances made in these new proposals for the tasks of geographical Ontology can only be strengthened by a more thorough recognition of the breadth of evidence that linguistic phenomena in fact provide, which will be a topic that we take up again after the foundation for linguistic ontologies offered in our Deliverable D3.

12 Conclusions and Recommendations

We conclude this tour of approaches to space, spatial representations and entities situated in space by setting out a preliminary framework for integration. This will certainly be modified as investigations proceed and more details of the integrated accounts are taken into account or refined. For the present, it sets out the main frame of reference within our integrative work is being carried out and raising the particular research questions that need to be addressed.

At a number of points in the above discussions the issue of the perspective taken by an ontology, or by someone creating an ontology, has been raised. This is fundamental in several ways: as we have seen in the account of Smith and colleagues, there is ample evidence that there may not be a single ontological perspective that is sufficient for all tasks. This was echoed again by Cohn & Hazariki (2001) with respect to qualitative spatial reasoning. Moreover, within Geographic Information Systems, a range of perspectives has always been required and a major issue is that of reconciling these perspectives and providing mappings across them. We need, therefore, to be able to respect the requirement of multiple perspectives without allowing the entire ontological enterprise to unravel—it is not the case that 'just any' representation is going to be as appropriate as any other. Here we must consider the options very carefully. We present our current position on how perspective choice is to be incorporated into the ontological account in the first subsection below.

We have also seen at numerous points in the above discussion tendencies by ontology designers to make appeals to natural language. We state here once more, therefore, why we believe it important to maintain a distance between our spatial ontology and linguistic issues concerning space and spatial language.

Finally, we need to establish our foundational framework within which, or with respect to which, ontologically viable organizations for space, spatial objects and spatial relationships can be placed: here difficult issues of re-use and modularization are inevitable for an effective treatment; we will address this more fully in deliverable D4 and so restrict our discussion here to a suggestive sketch of a direction to follow, stating proposals that will require subsequent formalization.

12.1 Observer's viewpoint

The perspective taken in ontology-construction is sometimes given a central role. The observer's viewpoint can have consequences for many modelling decisions that are made. We saw this above in Frank's position on the role of the observer. Further examples are provided by, Sowa (2000), for example, who argues that there is a difference between the case of the distinction physical/abstract, which does not depend on an observer, and that of the distinction continuant/occurent, which "depends on the choice of time scale" selected:

> "On a scale of minutes, a glacier is a continuant, and an avalance is an occurrent. But on a scale of centuries, the glacier is a process whose character may be transformed beyond recognition. The changes in a person's facial features are slow enough that friends can recognize an individual as 'the same' over the course of a lifetime. Yet each person gains and loses molecules with every bite of food and every breath of air." (Sowa 2000, p71)

We will take the approach of Sowa's indicated here as a counterpoint to the treatment of perspective and observer that we adopt here. This usefully reflects certain differences in orientation that will serve to make explicit just how our ontological framework is to be constructed.

The linking of ontological category—and for Sowa continuant/occurent is one of the highest distinctions made—to a question of 'choice of time-scale' appears to give the individual observer an important role. But then: just how much of 'ontology' becomes a matter for an observer to decide according to their preferred viewpoint? If someone chooses to 'see' a glacier as an occurrent, then is the glacer an occurrent? If an observer chooses to see the glacier as an abstract entity, although Sowa argues not, is it an abstract entity? Where is the line to be drawn between these cases? Clearly this path needs serious methodological (and ontological) constraint to keep the entire enterprise on track.

At least with respect to occurrents and continuants, the position to be taken can be made much stronger than Sowa's suggestion. In particular, we will follow the proposals of, for example, Simons (1987), Grenon & Smith (2004), Masolo et al. (2003) and others, to make the notion of occurrents *vs.* continuants crucially dependent on the relationship between the entities described and their 'unfolding' in time: is the entity 'entirely there' at a point in time or not? Thus, the glacier *per se* is at each point in time at which it exists entirely there, so is someone's face; this is entirely different to the flow of the glacier down the mountain, or someone's smiling, these are, in sharp contrast, at *no* point in time 'entirely there'. They unfold in time. This does not appear to us to be a matter of observer's perspective. Whether a human observer can see (measure) the flow of the glacier down the mountain is irrelevant: the glacier is there, it is a continuant; its flow (measured or not) is also there, but spread over time, it is an occurrent. Talking about *glacier* in both senses as if the 'same object' were under discussion can then be considered to invite confusion: a prime example of the importance of avoiding conflation of linguistic classifications and ontological categories.

The specification of necessary ontological organization for the areas covered will be the guiding methodology behind our account. Whereas there are possibilities for observer's choice, all of these possibilities must be subject to the ontological constraints that hold in general. It is thus important to be very careful about where and how 'observer's viewpoints' are allowed into the picture. We can talk about a glacier in many ways, we can talk about the flowing of the glacier in many ways, but that should not change its ontological categorization. It is more appropriate to clearly separate out descriptions of the ontology of domains and ways of talking about those domains. This is to move Sowa's selection of viewpoint back where we believe it belongs, i.e., to a question of language and discourse, not of ontology. The mixing of ways of talking about phenomenon and the phenomenon themselves is a particular danger for any ontology that also wishes to deal with semantics: linguistic categories, such as nouns, verbs, etc. do not have any automatic link with ontological distinctions: indeed, even their link with conceptual categories is rather more complex than is often assumed.

The phenomenon of viewpoint is done better justice by the approach to 'observer's perspective' suggested by **granularity**. From a distance, the beach looks like a smooth surface, but when lying on it, it is distinctly grainy; the surface of a table usually looks very smooth, but viewed in terms of its molecules, the boundary is very much harder to describe as anything resembling smooth. A scientific reductionist view of ontology according to which there is only one physical reality answers this question very clearly: everything apart from the reality of the atoms (and their subparticles) is mere appearance. As, however, Smith eloquently argues (e.g., Smith & Mark 2003), this does not really do any justice to the world in which humans live–this is not made up of collections of atoms, it is made up of tables and chairs and beaches. Observer's perspective is thus bound in *ontologically*: ontology concerns itself with reality as interacted with by humans and so the ontological categories involved are the categories out of which the human world is constructed. This essentially Gibsonian view is, again, not a simple matter of 'choice'. We cannot 'choose' whether the chair or table is real or not: they are real if we can sit on them or eat at them (or sell them or repair them). Thus the built-in observer-perspective is that of 'being human' rather than individual choice.

Naturally, within this space of possibilities, there are still choices: with modern scientific equipment, the chair can be seen as a collection of atoms of various kinds. This is indeed a choice of observer's perspective. But, crucially, it is not a choice which makes any difference to the ontological status of the categories involved: they are *all* real. This is what Smith describes as **perspectivalist**. Varying granularity is thus to be seen as a necessary mechanism that ontology must support, but not as a means of determining how larger elements are in fact composed of 'more real' smaller elements. The *fact of* differing observer's perspectives is thus pre-given ontologically; the observer's choice does not construct reality.

> "Partitions are at work ... whenever judgements are effected in relation to the empirical world of what happens and is the case. For a partition to do its work, it needs to have cells large enough to contain the objects that are of interest in the portion of reality which concerns the judging subject, but at the same time these cells must somehow serve to factor out the details which are of no concern. A partition ... is accordingly a device for focusing upon what is salient and for masking what is not salient. ... Thus, importantly, it does not in any way change the reality to which it is applied. This reality, and each of the objects within it, is what and where it is, and it has all its parts and moments, independently of any acts of human fiat and independently of our efforts to understand it theoretically." (Smith & Brogaard 2003)

And so, to take one common suggested counter-example to this position, the observer's choice does not 'construct' an electron as a particle or a wave, that choice merely determines which of the electron's ontological possibilities is revealed. While for Sowa, then, describing something as a 'chair' according to a commonsense granularity must be seen as an observational limitation, for the theory of granular partitions underlying BFO, 'chair' belongs to an ontological selection of partitions and is just as real as the quantum stuff that serves as its physical substrate according to an ontology partitioned according to the purposes of the atomic physicist.

This approach to granularity also applies to perspective shifts more generally. To take another example from the spatial area, if one asks whether a road is a one-dimensional, two-dimensional, or three-dimensional physical object, then the answer would usually depend on *purpose*. As illustrated well by Hobbs,

> "When we are planning a trip, we view it as a line. When we are driving on it, we have to worry about our placement to the right or left, so we think of it as a surface. When we hit a pothole, it becomes a volume for us." (Hobbs 1995, p820)

Sowa again sees this as an indication of choice and relates it to Peirce's Thirdness (to be taken up further in Deliverable D3 in our discussion of semiotics within linguistic ontologies), "since it depends on a triadic relationship between an object, an observer, and the observer's reason for ignoring or discarding certain details" (Sowa 2000, p122): this is categorized according to **Intentionality**.

So, consider a 'road': what is it ontologically? It is a physical object, and physical objects 'in reality' only come with three-dimensions—even though they may be very thin along one or two of those dimensions. But what relations can this entity enter in to? Spatially it can enter into metrical relations of width, length, etc., into topological relations of connection, it can have directedness, and shapes. All of these inhere in the physical object. But the 'road' can also be a way to get from A to B. As such, on Sowa's view, certain aspects of the road's reality can be ignored, only some of its properties are 'relevant': this is the relation to intentionality and purpose; the road may be 'viewed' as a line.

In contrast to this, we can take the position that it is not simply a matter of point of view, or that the road is a line or an approximation. There is an ontology (in fact, as we have seen, several ontologies) of spatial configurations involving directionality and line segments; such an ontology does not offer an approximation to a road, it is an important component of what a road is: ontologies of different granularities then serve to bring some of these aspects of what a road is into the foreground and to leave others in the background. For this reason, then, mapping the road may appropriately leave out aspects of its shape and precise curvatures: as long as sufficient connectivity information is maintained, the map will be judged as appropriate—and, in a very important sense, as *representing something real about the world*. Any account of the ontology of roads that does not include the fact that they can take you from A to B has missed something crucial: to talk of a line representation as an approximation does not therefore give the correct emphases. In contrast, perspectivalism on Smith's account demands that all of these aspects receive their due: they are all 'real' aspects of the world. This leads to a fundamental way of viewing the interconnectedness of the views constructed. They are not 'approximations' to the real reality, they are alternative accounts, with their own specific ontological commitments and requirements, as well as commitments to generic foundational ontological categories at large.

12.2 Spatial language

We have noted at various points above that something that has very explicitly *not* been suggested for integration so far is spatial language. This is not an oversight. The accounts that we have presented have been presented as far as possible in ways that are shorne of any commitments or pronouncements they may make about their relation to language. On the one hand, this is because we deal with language, and spatial language in particular, in more detail in later deliverables; but, on the other hand, it is because we take many proposals that have been made in the ontological and spatial representation tradition to be problematic. In essence the problem can be described as attempting to do too much with the wrong tools. While appropriate for spatial modelling, trying to bring in accounts of spatial language and its semantics within the same apparatus often does violence to the flexibility of observed language use. The consequence of this is a simplified semantics at the cost of an extremely bloated pragmatics, i.e., most cases of 'normal' language use, that is language that is not intended purely geometrically, turns out to require substantial pragmatic apparatus to arrive at actually intended meanings. We take this as a strong indication of a wrong turning and an over-simplification concerning where and how statements about language are to be made.

We maintain that language and ontology, even the semantics of language and ontology, are usually very much further apart than any easy combination of these domains of discourse suggests. This is one of the concerns we have with the ontology/semantics mixture of the approach of Bennett (Section 9 above): we strongly suspect that conflating the very different strata of semiotic abstraction involved creates as many problems as it might appear to solve. Untangling these relationships in a more effective fashion is precisely the task of our deliverables concerning linguistic ontologies and, subsequently, linguistic interaction in general.

We see a, for us, more convincing openness to the complexity of the relationships involved in the account of Casati & Varzi (1999). In their discussion of the problems inherent in such phrases (discussed in the spatial language literature at least since Herskovits (1986), Vandeloise (1985), Aurnague & Vieu (1993) and others) as:

"the fly is in the glass"	*vs.*	"the fly is in amber"
"the bulb is in the socket"	*vs.*	"the bulb is in the box"

Casati and Varzi define a substantial mereotopological apparatus that clarifies considerably the kinds of 'holes' and concave regions *in* which something can be. But they acknowledge that this in no way provides a sufficient explanation for even these very simple everyday usages of the prepositional modifiers 'in+NP':

> "It is apparent that these cases reveal the limits of the approach insofar as it is purely geometric: a full account calls for a step into other territories where pragmatics, or functional and causal factors at large, must be taken into account. Our point is that explicit reference to holes can mark an improvement as far as the *geometric part* of the story goes. True, only some holes count for the purpose of representing containment. But which holes do count is not a question for the geometric analysis of the problem." (Casati & Varzi 1999, p141)

The contributions to spatial ontology that we have collected together in this deliverable are to be seen in this light. We take up the other territories involved under the guise first (in Deliverable D3) of linguistic ontology and, later, for spatial language in particular.

When constructing discourses about entities in the world, we can do this drawing on any of the many alternative aspects of reality. These aspects must have a foundation in particular areas of ontology: that is, we see the semantics of such statements as drawing on how the world is organized. All such discourse can be 'true'. But language is not limited to the true: it is possible to construct discourses which are manifestly not supported by reality. Perhaps this is a necessary property of language—that it is equally able to construct discourses that correspond to the real and discourses that do not. It may even be a necessary property of language as such: i.e., either we have the ability to make non-veridical statements about the world or we do not have a semiotic system with the expressive power necessary for supporting human language. The mechanisms of language necessary for constructing true statements about the world may be sufficiently complex that it must necessarily include the possibility of being wrong. For this reason alone it would be wise to be wary of motivating too many ontological distinctions on linguistic evidence.

The relation between particular linguistic behavior and the partitions on reality that are taken up is a crucial one for our project.

> "Our judgements .... come along with partitions of reality of various sorts, whose type, granularity and scope depend on the contexts in which our judgements are made. ... This relation between judgement and partitions is a complex one ..." (Smith & Brogaard 2003, p38)

Because of this complexity, we consider linguistic and spatial ontologies as separate independent variables, whose precise interrelationship is to be clarified progressively for the duration of the rest of the project.

12.3 Towards a generic foundation for spatial ontologies

Finally we return to to the main task of this deliverable: establishing some groundwork for a foundational spatial ontology that can incorporate previous work in an integrative fashion. Such a foundational ontology needs to provide for the different perspectives that can be taken on phenomena in the world. For space, an excellent starting point appears to be offered by the DOLCE framework as described above (cf. Section 5.1). This is precisely because it *says very little* about how space is to be captured. We see this, just as indicated in our discussion above, as a way of freely parameterizing the particular ontology of space that is to be committed to. Such parameterization is necessary and will be made depending on particular purposes and intentions. Moreover, in order to make that parametric choice, further information may need to be associated with particular ontological perspective choices—**rationales** for adopting one perspective rather than another. This may be in terms of particular properties that a perspective then 'picks out' or generically in terms of overall computational properties of the concerned body of theory. We need to see all of these aspects as important components of a complete ontological account.

12.3.1 An outline of a framework

Thus, first, entities in the world will receive positions with respect to some specified spatial quality regions as described for DOLCE above. Here space is treated as a quality: particular physical endurants are related to categories of physical quality (PQ) including that of 'having a location'. This is a necessary category for physical endurants. The particular locations that may be had are related to entities organized within a space region (S), a subcategory of abstract (AB) physical regions (PR).

Spatial entities are thus bound into space by virtue of a spatial mutual specific dependency relationship (MSD_S: cf. Figure 33). Specific dependency is defined in terms of mutual disjointness and the necessary existence of a spatial dependency between particulars such that it is necessarily the case that those particulars are present in the same setting and at the same time. Being present simply requires that there be some spatial location but does not further specify what that might be. The relevant relations between object, locations and space regions are repeated here for convenience:

$$\text{PQ}\ (\texttt{l\#1}) \wedge \text{qt}\ (\texttt{l\#1}, \texttt{physical entity\#1})$$
$$\text{PR}\ (\texttt{location\#1}) \wedge \text{P}\ (\texttt{location\#1}, \texttt{space region}) \wedge \text{ql}\ (\texttt{location\#1}, \texttt{l\#1}, t)$$

The specification of possibilities within space region is then left free to be filled in by particular schemes.

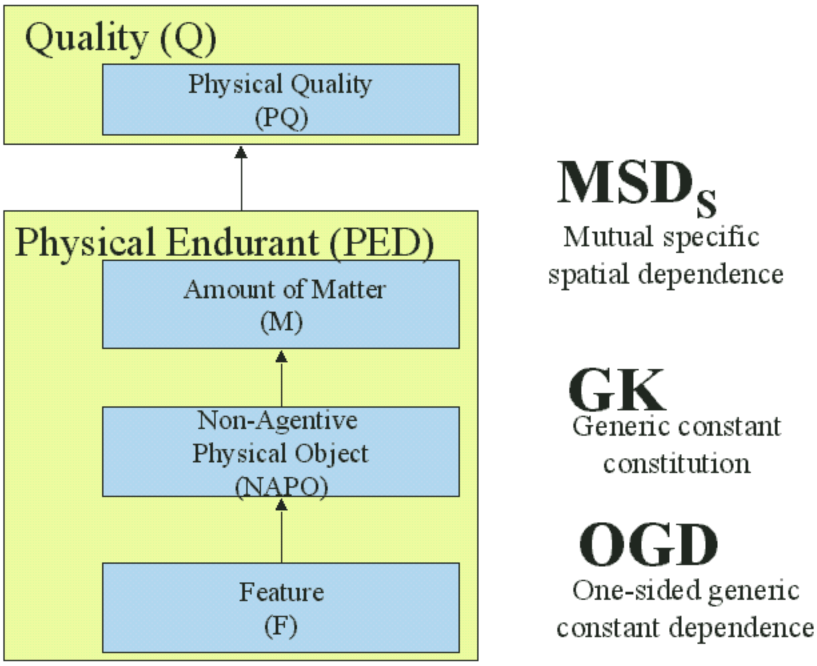

Figure 33: Inter-entity relations within Physical Endurants according to DOLCE

Therefore, while DOLCE does not commit to a particular view of space, it certainy provide the 'glue' for binding physical endurants to any particular perspective on space that we specify. The provision of space as physical regions (PR) makes the statement that space must be structured (at least minimally in that we can distinguish at least two elements in some abstract ordering) but does not place further restriction on that ordering. Candidates for location schemes then include all of the formal accounts of space, of topology, of regions and so on that we have seen above. Any account that specifies a structuring of regions is a candidate for a locational scheme within the foundational ontology. If that account is axiomatized in a manner compatible with the axiomatization provided for the rest of the foundational ontology, then it can directly constitute a parameterized ontology module. This is suggested graphically in Figure 34. The precise mechanisms for achieving this integration are explored in our deliverable D4 and draw extensively on current work elsewhere in the SFB/TR8 involving the formalization of spatial calculi.

For our present purposes, however, we need to follow this line of development a little further by considering just what categories are being linked by the notion 'located at'. It could be suggested that it is perhaps more accurate to state that it is matter and location that are strongly linked, rather than the physical objects that are constituted by matter. Then individuals would only achieve their spatial extension indirectly via their material substrate. An amount of matter always takes up some particular space, and that space can be located; the location would take place in the region layer proposed by Donnelly and Smith.

Such a position can certainly be criticized, however, as favoring a view inherited from physics rather than one that is cognitively or 'naively' motivated. The object that is most salient, e.g., a particular book lying on the table, would be being made less significant as far as location is

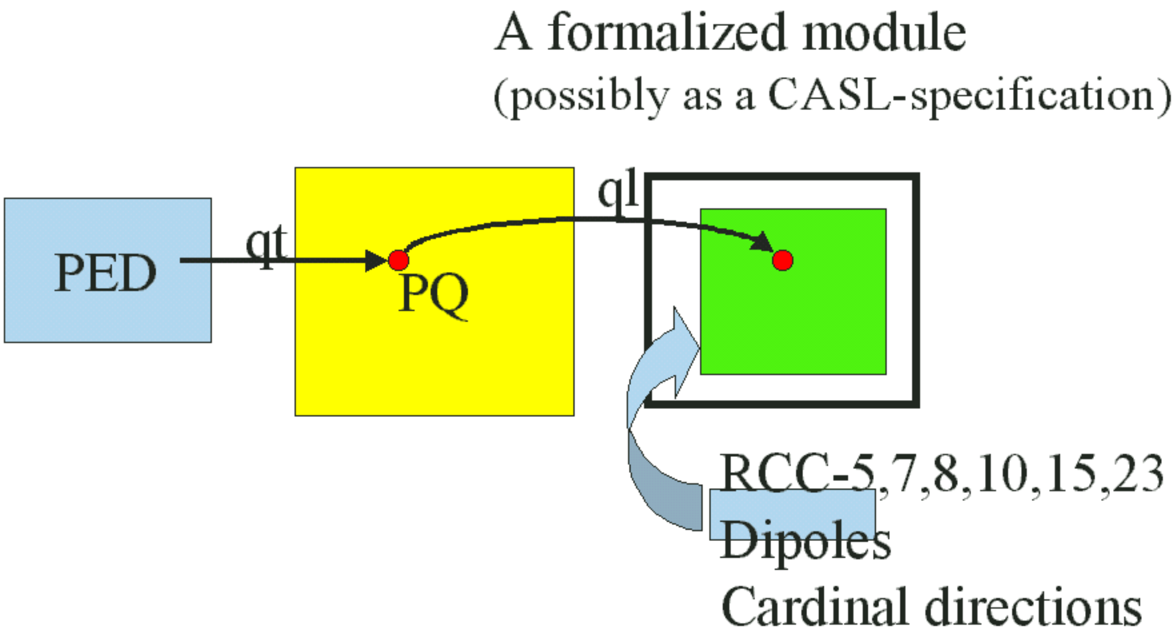

Figure 34: Using space regions as an ontological place for spatial ontology modules

concerned than the matter which constitutes the book. There is then no doubt much in favor of following DOLCE's lead in trying to bind objects and their constituting matter together ontologically sufficiently tightly that we can still sensibly talk of the book's location rather than the location of the matter out of which the book is constituted—which must come out in the axiomatization as equivalent. There are also, however, objects where it is much less clear what their constituting matter actually is—as illustrated with respect to a geographic object, Mount Everest, in the following:

> "An object ... which was delineated, marked out, demarcated, set into relief in this fashion, and thereby also named, Mount Everest exists only as a result of human beliefs and habits. In this sense Mount Everest is, like downtown Sta Barbara—and like professions, religions, human ethnic groups and other similar phenomena ...—a product of socially established beliefs and habits. It is a fiat object. As a portion of geophysical reality, in contrast, that is to say as an aggregate of molecules connected together in space in just this fashion, Mount Everest exists entirely independently and had already existed for a long period of time before we developed our current cognitive abilities." (Smith & Mark 2003, p15)

We do not want to make statements of the location of Mount Everest conditional on having fixed the problem of precisely which amounts of matter are actually meant. This is what is handled in Bennett's approach by adopting a supervaluationist logic that enables statements and inferences to be made even in the face of underlying vagueness. For the moment we will assume that we can move between a physical object and its constituting matter sufficiently unproblematically in order to focus primarily on the nature of their relationships with location.

With respect to the DOLCE indirect relationship between objects and places within a space region, both halves of the relationship appear in need of further refinement but nevertheless offer a good starting point. Although it is not yet clear how these relationships should be formalized in order to bring out the difference between, e.g., colors and locations, the indirection remains useful. Consider, first, BFO's notion of a site. We saw above how BFO allows a specification of location in terms of objects functioning as locations (sites: e.g., a room), suggests distinctive relationships between objects and sites and between sites and locations, and already appears to suggest a characterization of space in terms of a three dimensional space from which extracts of lower dimensionality can be extracted and maintains space as a distinct high-level ontological category in its own right. Site therefore appears to be used in the sense of an object that is functioning as a location; this is not how it is introduced in the BFO description where it is simply stated that there are entities in the world which are sites. But if any object that can be a container is automatically also a site, then this appears to require entities to change their ontological status depending on how we are considering them. This is clearly not a desirable feature of an well-defined ontology.

To avoid this we take the following path. Clearly, physical objects with spatial extent can provide a setting for various kinds of activities. But the fact that this occurs is to a certain extent independent of the entities involved. A church as a building may be used as a church for worshipping activities but the physical object itself does not necessitate those activities. We therefore separate out, first, the notion of X-as-object and X-as-site—the former is a recogniseable physical endurant with physical properties, the latter is a functional location scheme for placing activities. Following this, we can then go further and make any object that has spatial extent potentially serve as a component of a scheme for structuring location that can be utilized for qualitative coordinates. This is then one way of filling out the parameterizable space regions left open in DOLCE. The qualitative coordinates may draw to a greater or lesser extent from the particular activities that the corresponding site supports.

Such objects are also subject to granularity selections, in that a description can be 'in the room', 'in the corner of the room', 'in the drawer of the desk in the corner of the room', etc. Just as with the case with quality regions and color, the *labels* for the qualities are drawn from the quality region (e.g., 'red'): thus the particular location descriptions are similary drawn from the make-up of a space region. We can also at this point move the DOLCE dependent category of places (a type of feature) including such induced locations as 'underneath the table', 'in front of the house', etc. to play a simlar role to sites within a structuring scheme for locations.

These developments relate the description of possible space regions to traditional talk of landmarks and relators: the site is simply an extended kind of landmark that is adopted (at a particular granularity) and the relators are simply the qualitative spatial relations provided (at a particular granularity). Sites appear to be 'container'-landmarks. These characteristics taken together constitute a **location scheme**. We suggest that making extended objects take a primary role in the structuring of location schemes assigns a cognitively plausible central role to objects when considering locations and space. The regions which are taken up by such objects are more implicit (as is the particular color of the rose which, while perceptually immediate, cannot have a simple form of description).

Space is also necessarily structured; but in a way that perhaps differs from the structuring of a conceptual space such as color. There are no observables that can be measured directly and so we prefer to see the structure of space as purely Gibsonian in that it is measured by how we can move in it and how objects can be placed within it. It is only recently, with the advent of global positioning, that location has been made into an 'observable' in this sense for anyone who cares to purchase the necessary devices; previously, this would only be accessible to those trained in astronavigation and is unlikely to have had any strong effect on the human conception of space. Whereas particular colors that may exist (such as the particular red of this particular rose) depend on their bearing objects for their existence, locations depend on their locational scheme and given a scheme all reachable locations exist. We take this to be the case for both Newtonian/Galilean and Leibnitzian views of space. Thus any physical object will always be placed within an entire framework of spatial relationships and the precise characterization of that placement depends on the locational scheme adopted.

From this we can see what is partcularly suggestive for us in BFO's definition of niches as *functional space* (cf. Section 6.1.4); a niche is not just a location in which something is located but instead it brings out functional characteristics of the spatial environment. Indeed, the very same physical space can serve differently in overlapping niches. If space is to be considered in Gibsonian terms, then this would argue that all of space be structured in terms of niches. Since a niche necessarily has a tenant, this would also serve to anchor objects into their locations in a rather stronger fashion than simply being 'located at'. Here the relationship between niches, settings in DOLCE, and the situations/situoids of Heller & Herre (2003) introduced in Deliverable D1 needs to be carefully explored. Niches and layers will also play a role for considerations of movement such as is necessary when we come to discuss routes and actions.

12.3.2 Locating entities

We now have numerous components that are to be combined in a complete model that nevertheless maintains sufficient flexibility to employ a variety of perspectives on space and the world. In order to specify location of an entity e_1 we need the following:

- a selection of an appropriate granular partition of the world that picks out the entity that we wish to locate with respect to other relevant entities, to purpose, and to appropriately related layers;
- a selection of an appropriate space region formalization that brings out or makes available relevant spatial relationships;
- a selection of an appropriate partition over the space region (e.g., from RCC to RCC-8 or RCC-5: cf. Figure 24 above);
- the identification of the location of the entity with respect to the selected space region description.

Moreover, in order to describe a location of some additional entity e_2 or further location in terms of the previously specified entity e_1, we need to select a fitting relation from those

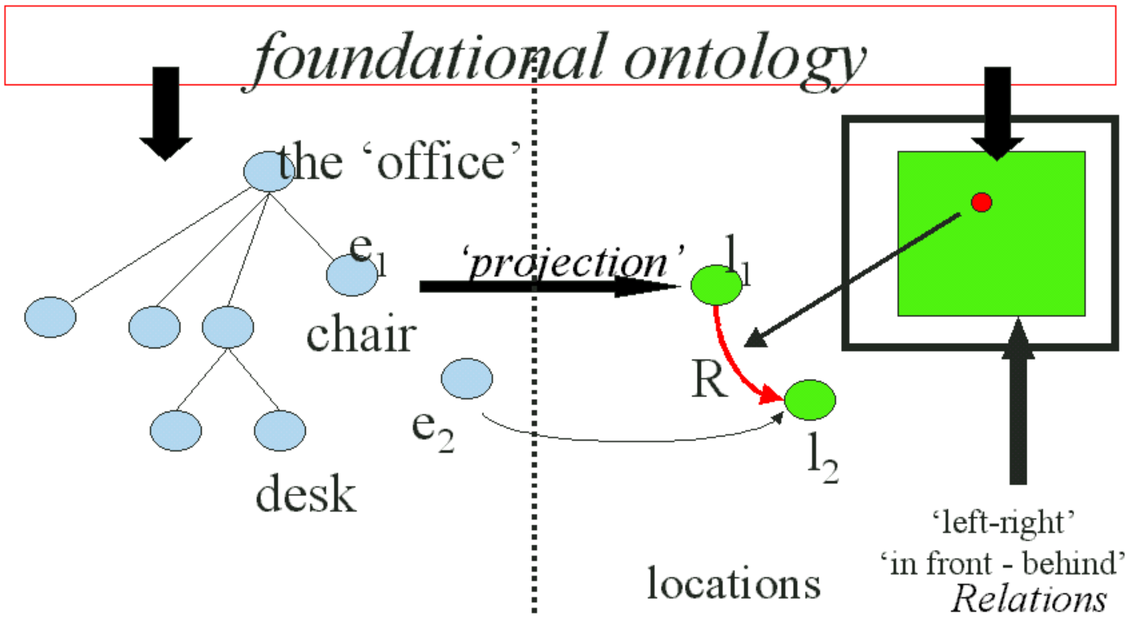

Figure 35: Qualitative entity location in terms of parameterized space regions and relations

made available in the space region. All of these selections may involve more or less precise information as determined by the granularity of the partitions selected and their inherently qualitative nature. This can profitably be seen in terms of Bittner & Smith's (2001*a*) notion of **qualitative coordinates** and is depicted graphically in Figure 35.

This uses as an example an office environment as envisaged for several of the SFB application scenarios. The granularity of the selected partition picks out an office and some items of office furniture, such as chairs and desks. Similarly, the space region is decomposed according to a selected specification involving spatial relations. We can then give relative spatial specifications using those relations (e.g., 'in front of' / 'behind' or connects, etc.). We can also identify locations (e.g., l_1) using the objects that inhabit particular locations: such as the 'chair' (e_1). Then, in order to describe the location of some additional entity that may not already be ordered within the adopted partition (e.g., e_1), we can select a relation (R) from the space region in order to capture that location (l_2). This can then stand behind descriptions such as:

The e_2 is in front of the chair
The e_2 is near the chair

depending on how the space region is structured. All of the categories used are of course to be grounded in the basic categories of the foundational ontology, such as physical endurant, etc., as taken from a broad coverage general ontology of the kind explored in Deliverable D1, extended as necessary by the modular spatial extensions that we have seen here.

References

Allen, J. F. (1984), 'Towards a general theory of action and time', *Artificial Intelligence* **23**, 123–154.

Asher, N. & Vieu, L. (1995), Towards a geometry of common sense: a semantics and a complete axiomatisation of mereotopology, *in* 'Proceedings of the 15th. International Joint Conference on Artificial Intelligence (IJCAI'95)', Morgan Kaufmann, San Mateo, CA, pp. 846–852.

Aurnague, M. & Vieu, L. (1993), A three-level approach to the semantics of space, *in* C. Zelinski-Wibbelt, ed., 'Semantics of Prepositions: From Mental Processing to Natural Language Processing', number 3 *in* 'Natural Language Processing', Mouton de Gruyter, Berlin, pp. 393–439.

Bateman, J. A. (1995), 'On the relationship between ontology construction and natural language: a socio-semiotic view', *International Journal of Human-Computer Studies* **43**(5/6), 929–944.

Bateman, J. A. (to appear), Linguistic interaction as ontological mediation, *in* D. Zaefferer & A. Schalley, eds, 'Ontolinguistics. How ontological status shapes the linguistic coding of concepts', Mouton de Gruyter.

Bateman, J. A. & Farrar, S. (2005), Modelling models of robot navigation using formal spatial ontology, *in* C. Freksa, M. Knauff, B. Krieg-Brückner, B. Nebel & T. Barkowsky, eds, 'Spatial Cognition IV: Reasoning, Action, Interaction. International Conference Spatial Cognition 2004, Frauenchiemsee, Germany, October 2004, Proceedings', Springer, Berlin, Heidelberg, pp. 366–389.

Bennett, B. (1994), Spatial reasoning with propositional logic, *in* 'Proceedings of Principles of Knowledge Representation and Reasoning: Proceedings of the Fourth International Conference (KR94)', Morgan Kaufmann, San Mateo, CA, pp. 51–62.

Bennett, B. (2001*a*), Space, time, matter and things, *in* C. Welty & B. Smith, eds, 'Formal Ontology in Information Systems: Proceedings of the 2nd International Conference on Formal Ontology in Information Systems', ACM Press, New York, pp. 105–116.
*ftp://agora.leeds.ac.uk/scs/doc/srg/bennett-fois01.ps

Bennett, B. (2001*b*), 'What is a forest? on the vagueness of certain geographic concepts', *Topoi* **20**(2), 189–201.

Bennett, B. (2002), Physical objects, identity and vagueness, *in* D. Fensel, D. McGuinness & M.-A. Williams, eds, 'Principles of Knowledge Representation and Reasoning: Proceedings of the Eighth International Conference (KR2002)', Morgan Kaufmann, San Francisco, CA.

Bennett, B., Cohn, A., Torrini, P. & Hazariki, S. (2000), A foundation for region-based qualitative geometry, *in* W. Horn, ed., 'Proceedings of ECAI-2000', pp. 204–208.

Bennett, B. & Galton, A. (2001), A versatile representation for time and events, *in* 'Proceedings of the 5th Symposium on logical formalizations of commonsense reasoning (Commonsense 2001)', pp. 43–52.

Bennett, B., Wolter, F. & Zakharyaschev, M. (2002), 'Multi-dimensional modal logic as a framework for spatio-temporal reasoning', *Applied Intelligence* **17**(3), 239–251.

Bittner, T. & Smith, B. (2001*a*), A taxonomy of granular partitions, *in* 'Proceedings of the Conference on Spatial Information Theory - COSIT 2001', Lecture Notes in Computer Science, Springer-Verlag, Berlin-Heidelberg, pp. 28–43.
*http://people.ifomis.uni-leipzig.de/thomas.bittner/tp.pdf

Bittner, T. & Smith, B. (2001*b*), Vagueness and granular partitions, *in* 'Proceedings of the Conference on Formal Ontology in Information Systems – FOIS2001', Sheridan Press, pp. 309–321.
*http://people.ifomis.uni-leipzig.de/thomas.bittner/bittner_smith_fois2001.pdf

Bittner, T. & Smith, B. (2003*a*), Granular spatio-temporal ontologies, *in* '2003 AAAI Symposium: Foundations and Applications of Spatio-Temporal Reasoning (FASTR)'.
*http://people.ifomis.uni-leipzig.de/thomas.bittner/SSS303TBittner.pdf

Bittner, T. & Smith, B. (2003*b*), A theory of granular partitions, *in* M. Duckham, M. F. Goodchild & M. F. Worboys, eds, 'Foundations of Geographic Information Science', Taylor and Francis, London, pp. 117–151.
*http://wings.buffalo.edu/philosophy/faculty/smith/articles/partitions.pdf

Bittner, T. & Stell, J. G. (2002), 'Vagueness and rough location', *Geoinformatica* **6**, 99–121.
*http://people.ifomis.uni-leipzig.de/thomas.bittner/BittnerStellGEIN101-99.pdf

Borgo, S., Guarino, N. & Masolo, C. (1996*a*), A pointless theory of space beased on strong congruence and connection, *in* L. Aiello & J. Doyle, eds, 'Principles of knowledge representation and reasoning, Proceedings of the 5th International Conference KR96', Morgan Kaufmann, San Mateo, CA.

Borgo, S., Guarino, N. & Masolo, C. (1996*b*), Stratified ontologies: the case of physical objects, *in* 'Proceedings of the Workshop on Ontological Engineering at ECAI'96', Budapest, Hungary.

Borgo, S., Guarino, N. & Masolo, C. (1997), An ontological theory of physical objects, *in* L. Ironi, ed., 'Proceedings of 11th international workshop on qualitative reasoning (QR'97)', Cortona, Italy.

Casati, R. & Varzi, A. C. (1994), *Holes and other superficialities*, MIT Press (Bradford Books), Cambridge, MA and London.

Casati, R. & Varzi, A. C. (1999), *Parts and places: the structures of spatial representation*, MIT Press (Bradford Books), Cambridge, MA and London.

Clementini, E., Felici, P. & Hernandez, D. (1997), 'Qualitative representation of positional information', *Artificial Intelligence* **95**, 317–356.

Cohn, A. G., Bennett, B., Gooday, J. & Gotts, N. M. (1997), Representing and reasoning with qualitative spatial relations, *in* O. Stock, ed., 'Spatial and temporal reasoning', Kluwer Academic Publishers, Dordrecht, pp. 97–132.

Cohn, A. & Hazariki, S. (2001), 'Qualitative spatial representation and reasoning: an overview', *Fundamenta Informaticae* **43**, 2–32.

Cohn, A., Randell, D. & Cui, Z. (1995), 'Taxonomies of logically defined qualitative spatial relations', *International Journal of Human-Computer Studies* **43**(5/6), 831–846.

Cristani, M. (1999), 'The complexity of reasoning about spatial congruence', *Journal of Artificial Intelligence Research* **11**, 361–390.

Cycorp (2004*a*), Cyc tutorial, Technical report.
*http://www.cyc.com/cycdoc/course/collections-module.html

Cycorp (2004*b*), Opencyc 0.7.0, Technical report.
*http://www.opencyc.org

Davis, E., Gotts, N. & Cohn, A. (1999), 'Constraint networks of topological relations and convexity', *Constraints* **4**(3), 241–280.

de Kleer, J. & Brown, J. S. (1985), A qualitative physics based on confluence, *in* J. R. Hobbs & R. C. Moore, eds, 'Formal theories of the commonsense world', Ablex Publishing Corporation, Norwood, New Jersey, pp. 109–184.

de Saussure, F. (1959/1915), *Course in General Linguistics*, Peter Owen Ltd., London. (translated by W.Baskin).

Donnelly, M. & Smith, B. (2003), Layers: A new approach to locating objects in space, *in* W. Kuhn, M. F. Worboys & S. Timpf, eds, 'Spatial Information Theory: Foundations of Geographic Information Science', number 2825 *in* 'Lecture Notes in Computer Science', Springer, Berlin, pp. 50–65. *http://ontology.buffalo.edu/geo/Layers.pdf

Düntsch, I., Wang, H. & McCloskey, S. (1998), A relation-algebra approach to the region connection calculus, Technical report, School of Information and Software Engineering, University of Ulster.

Dylla, F. & Wallgrün, J. O. (n.d.), On generalizing orientation information in $\mathcal{OPRA}_{\Updownarrow}$, *in* 'Proceedings of the 29th German Conference on Artificial Intelligence (KI 2006)', Lecture Notes in Artificial Intelligence, Springer, Berlin.

Egenhofer, M. (1991), Reasoning about binary topological relations, *in* O. Günther & H. Schek, eds, 'Proceedings of the Second Symposium on Large Scaled Spatial Databases', Springer, Berlin, pp. 143–160.

Egenhofer, M. & Herring, J. (1991), Categorizing binary topological relationships between regions, lines and points in geographic databases, *in* M. Egenhofer, J. Herring, T. Smith & K. Park, eds, 'A framework for the definition of topological relationships and an algebraic approach to spatial reasoning within this framework', number 91-7 *in* 'NCGIA Technical Report', National Center for Geographic Information and Analysis, Santa Barbara, CA.

Egenhofer, M. & Mark, D. (1995), Naive geography, *in* A. U. Frank & W. Kuhn, eds, 'Spatial Information Theory: a theoretical basis for GIS', number 988 *in* 'Lecture Notes in Computer Science', Springer, Berlin/Heidelberg/New York, pp. 1–15. International Conference COSIT'95.

Egenhofer, M. & Rodríguez, A. (1999), 'Relation algebras over containers and surfaces: An ontological study of a room space', *Spatial Cognition and Computation* **1**(2), 155–180.

Eschenbach, C. (2001), Viewing composition tables as axiomatic systems, *in* C. Welty & B. Smith, eds, 'Formal Ontology in Information Systems. Collected Papers from the Second International conference', ACM Press, New York, pp. 93–104.

Eschenbach, C. & Heydrich, W. (1995), 'Classical mereology and restricted domains', *International Journal of Human-Computer Studies* **43**, 723–740.

Eschenbach, C. & Kulik, L. (1997), An axiomatic approach to the spatial relations underlying *left-right* and *in front of-behind*, *in* G. Brewska, C. Habel & B. Nebel, eds, 'KI-97 Advances in Artificial Intelligence', Springer, Berlin, pp. 207–218.

Eschenbach, C., Tschander, L., Habel, C. & Kulik, L. (2000), Lexical specification of paths, *in* C. Freksa, W. Brauer, C. Habel & K. Wender, eds, 'Spatial Cognition II - An interdisciplinary approach to representing and processing spatial knowledge', Springer, Berlin, Heidelberg, pp. 127–144.
*http://link.springer-ny.com/link/service/series/0558/tocs/t1849.htm

Fine, K. (1975), 'Vagueness, truth and logic', *Synthese* **30**, 265–300.

Fine, K. (1995), 'Ontological dependence', *Proceedings of the Aristotelian society* **95**, 269–290.

Fonseca, F., Egenhofer, M., Davis, C. & Câmara, G. (2002), 'Semantic granularity in ontology-driven geographic information systems', *Annals of Mathematics and Artificial Intelligence* **36**(1–2), 121–151.

Fonseca, F. T., Egenhofer, M. J., Agouris, P. & Câmara, G. (2002), 'Using ontologies for integrated geographic information systems', *Transactions in GIS* **6**(3).

Frank, A. (1991), Qualitative spatial reasoning with cardinal directions, *in* 'Proceedings of the Seventh Austrian Conference on Artificial Intelligence', Springer, Berlin, Heidelberg, New York.

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

Frank, A. (1996), 'Qualitative spatial reasoning: cardinal directions as an example', *IJGIS* **10**(3), 269–290.

Frank, A. (2003*a*), A linguistically justified proposal for a spatio-temporal ontology', *in* 'Proceedings of the COSIT03 international conference'. Position paper in COSIT-03 workshop.

Frank, A. (2003*b*), Ontology for spatio-temporal databases, *in* M. Koubarakis, ed., 'Spatiotemporal Databases: The Chorochronos Approach', Lecture Notes in Computer Science, Springer-Verlag, Berlin, pp. 9–78.

Frank, A. & Kuhn, W. (1986), Cell graphs: a provably correct method for the storage of geometry, *in* 'Proceedings of the 2nd. international symposium on spatial data handling', Seattle, Washington, pp. 411–436.

Frank, A. & Kuhn, W. (1999), A specification language for interoperable GIS, *in* M. Goodchild, M. Egenhofer, R. Fegeas & C. Kottmann, eds, 'Interoperating geographic information systems', Kluwer, Norwell, MA, pp. 123–132.

Frank, A. U. (1997), Spatial ontology: a geographical information point of view, *in* O. Stock, ed., 'Spatial and temporal reasoning', Kluwer Academic Publishers, Dordrecht, pp. 135–153.

Frank, A. U. (2001), 'Tiers of ontology and consistency constraints in geographical information systems', *International Journal of Geographical Information Science* **15**(7), 667–678.

Freksa, C. (1982), Linguistic description of human judgements in expert systems and in the 'soft' sciences, *in* M. Gupta & E. Sanchez, eds, 'Approximate reasoning in decision analysis', North-Holland Publishing Company, pp. 297–305.

Freksa, C. (1991), Conceptual neighborhood and its role in temporal and spatial reasoning, *in* M. Singh & L. Travé-Massuyès, eds, 'Decision Support Systems and Qualitative Reasoning', North-Holland, Amsterdam, pp. 181–187.

Freksa, C. (1992*a*), 'Temporal reasoning based on semi-intervals', *Artificial Intelligence* **54**, 199–227.

Freksa, C. (1992*b*), Using orientation information for qualitative spatial reasoning, *in* A. U. Frank, I. Campari & U. Formentini, eds, 'Theories and methods of spatio-temporal reasoning in geographic space', Vol. 639 of *LNCS*, Springer, Berlin, pp. 162–178.

Gärdenfors, P. (2000), *Conceptual spaces: the geometry of thought*, MIT Press, Cambridge, MA.

Garfinkel, H. (1972), Studies in the routine grounds of everyday activities, *in* D. Sudnow, ed., 'Studies in Social Interaction', The Free Press, New York, pp. 1–30.

Gerevini, A. & Renz, J. (1998), Combining topological and qualitative size constraints for spatial reasoning, *in* 'Proceedings of the 4th international conference on principles and practice of constraint programming (CP'98)', Springer-Verlag, Berlin.

Gibson, J. (1977), The theory of affordances, *in* R.Shaw & J.Brandsford, eds, 'Perceiving, Acting, and Knowing: Toward and Ecological Psychology', Erlbaum, Hillsdale, NJ, pp. 62–82.

Gottfried, B. (2004), Reasoning about intervals in two dimensions, *in* 'IEEE International Conference on Systems, Man and Cybernetics', Omnipress, The Hague, pp. 5324–5332.

Gottfried, B. (2005), Global feature schemes for qualitative shape descriptions, *in* 'Proceedings of the IJCAI-05 Workshop on spatial and temporal reasoning.', Edinburgh, Scotland.

Grenon, P. (2003), Tucking RCC in Cyc's ontological bed, *in* G. G. & W. T., eds, 'IJCAI-03 – Eighteenth International Joint Conference in Artificial Intelligence', Morgan Kaufmann, San Francisco, pp. 894–899.
*http://people.ifomis.uni-leipzig.de/pierre.grenon/Downloads/Grenon-IJCAI03.pdf

Grenon, P. & Smith, B. (2004), 'SNAP and SPAN: Towards dynamic spatial ontology', *Spatial Cognition and Computation* **4**(1), 69–103.
*http://ontology.buffalo.edu/smith/articles/SNAP_SPAN.pdf

Grigni, M., Papadias, D. & Papadimitriou, C. (1995), Topological inference, *in* 'Proceedings of the 14th Int. Joint Conf. of Artificial Intelligence (IJCAI '95)'.

Guarino, N. (1998), Formal ontology and information systems, *in* N. Guarino, ed., 'Formal Ontology in Information Systems', IOS Press, Amsterdam, pp. 3–18.

Guarino, N. & Welty, C. (2002), 'Evaluating ontological decisions with OntoClean', *Communications of the ACM* **45**(2), 61–65.

Guarino, N. & Welty, C. (2004), An overview of OntoClean, *in* S. Staab & R. Studer, eds, 'Handbook on Ontologies', Springer-Verlag, Heidelberg and Berlin.

Gupta, A. (1980), *The logic of common nouns: an investigation in quantified modal logic*, Yale University Press, New Haven.

Habel, C. & Eschenbach, C. (1997), Abstract structures in spatial cognition, *in* C. Freksa, M. Jantzen & R. Valk, eds, 'Foundations of Computer Science. Potential – Theory – Cognition', Springer, Berlin, pp. 369–378.
*http://link.springer-ny.com/link/service/series/0558/tocs/t1337.htm

Hayes, P. J. (1979), The naive physics manifesto, *in* D. Michie, ed., 'Expert systems in the microelectronic age', Edinburgh University Press, Edinburgh, Scotland.

Heller, B. & Herre, H. (2003), Formal ontology and principles of GOL, OntoMed Report 1, Institute for Medical Informatics, Statistics and Epidemiology (IMISE), University of Leipzig, Germany.
*http://www.onto-med.de

Herskovits, A. (1986), *Language and Spatial Cognition: an interdisciplinary study of the prepositions in English*, Studies in Natural Language Processing, Cambridge University Press, London.

Hobbs, J. R. (1995), 'Sketch of an ontology underlying the way we talk about the world', *International Journal of Human-Computer Studies* **43**(5/6), 819–830.

Hornsby, K. (1999), Identity-based reasoning about spatial-temporal change, Technical report, Spatial Information Science and Engineering, University of Maine, Orono.

Isli, A. & Cohn, A. G. (1998), An algebra for cyclic ordering of 2d orientations, *in* 'Proceedings of AAAI-98', MIT Press, Madison, WI, pp. 643–649.

Isli, A., Museros, L., Cabedo, L. M., Barkowsky, T. & Moratz, R. (2000), A topological calculus for cartographic entities, *in* C. Freksa, W. Brauer, C. Habel & K. Wender, eds, 'Spatial Cognition II - Integrating Abstract Theories, Empirical Studies, Formal Methods, and Practical Applications', Springer, Berlin, pp. 225–238.
*http://link.springer-ny.com/link/service/series/0558/tocs/t1849.htm

Johnson, M. (1987), *The body in the mind*, University of Chicago Press, Chicago, Il.

Jordan, T., Raubal, M., Gartrell, B. & Egenhofer, M. (1998), An affordance-based model of place in GIS, *in* 'Proceedings of the 8th International Symposium on Spatial Data Handling (SDH'98)', Vancouver, Canada.
*citeseer.ist.psu.edu/jordan98affordancebased.html

Kent, R. E. (2000), The information flow foundation for conceptual knowledge organization, *in* 'Proceedings of the 6th. international conference of the international society for knowledge organization (ISKO)', Toronto, Canada.

Knauff, M., Rauh, R. & Renz, J. (1997), A cognitive assessment of topological spatial relations: Results from an empirical investigation, *in* 'Proceedings of the 3rd International Conference on Spatial Information Theory (COSIT97)'.

Krieg-Brückner, B., Frese, U., Lüttich, K., Mandel, C., Mossakowski, T. & Ross, R. (2005), Specification of an Ontology for Route Graphs, *in* C. Freksa, M. Knauff, B. Krieg-Brückner, B. Nebel & T. Barkowsky, eds, 'Spatial Cognition IV: Reasoning, Action, Interaction. International Conference Spatial Cognition 2004, Frauenchiemsee, Germany, October 2004, Proceedings', Springer, Berlin, Heidelberg, pp. 390–412.

Kuhn, W. (2001), 'Ontologies in support of activities in geographical space', *International Journal of Geographical Information Science* **15**(7), 679–687.

Kuhn, W. (2003), 'Semantic reference systems', *International Journal of Geographical Information Science* **17**(5), 405–409.

Kuhn, W. & Raubal, M. (2003), Implementing semantic reference systems, *in* M.F.Gould, ed., '6th AGILE Conference on Geographic Information Science'.

Kuipers, B. (1978), 'Modelling spatial knowledge', *Cognitive Science* **2**(2), 129–154.

Kuipers, B. (1998), A hierarchy of qualitative representations for space, *in* C. Freksa, C. Habel & K. Wender, eds, 'Spatial Cognition I - An interdisciplinary approach to representing and processing spatial knowledge', Springer, Berlin.

Kuipers, B. J. (1977), 'Representing knowledge of large-scale space'.

Kuipers, B. & Levitt, T. (1988), 'Navigation and mapping in a large-scale space', *AI Magazine* **9**(2), 25–43. Reprinted in *Advances in Spatial Reasoning*, Volume 2, Su-shing Chen (Ed.), Norwood NJ: Ablex Publishing, 1990.

Lehmann, F. & Cohn, A. G. (1994), The EGG/YOLK reliability hierarchy: semantic data integration using sorts with prototypes, *in* 'Proceedings of the Conference on Information Knowledge Management', ACM Press, pp. 272–279.

Levinson, S. (1996), Frames of reference and Molyneux's question: Crosslinguistic evidence, *in* P. Bloom, M. Peterson, L. Nadel & M. Garrett, eds, 'Language and Space', MIT Press, Cambridge, MA, pp. 109–169.

Levinson, S. C. (2003), *Space in language and cognition: explorations in cognitive diversity*, Cambridge University Press, Cambridge.

Lewis, D. K. (1986), *The plurality of worlds*, Blackwell, Oxford.

Lizogat, G. (1998), 'Reasoning about cardinal directions', *Journal of visual languages and computing* **9**, 137–168.

Loebe, F. (2003), An analysis of roles: towards ontology-based modelling, OntoMed Report 6, Institute for Medical Informatics, Statistics and Epidemiology (IMISE), University of Leipzig, Germany. *http://www.onto-med.de

Longley, P., Goodchild, M., Maguire, D. & Rhind, D. (2001), *Geographic Information Systems and Science*, Wiley and sons, Chichester, U.K.

Masolo, C., Borgo, S., Gangemi, A., Guarino, N. & Oltramari, A. (2003), Ontologies library (final), WonderWeb Deliverable D18, ISTC-CNR, Padova, Italy.

Masolo, C., Borgo, S., Gangemi, A., Guarino, N., Oltramari, A. & Schneider, L. (2002), The WonderWeb library of foundational ontologies: preliminary report, WonderWeb Deliverable D17, ISTC-CNR, Padova, Italy.

Masolo, C. & Vieu, L. (1999), Atomicity vs. infinite divisibility of space, *in* C. Freksa & D. Mark, eds, 'Spatial Information theory. Proceedings of COSIT'99', number 1661 *in* 'LNCS', Springer-Verlag, Berlin, pp. 235–250.

Moratz, R., Dylla, F. & Frommberger, L. (2005), A Relative Orientation Algebra with Adjustable Granularity, *in* 'Proceedings of the Workshop on Agents in Real-Time and Dynamic Environments (IJCAI 05)'.

Moratz, R., Renz, J. & Wolter, D. (2000), Qualitative spatial reasoning about line segments, *in* W. Horn, ed., 'Proceedings of the 14th European Conference on Artificial Intelligence (ECAI00)', IOS Press, Amsterdam.

Mulligan, K. (1999), Perception, predicates and particulars, *in* D. Fisette, ed., 'Consciousness and intentionality: models and modalities of attribution', Kluwer, Dordrecht / Boston / London, pp. 163–194.

Nebel, B. (1995), Computational properties of qualitative spatial reasoning: First results, *in* I. Wachsmuth, C.-R. Rollinger & W. Brauer, eds, 'KI-95: Advances in Artificial Intelligence', Springer, Heidelberg, pp. 233–244.
*http://www.informatik.uni-freiburg.de/ sppraum/

Niles, I. & Pease, A. (2001*a*), Origins of the Standard Upper Merged Ontology: A Proposal for the IEEE Standard Upper Ontology, *in* 'Working Notes of the IJCAI-2001 Workshop on the IEEE Standard Upper Ontology', pp. 37–42.
*http://projects.teknowlege.com/IJCAI01/index.html

Niles, I. & Pease, A. (2001*b*), Toward a standard upper ontology, *in* C. Welty & B. Smith, eds, 'Proceedings of the 2nd International Conference on Formal Ontology in Information Systems (FOIS-2001)', Association for Computing Machinery, Ogunquit, Maine.
*http://projects.teknowledge.com/HPKB/Publications/FOIS.pdf

Poli, R. (1998), 'Levels', *Axiomathes* **9**(1-2), 197–211.

Poli, R. (2001), 'The basic problem of the theory of levels of reality', *Axiomathes* **12**, 261–283.

Pullar, D. & Egenhofer, M. (1988), Toward formal definitions of topological relations among spatial objects, *in* 'Proceedings of the 3rd International Symposium on Spatial Data Handling', Sydney, Australia, pp. 165–176.

Randell, D., Cohn, A. & Cui, Z. (1992), Computing transitivity tables: A challenge for automated theorem provers, *in* 'CADE 1992', pp. 786–790.

Randell, D., Cui, Z. & Cohn, A. (1992), A spatial logic based on regions and connection, *in* 'Proceedings of the 3rd. International Conference on Knowledge Representation and Reasoning', Morgan Kaufmann, San Mateo, pp. 165–176.

Rashid, A., Sharif, B., Egenhofer, M. & Mark, D. M. (1998), 'Natural-language spatial relations between linear and areal objects: the topology and metric of English-language terms', *International Journal of Geographical Information Science* **12**(3), 215–246.

Renz, J. (1999), Maximal tractable fragments of the region connection calculus: A complete analysis, *in* T. Dean, ed., 'Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99)', Morgan Kaufman Publishers, San Francisco, CA, pp. 448–454.

Renz, J. & Mitra, D. (2004), Qualitative direction calculi with arbitrary granularity, *in* C. Zhang, H. W. Guesgen & W. K. Yeap, eds, 'Proceedings of the 8th Pacific Rim International Conference on Artificial Intelligence on AI (PRICAI)', Springer-Verlag, pp. 65–74.

Renz, J. & Nebel, B. (1998), Efficient methods for qualitative spatial reasoning, *in* H. Prade, ed., 'Proc. of the 13th European Conference on Artificial Intelligence (ECAI-98)', John Wiley & Sons, Chichester, pp. 562–566.

Renz, J. & Nebel, B. (1999), 'On the complexity of qualitative spatial reasoning: a maximal tractable fragment of the region connection calculus', *Artificial Intelligence* **108**(1–2), 69–123.

Rigaux, P. & Scholl, M. (1995), Multi-scale partitions: applications to spatial and statistical databases, *in* M. Egenhofer & J. Herrings, eds, 'Advances in Spatial databases (SSD'95)', number 951 *in* 'Lecture Notes in Computer Science', Springer-Verlag, Berlin.

Schlieder, C. (1995), Reasoning about ordering, *in* A. Frank & W. Kuhn, eds, 'Proceedings of the 3rd Conference on Spatial Information Theory (COSIT)', Springer, Heidelberg, Berlin, pp. 341–349.

Schlieder, C. (1996*a*), Ordering information and symbolic projection, *in* S. C. et al., ed., 'Intelligent image database systems', World Scientific, Singapore, pp. 115–140.

Schlieder, C. (1996*b*), Qualitative shape representation, *in* A. Frank, ed., 'Spatial conceptual models for geographic objects with undetermined boundaries', Taylor & Francis, London, pp. 123–140.

Schlieder, C., Vögele, T. & Werner, A. (2001), Location modeling for intentional behaviour in spatial partonomies, *in* 'Proceedings of the Workshop on Location Modeling for Ubiquitous Computing (Ubicomp 2001)', Atlanta, Georgia.

Schulz, S. & Hahn, U. (2001), Mereotopological reasoning about parts and (w)holes in bio-ontologies, *in* C. Welty & B. Smith, eds, 'Proceedings of the 2nd International Conference on Formal Ontology in Information Systems', ACM Press, New York, pp. 210–221.

Schulz, S., Hahn, U. & Romacker, M. (2000), Modeling anatomical spatial relations with description logics, *in* J. Overhage, ed., 'AMIA-2000 - Proceedings of the annual symposium of the American Medical Informatics Association. Converging Information, Technology and Health Care', Hanley and Belfus, Philadelphia, PA, pp. 779–783.

Scivos, A. & Nebel, B. (2001), Double-crossing: Decidability and computational complexity of a qualitative calculus for navigation, *in* 'Proc. COSIT-2001', Springer.
*http://www.informatik.uni-freiburg.de/ sppraum/

Simons, P. (1987), *Parts: a study in ontology*, Clarendon Press, Oxford.

Smith, B. (1996), 'Mereotopology: a theory of parts and boundaries', *Data and knowledge engineering* **20**, 287–303.

Smith, B. (2001), 'Fiat objects', *Topoi* **20**(2), 131–148.
*http://ontology.buffalo.edu/smith/articles/fiat.htm

Smith, B. & Brogaard, B. (2002), 'Quantum mereotopology', *Annals of Mathematics and Artificial Intelligence* **35**(1-2), 153–175.
*http://ontology.buffalo.edu/smith/articles/QM.htm

Smith, B. & Brogaard, B. (2003), 'A unified theory of truth and reference', *Loqique et Analyse* **43**(169-170), 49–93.
*http://ontology.buffalo.edu/smith/articles/truthandreference.htm

Smith, B. & Grenon, P. (2004), 'The cornucopia of formal-ontological relations', *Dialectica* **58**(3), 279–296.
*http://wings.buffalo.edu/philosophy/ontology/smith/articles/cornucopia.pdf

Smith, B. & Mark, D. M. (2003), 'Do mountains exist? towards an ontology of landforms', *Environment and Planning B (Planning and Design)* **30**(3), 411–427. *http://ontology.buffalo.edu/smith/articles/Mountains.htm

Smith, B. & Varzi, A. C. (1999), 'The Niche', *Noûs* **33**(2), 214–238.

Smith, B. & Varzi, A. C. (2000), 'Fiat and bona fide boundaries', *Philosophy and phenomenological research* **60**(2), 401–420.

Sowa, J. F. (2000), *Knowledge Representation: logical, philosophical, and compuational foundations*, Brooks/Cole, Pacific Grove, CA.

Van de Weghe, N., Cohn, A., Maeyer, P. D. & Witlox, F. (2005), 'Representing moving objects in computer based expert systems: The overtake event example', *Expert Systems with Applications* **29**(4), 977–983.

Van de Weghe, N., Cohn, A., Tré, G. D. & Maeyer, P. D. (2005), 'A qualitative trajectory calculus as a basis for representing moving objects in geographical information systems', *Control and Cybernetics* .

Van de Weghe, N., Kuijpers, B., Bogaert, P. & Maeyer, P. D. (2005), A qualitative trajectory calculus and the composition of its relations, *in* M. A. Rodríguez, I. F. Cruz, M. J. Egenhofer & S. Levashkin, eds, 'GeoS 2005: GeoSpatial Semantics, First International Conference, Proceedings', Vol. 3799 of *Lecture Notes in Computer Science*, Springer, Berlin, Heidelberg, pp. 60–76.

Vandeloise, C. (1985), *Description of Space in French*, number 150 *in* 'A', L.A.U.D.T.

Vieu, L. (1997), Spatial representation and reasoning in artificial intelligence, *in* O. Stock, ed., 'Spatial and temporal reasoning', Kluwer Academic Publishers, Dordrecht, pp. 5–41.

Visser, U., Stuckenschmidt, H. & Vögele, T. (2001), 'Ontologies for geographic information processing', *Computers and Geoscience* .

Wierzbicka, A. (1996), *Semantics – Primes and Universals*, Oxford University Press, Oxford.