Tucking RCC in Cyc's Ontological Bed

Pierre Grenon
Institute for Formal Ontology and Medical Information Science
University of Leipzig
Haertelstrasse 16-18, D-04107, Leipzig, Germany
pgrenon@ifomis.uni-leipzig.de

Abstract

Formalisms and axiomatic theories are designed to support reasoning, they are often intended with a preferred interpretation and a targeted ontology. Questions of proper interpretations and of the possible challenge of an intended interpretation arise when integrating a particular theory in pre-existing formal and ontological settings. This paper reports on an instance of this general problem of ontological engineering. The case study is that of the integration of the Region Connection Calculus for spatial reasoning in the Cyc knowledge base. We show that given the assumptions on the Cyc ontology, RCC had to be interpreted within a substantivalist metaphysic of space as a Boolean algebra of spatial regions which are distinct from their occupants. The RCC literature suggests such an intended interpretation, and this paper intends to show that this was a necessary condition of integration in Cyc's ontology. This led to the enrichment of the Cyc knowledge base, rather than to a radical modification of the upper-level ontology.

1 Introduction

Our standpoint is that of ontological engineers attempting to integrate a comparatively small and allegedly domain specific theory within a larger, multipurpose formal framework. Bluntly put, our ontological interest is concerned with types of entities admitted by the respective theories (and their domains of variables), the intended interpretations of these theories, and issues of relative coherence and soundness. We consider that ontological engineering in this sense is distinct from a work of implementation. Indeed, the power of information science ontologies lies in their semantic aspects and their ontological foundations. These elements will guide, but ought to remain independent from, a work of implementation within a peculiar computing system. In this context, it clearly appears that in order to use an existing axiomatic theory, it is important to know which are its possible interpretations and, especially, if one is the intended one. It is indeed crucial to evaluate how the theory has to be interpreted in order for it to fit within an existing ontology. Two main questions arise: (i) does the formal theory support an interpretation compatible with that intended for the embedding framework? (ii) is it possible to bend the theory and its intended interpretation such as to allow maximal integration? This may lead either to non-trivial changes in the embedding ontology or to the choice of a peculiar interpretation for the theory to be integrated. In turn, it becomes necessary to check whether the chosen interpretation is really allowed by the integrated theory (given its axioms) and that it does not generate ontological monsters, so to speak, according to the standards of the larger ontology. When starting a work of integration, one tends to be maximally conservative from the standpoint of the embedding framework (this is, at least, our experience), for the obvious reason that, *a priori,* it might be simpler and more efficient to adapt a specific theory to an existing larger ontology rather than the converse. Changes in the embedding ontology are always perilous and their consequences are not always easily anticipated. Moreover, from the standpoint of an allegedly total or universal ontology, it is very tempting when integrating presumably specific theories to generalize their relevance by extending their domain of interpretation, and even relaxing not only the intended interpretation, but also the axiomatization. The present paper can be taken as a reminder to be cautious when making a methodological maxim of this tendency.

The axiomatic theory of interest is the so-called Region Connection Calculus (RCC as put forward in [Randell *et al.,* 1992]), a theory prototypically applied to spatial reasoning. The embedding ontology in which we introduced RCC is that of the Cyc knowledge base (Cyc KB), as described by [Lenat and Guha, 1990], and which stands as a larger theory of reality. In part 2, we give some elements regarding the background of Cyc's spatial ontology. In part 3, we present RCC's underlying metaphysics as informally rendered in the most significant sources. Part 4 exposes how we were led to question RCC's ontological significance in our work of integration. We present in part 5 the resulting ontological settings and some extensions of the basic theory that we were able to introduce in a Cyc framework.

2 Cyc's original spatial ontology

The Cyc KB is intended to serve as an encyclopedic repository of all human knowledge, primarily common-sense knowledge. It purports to provide a medium for the representation of facts and the inscription of rules about all existing and imaginable things. Cyc's knowledge is represented in a language called CycL, described in [Cycorp, 2002], which here can be seen as a first order language. The Cyc KB is structured and compartmented into theories (called microtheories), additional theories can be introduced *a volo* in order to account for the specificity of a given domain or context. Since our objective was to produce an integration of RCC supporting spatial reasoning in the Cyc system, we will focus on Cyc spatial ontology.

The relevant fragment of the KB is nonetheless extremely broad. Cyc's declared goal of achieving common-sense reasoning has, as an objective, carved up the ontology in a metaphysically determinate way. In particular, "spatial reasoning" has been understood primarily as a type of reasoning about objects in space, and not primarily about space itself. Roughly, this view can be outlined by the position that in the world there are objects only, in particular physical ones, delimitating their 'place'. In this context, spatial reasoning as such refers to the relations that exist between such objects, say, in virtue of being in such and such relative positioning. We shall refer hereafter to this understanding as the 'relational' approach. Cyc's spatial ontology was primarily relational in this sense. Its ontology's top spatial category, #$SpatialThing, is defined in [Cycorp, 1997] as that of entities "with a spatial extent or location relative to some other [instance of #$SpatialThing]". A large part of the hierarchy under #$SpatialThing was covered by the category so-called #$PartiallyTangible which is, in a word, that of objects with a material constituent, concrete things. In addition, #$SpatialThing also includes certain events. But many specializations were geometrical or otherwise abstract in some sense (so-called #$Intangible). Meanwhile, Cyc had already a large body of spatial and topological relations, most of them over #$PartiallyTangible. These were typically documented as relations holding between objects in virtue of the relations between these objects' spatial extents. Indeed, many of the spatial entities in Cyc could be handled as spatial extents, metonymically, but were not properly speaking identical with their extents. This is notably the case with geopolitical entities (on the other hand, Cyc has a microtheory in which geopolitical entities are geographical regions, not merely located at such a region), but also with persons (which are examples of composite entities, with a tangible part, a body, and an intangible part), or even books (composed of a material support and an informational content). However, spatial extents (intuitively, spatial regions) were not 'reified' nor mentioned in assertions. Their ontological status was, so to speak, implicit and mostly intuitive. On the other hand, still in [Cycorp, 1997], Cyc had a notion of (pure) spatial region used essentially to define contextually empty space. But then, for an entity, its being an instance of this collection was overtly and essentially context dependent. The documentation discusses as an example a piece of atmosphere which can be seen as empty space in a particular context, not in another. A substantivalist view of space considers space and regions of space as substances independent for their existence of what occupies them. Under this view, it is not contextual for an entity to be a region of space, it is essential. In other words, regions of space as portions of an existentially independent (spatial) substance did not exist in Cyc. We will not delve further into the Cyc ontology. For our purpose, it suffices to observe that attempting to integrate RCC in Cyc meant starting to look at RCC from an essentially relational spatial ontology.

3 RCC and its canonical interpretation

Was this approach allowed by RCC itself? RCC is technically referred to as a mereotopology, that is, a theory with both a mereological fragment (a theory of part and whole) and a topological fragment articulated together. The basic part of the theory, so-called RCC-8, has only one topological primitive: Connects-With (or C). We will assume familiarity with the basics of RCC which can be gathered in [Randell *et al.,* 1992] and, informally, in the introduction of [Cohn *et ai,* 1997]. The calculus can be assimilated to a Boolean algebra, which, as we shall see, is not without ontological consequences. Our interest here is mostly in the ontological significance of the theory. RCC is said to support both a spatial and a temporal interpretation. Here, only a putative spatial interpretation is of interest, so regions will be spatial regions. *RCC* is, with respect to spatial reasoning, a theory of regions of space and of their relations, not a theory of physical objects nor of their relations to regions. It is already an other form of spatial reasoning than the one we encountered in Cyc. Characteristically, regions are considered as extended (RCC is originally introduced in analogy to Allen's interval logic). In purely model-theoretic terms, intended models are regular open sets of R^n. For common-sense reasoning, n is equal to either 1, 2 or 3. Going further in the specification of regions would be mostly irrelevant, as specializations derive straightforwardly from the theory. 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 arc 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.

It is useful to remark that the relations in RCC are required to hold between regions of equal dimensionality. This indicates the possibility of treating uni-dimensional and bi-dimensional regions of space, even if the respective domains are independent. This means an acceptance of objects of lower dimensionality than that of physical objects (which are uncontroversially tri-dimensional). Thus, if RCC is a fragment of a theory of reality, its underlying ontology probably is not restricted to physical objects and assimilated things, or, in Cyc's terms, #$PartiallyTangible. However, so far it remains open whether one can conceive in a RCC-like framework objects of lower dimensions that are not spatial regions or, alternatively, only regions of lower dimensions at-which material objects are somehow located. The rejection of boundaries from RCC after [Randell and Cohn, 1989] is yet another evidence that the theory is about relations among regions of a given non-null dimension.

More generally, in RCC, the distinction between regions and their putative occupants is almost systematic in the informal litterature, especially in [Randell and Cohn, 1989; Randell *et al,* 1992; Cohn *et aL,* 97]. These sources suggest the acceptance of a substantivalist view of space; that is, the view that space consists in a substance in its own rights distinct from the objects that occupy portions of it. Formally, the issue is scarcely addressed before [Bennett, 2001] which represents a dramatic broadening of the scope of application of the theory (in addition to the stronger geometrical approach it presents). In multi-sortal axiomatizations, the domain is always restricted to the sort REGION (and secondarily to the sort NULL), a sort PhysObj for occupants of regions is mentioned only to limit the scope of the theory. The first order axiomatizations are seen as fragments of more complete theories. It is not clear how much this is, or is not, an artifact of the sortal approach which is itself an artefact of a putative implementation. Presumably, however, first-order axioms should incorporate quasi-sortal restrictions implicit in the multi-sortal settings.

4 Path to Integration

Cyc's general maxim for ontological integration is generalization. On the one hand, the foregoing considerations on the use of sorts in RCC provide a starting point for expanding the RCC schema by generalizing over sortal restrictions. Overlooking any sortal restriction, on the other hand, is tantamount to claiming generalization of RCC to putatively all entities, namely, #$Thing in Cyc. This is clearly not desirable: the scope has to be at least restricted to spatial things. For Cyc, objects in space (and physical events) are primitives. If RCC is generalizable to #$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. After all, there seems to be no *prima facie* objection to this in the RCC literature. Indeed, the intended interpretation for RCC is as a calculus of 'spatial regions'. But the sources are all starting from the introduction of regions as primitives. There seems to be no claim in the other sense, namely that RCC means, specifically and necessarily, spatial regions, and not, more generally, spatial things in the sense of the Cyc KB.

Let us assume that generalization is possible. C is then to be interpreted as a relation between objects such that these spatially connect along the RCC lines ([Cohn and Varzi, 1996] situates RCC's C among a variety of interpretations). Modulo the issues with openness of regions in RCC, Cyc had a relation of touching among #$PartiallyTangible which intuitively was close enough to C to suggest the possibility of a sound generalization at the level of #$SpatialThing. Most RCC-8 relations (the basic set of binary relations in RCC) had proxies in Cyc (with exception of inverses for implementation reasons). More precisely, there were relations at the level of #$PartiallyTangible which were holding true in virtue of their relata's putative spatial extents in RCC-like relations themselves. In particular, Cyc already had a relation of co-spatiality, a relation holding true of things with identical spatial extents. Meanwhile, [Randell *et al,* 1992] defines = as identity between regions. However, in Cyc, co-spatiality was not identity. There were conceivable cases of entities with the same (spatial or geographical) extent and which were not identical as alluded to in Part 2. The question of whether a generalization of RCC's relations was possible could be rephrased as whether = in RCC was indeed identity or could be interpreted as a co-spatiality relation on #$SpatialThing. If the former, we would have applied RCC straightforwardly to the Cyc ontology. We will see that the status of = in the RCC literature is something of a puzzle.

In [Randell and Cohn, 1989; Randell *et al,* 92], we find the following axiom:

$$(1)\ x = y \equiv (P(x,y) \wedge P(y,x)),$$

where P is "Part-Of. That is, x is identical to y if and only if x is a part of y and y is a part of x. (1) simply states the anti-symmetry of P, i.e., if two objects of the domain are part of the other, they are actually one and the same. Observe that given the definition of P:

$$(2)\ P(x,y) \equiv \forall z\ (C(z,x) \rightarrow C(z,y)),$$

that is, x is a part of y if and only if everything that connects with x connects with y, we can infer the extensionality of C:

$$(3)\ x = y \leftrightarrow \forall z\ (C(z,x) \leftrightarrow C(z,y)).$$

There is however a progressive terminological and ideological shift in the RCC literature concerning the status of =. In [Gotts, 1994], 'EQ' is the privileged notation for the relation of identity and '=' is only mentioned as an alternative notation. With [Bennett, 1995], the matter becomes allegedly more than a notational issue. [Bennett, 1995] argues in section 3 that an axiom of extensionality for C (and thus the anti-symmetry of P) is optional in RCC. The choice depends on the way we construe (1), either as an axiom or as a definition. A similar claim is made with the note 2 of [Bennett, 1996]:

> "[Randell *et al,* 1992] define[s] the relation EQ(x,y) as equivalent to $P(x,y) \wedge P(y,x)$. If this definition is regarded as an axiom rather than a definition and if EQ is treated as logical equality then this axiom is equivalent to Cext."

Cext is the right to left implication in (3), the use of 'EQ' for '=' is [Bennett, 1995]'s import. Using '=' is just a convenient abuse of notation, this symbol is not really that for equality. In order to follow the rationale, let us adopt momentarily this shift of notation, and reserve '=' as a symbol for equality. The thesis is radically defended by [Stell, 2000] that because EQ is allegedly defined and not simply introduced as a logical primitive, EQ is not equality. EQ would simply be an equivalence relation. [Stell, 2000] mentions the note 11 in [Cohn *et* al., 1997]:

> "For notational convenience we will sometimes write x = y rather than EQ(x, y); technically the latter is preferable, since EQ is a relation defined in terms of C rather than true logical equality."

[Stell, 2000] claims that this note provides evidence supporting intuitively his claim. In our opinion, this is a misdiagnosis and the note warrants a better interpretation. Namely, it is an indication that the introduction of EQ (originally, =) by way of definition in RCC reflects a desire of reducing the number of primitive relations in the theory. The equality relation is defined via P, P is defined via C; thus far, there is only one primitive in the theory, C, instead of two, C and = Interestingly, although not addressing this point of detail, [Smith, 1996] stigmatizes this general tendency in RCC as inherited from the classical schools of mereology.

There is a more serious rationale than such footnotes exegesis for not endorsing the [Bennett, 1995; Stell, 2000] alternative. It seems now that the contentious part of the argument rests on the claim that EQ is capable of taking different interpretations according to the way (1) is conceived. Let us assume that *EQ* is not = (logical identity) and rewrite (1), (2) and (3) with EQ in place of = (call the resulting propositions (T), (2') and (3')). Let us see the possible consequences for the ontology. This choice apparently enables the generalization of the intended domain of interpretation from the domain of pure space to that of objects in space. EQ is merely an equivalence relation, "EQ(x,y)" means that x and y have the same region, the same spatial extent. This corresponds (again, abstracting away fine details of the topology) to the co-spatiality relation in Cyc. Note how accommodating this interpretation might be. If we can still include in our ontology regions of space and preserve the ability to interpret on this restricted domain EQ as =, EQ gives the identity criterion for space regions (= entails EQ and, on regions, EQ entails =). So, it seems that choosing a weak interpretation is rather positive and optimal for the resulting ontology.

But there is a drawback. If EQ is but an equivalence relation over a broader domain, then P, which is now a relation on objects (such that the extent of one is part of the other's extent), is no longer anti-symmetric. This means that RCC no longer warrants that two objects with the same (spatial) extents are one and the same thing. In particular, there could be several distinct entities which are all the sums of a given set of objects, say three bricks. Without (1) and only (1), the sum of brick 1 and of the sum of brick 2 and 3 is distinguishable from the sum of brick 3 and of the sum of brick 1 and 2. What is crucial in deciding of the possibility of generalizing the full RCC calculus (RCC-8 and additional Boolean operators) is the behavior of operators in the absence of (1) or equivalently the extensionality of C. The operator for sum, for instance, is defined by [Randell *et al,* 1992] by:

(4) $\mathrm{sum}(x,y) =_{def} \iota w\ (\forall z\ (C(z,w) \leftrightarrow (C(z,x) \vee C(z,y))))$.

RCCs literature generates a new puzzle for interpreting (4). [Randell and Cohn, 1989; Randell *et al* 1992] suggest that while (4) uses the iota operator of definite description, its method of elimination is not necessarily, although possibly, to follow Russell's. Rather, it is intended to be eliminated such as to produce the following assertion:

(5) $\forall x \forall y \forall z\ (C(z,\mathrm{sum}(x,y)) \leftrightarrow (C(z,x) \vee C(z,y)))$.

This has the effect of virtually allowing plural reference, if EQ is to be interpreted otherwise than as equality (since without (1), (1) is not sufficient to rule out that there may be more than one entity substitutable to the sum of two entities in (5) and thus more than one candidate for being such a sum). The trouble with operators which may lose their intended functionality, goes seemingly unnoticed. Moreover, (5) does not even secure the existence of a value. By suggesting this transcription, the authors merely seek to ensure that the Boolean operators are totally defined on their domain. For this purpose, providing for a conventional value in cases where the operators do not properly denote, the introduction of the NULL sort, or, in other words, a notion of null region, is required. It is clear, even for these authors, that this move is not ontologically defendable (see also page 13 of [Simons, 1987] for its vehement dismissal). Rather, this is an algebraic desiderata and an implementational requirement (in short, codes have to return a value). In the context of a putative generalization, this is even more puzzling. What could possibly be the mereological product of two people? A null person (or null person's part)? Such a question leads to nonsense from Cyc's common-sense standpoint. Cyc's uneasiness with the ontological prodigality of RCCs Boolean operators goes further. If Boolean operators denote, they denote objects in space. The sum of this cup and the spill of coffee on this desk is constructible in the theory as a spatial thing. Both the cup and the puddle are in space, that is consensual. The question is whether it is desirable (and needed) to consider that there are sums of such arbitrarily united objects. Ontologically speaking, there were no incentives to go after a notion of null object and totally defined operators, although even in Cyc, the implementation would eventually require an apparatus mimicking the function of the null region.

So, conversely, what was the meaning of a Russellian elimination of definite description? This would result in the following transcription of (4):

(6) $\exists w\ (\forall z\ (C(z,w) \leftrightarrow (C(z,x) \vee C(z,y)))$
$\wedge\ \forall w'\ (\forall z\ (C(z,w') \leftrightarrow (C(z,x) \vee C(z,y))) \rightarrow w = w'\ ))$.

Observe that (5) is indeed entailed by (6). But there is something new in (6). Namely, = is indubitably logical identity and we verify that (6) entails the uniqueness of sum. If EQ was not =, there could be two objects, a and b, in the domain such that EQ(a,b) but not a=b. However, by (3'), these two objects would connect the same objects. Since trivially any object is a sum, any object in the EQ relation with it connects its parts (by (3') and the first conjunct of (6)). But then a and b are identical (by the second conjunct of (6)). EQ is therefore -. This, in turn, entails the extensionality of C.

So, if the iota operator is eliminated a la Russell, C is indubitably extensional, P is anti-symmetric, and EQ is =. We took this to be a reductio ad absurdum of the [Bennett, 1995; Stell, 2000] thesis, and, in fine, as an indication that Cyc's co-spatiality was not an RCC relation. In other words, we had to interpret RCC as a calculus of regions of a substantival space at which entities of various kinds could be located, and not as a theory of these entities. The alternative would have been either going forward with a non extensional mereo(topo)logy of objects-in-space departing significantly from both RCC and Cyc or going back precisely to the more restricted domain of space regions.

5 Extending Cyc's spatial ontology

We thus chose to go for an interpretation according to which RCC was a theory of regions of substantival space. We therefore developed an ontology of space as an entity with an independent existence, choosing to rework and have #$SpatialRegion for domain of RCC relations. In order to bridge the gap between spatial things and spatial regions, we developed a locative apparatus inspired largely by [Casati and Varzi, 1996]. This enabled us to formalize the meaning of object level spatial categories and relations, and make explicit their metonymical character: spatial relations among objects were holding in virtue of similar relations between these objects' spatial extents. Originally, only spatial objects were in the ontology. We introduced (pure) spatial regions, RCC relations among some of these regions, and relations of location between spatial objects and regions. We used these in order to explicitly define spatial relations. This resulted in a spatial reasoning language with greater expressive power, clearer semantics, and an enriched ontology. It also facilitated spatial reasoning with access to a greater variety of data. We could use spatial regions for instance as handles when manipulating datasets for GIS based reasoning. There is unfortunately not enough room here to give specifics.

In a dedicated microtheory, we put some basic constraints on our calculus following RCC's intended interpretation, these motivated in particular further categorial distinctions. To begin with, RCC is region based: relationships are understood as holding among extended regions, i.e., neither points or scattered sets of points are in the domain. We interpreted the regularity of the open sets in the alleged models of RCC as founding yet another category of spatial regions, namely regions homogeneous in dimensionality. We wanted to apply RCC to regions of the three basic spatial dimensionalities (namely, one, two and three dimensional regions), consistently with Cyc's common-sense reasoning leanings. We chose to introduce dimensionality as primitive notions (at the topological level) and as instances of #$Attribute Value (properties of entities as abstract particulars). This afforded us to define three specializations of extended spatial regions which were homogeneous in dimensionality, i.e., regular in the intended sense. We were only requiring our relations to take as arguments homogeneous extended regions and were axiomatically enforcing constraints on co-dimensionality of the relata. So, in the end, we had a single set of RCC relations and operators which were applicable within each dimension independently of other dimensions.

Anecdotally, we carried on in a specific microtheory an atomic extension of RCC along the lines suggested in [Randell *et al,* 1992] by imposing categorial restrictions on the relevant axioms. Despite RCC being a region based calculus, a larger ontology such as Cyc could make use of point-based reasoning and zero dimensional elements. We developed a variant of RCC in a specific microtheory introducing the category of #$SpatialPoint as a primitive and a relation of incidence of points in a region as suggested by [Randell *et al.,* 1992]. We refined this elementary treatment with more topologically oriented primitives of interior and boundary point incidence. Additionally, we developed a set of trans-dimensional relations among entities of other dimensions. If material objects are always located at a three dimensional region, they can be geographically located at a given parallel for instance. Moreover, their footprints (geographical projection) are two-dimensional extents with similar locational properties. So, on the one hand, we allowed partial location of entities at regions of lower dimensionality, possibly at points. On the other hand, we designed projective locations, typically intended to account for geographical positioning. More generally, this allowed approaching regional boundaries in terms of incidence.

6 Concluding Remarks

There exist alternatives in the metaphysics of space. The main distinction being between substantival and relational space. This paper intended to provide evidence that the RCC contributors are 'right' in putting forward certain ontological basics as a frame for using their theory, namely a substantival view of space. The reason why it could not be relational is that the distinction between co-spatial and identical objects would be lost, while co-spatiality is not an identity criterion of objects in space. Once the ontological independence of spatial regions is endorsed, however, two solutions exist. On the one hand, objects are themselves independent from the regions they occupy or they are not. If they are, as in Cyc, RCC has to be interpreted as we put forward here. However, there might be a way in which we could mimic if not secure a relational formalism. RCC predicates could be predicated of objects and events in space, nonetheless, the denotations of operators such as sum, for instance, provided they exist, would be spatial extents. In that context, the sum of two entities in the RCC sense would be the sum of their extents. As the extensionality of C would be preserved by this functionality of sum, this would be identical to the extent of any of their putative sums (in a purely mereological sense). One can then endorse the thesis that such extents are parts of some kind (namely, spatial) of entities in space. Co-spatiality becomes a form of overlapping, it is identity of spatial parts. However, this solution requires a broad and strong mereological treatment at the upper-level of the ontology which goes far beyond spatiality. Such a solution which was not affordable in the terms of Cyc's ontology suggests the possibility of alternative interpretations for RCC provided it is to be embedded in an ontology with a more general and primitive mereological framework.

Acknowledgements

The redaction of this paper was supported by the Alexander von Humboldt Foundation under the auspices of its Wolfgang Paul Program. This paper stems from research done while working at Cycorp. Such work, in particular that alluded to in section 5, could never have been accomplished without the participation of James Masters, many thanks go to him for the many discussions and debates we had there. Additional thanks go to James Zaiss, Pierluigi Miraglia, Michael Pool, Fritz Lehmann and William Jarrold for helpful comments and discussions during this period. Thanks go to Brandon Bennett for discussions during his visit at IFOM1S, and to finish with to the anonymous reviewers of IJCAI for their insightful comments. They are not responsible for any errors which remain.

References

[Bennett, 1995] Brandon Bennett. Carving up Space: existential axioms for logical theories of spatial regions. In *Proceedings of the 1995 International Joint Conference on Artificial Intelligence workshop on Spatial and Temporal Reasoning,* Montreal, Canada, 1995.

[Bennett, 1996] Brandon Bennett. Carving up Space: steps towards construction of an absolutely complete theory of spatial regions. In *Proceedings ofJELIA '96,* pages 337-353, 1996.

[Bennett, 2001] Brandon Bennett. Space, Time, Matter and Things, *Proceedings of the Second International Conference on Formal Ontology in Information Systems (FOIS'01),* pages 105-116, 2001.

[Casati and Varzi, 1996] Roberto Casati and Achille C. Varzi. The structure of Spatial Localization. *Philosophical Studies,* 82, pages 205-239.

[Cohn and Varzi, 1996] Anthony G. Cohn and Achille C. Varzi. Modes of Connection. In *Spatial Information Theory - Cognitive and Computational Foundations of Geographic Information Science* (Lecture Notes in Computer Science 1661), pages 299-314, Springer-Vcrlag, 1999.

[Cohn *et ah,* 1997] Anthony Cohn, Brandon Bennett, J. Gooday, and N. M. Gotts. RCC: a calculus for Region based Qualitative Spatial Reasoning. *GeoInformatica,* 1:275-316, 1997.

[Cycorp, 1997] http://www.cyc.com/cyc-2-l/toc.html, Cycorp, Inc.

[Cycorp, 2002] http://www.cyc.com/cycl.html, Cycorp, Inc.

[Gotts, 1994] N. M. Gotts. How far Can We 'C? Defining a 'Doughnut' Using Connection Alone. In *Proceedings of the Fourth International Conference on Knowledge Representation,* Morgan Kaufmann, pages 246-257, 1994.

[Lenat and Guha, 1990] Douglas B. Lenat and R. V. Guha. *Building Large Knowledge Based Systems.* Reading, Massachusetts: Addison Wesley, 1990.

[Randell and Cohn, 1989] David A. Randell and Anthony G. Cohn. Modelling Topological and Metrical Properties in Physical Processes. *Proceedings of the First International Conference on the Principles of Knowledge Representation and Reasoning,* pages 55-66, Morgan Kaufmann, Los Altos, 1989.

[Randell *et al,* 1992] David A. Randell, Zhan Cui, and Anthony G. Cohn. A Spatial Logic Based on Regions and Connection. *Proceedings of the Third International Conference on Knowledge Representation and Reasoning,* pages 165-176, Morgan Kaufrnann, San Mateo, 1992.

[Simons, 1987] Peter Simons. *Parts: a study in ontology.* Oxford University Press, 1987.

[Smith, 1996] Barry Smith. Mereotopology: A theory of parts and boundaries. In *Data & Knowledge Engineering,* 20:287-303, 1996.

[Stell, 2000] John G. Stell. Boolean Connection Algebras: A New Approach to the Region-Connection Calculus. *Artificial Intelligence,* 122:11-136, 2000.