Modelling Models of Robot Navigation Using Formal Spatial Ontology*

John Bateman and Scott Farrar

University of Bremen,
FB10, Sprach- und Literaturwissenschaften,
Bibliothekstr. 1, Bremen 28359, Germany
{`bateman, farrar`}`@uni-bremen.de`
`http://www.sfbtr8.uni-bremen.de/i1`

Abstract. In this paper we apply a formal ontological framework in order to deconstruct two prominent approaches to navigation from cognitive robotics, the Spatial Semantic Hierarchy of Kuipers and the Route Graph of Krieg-Brückner, Werner and others. The ontological framework is based on our current work on ontology specification, where we are investigating Masolo *et al.*'s Descriptive Ontology for Linguistic and Cognitive Engineering (DOLCE) extended particularly for space and navigation by incorporating aspects of Smith *et al.*'s Basic Formal Ontology (BFO). Our conclusion is that ontology should necessarily play an important role in the design and modelling of cognitive robotic systems: comparability between approaches is improved, modelling gaps and weaknesses are highlighted, re-use of existing formalisations is facilitated, and extensions for interaction with other components, such as natural language systems, are directly supported.

1 Introduction

The use of formal ontology in the field of cognitive robotics has until recently been quite limited. We argue in this paper, however, that the sophistication required of current cognitive models, the functionalities required of cognitive robots, and the state of the art in formal and computational ontology all combine to suggest that a closer interaction between ontology and cognitive robotic modelling is now appropriate. The explicit adoption of computational ontology brings a stronger set of modelling constraints to bear on the necessary issues, and also provides a much richer set of re-usable building blocks for modelling.

In general one can envision at least three scenarios for incorporating ontologies into cognitive robotics. First, ontology can be used to enhance the design

---

* The work reported in this paper was supported by the Deutsche Forschungsgemeinschaft (DFG) as project I1-[OntoSpace] of the Transregional Collaborative Research Center SFB/TR 8 Spatial Cognition. The authors also thank particularly the anonymous reviewers for their extensive corrections and comments, as well as the other members of the SFB/TR 8.

C. Freksa et al. (Eds.): Spatial Cognition IV, LNAI 3343, pp. 366–389, 2005.

of robot knowledge representations; that is, ontology can help to clarify the relations among various representational levels and to provide a semantically coherent account of the entities used in symbolic reasoning. Second, ontology can be used to develop and constrain a sharable conceptualization of the environment in the interaction of intelligent agents, including both robot-robot or human-robot interaction. Third, ontology can contribute to solutions for the problem of partial information at the sensory-symbol interface; that is, partial sensory input can be augmented with knowledge from an ontology to build a more accurate symbolic representation. For example, the laws of mereotopology (see below) can be leveraged to verify the existence of necessary parts from poor or incomplete sensory data.

In this paper we focus primarily on the first and second scenarios by investigating the use of spatial ontologies in the modelling of navigational capabilities for cognitive robotics. We do this concretely with respect to two navigational models currently being developed and used in cognitive robotics: the Spatial Semantic Hierarchy of Kuipers [1] and the Route Graph of Werner, Krieg-Brückner and others [2, 3]. These were chosen on the one hand due to their importance for robot navigation and, on the other, because they have already made considerable moves towards compatibility with ontology-based design. We will show how such models can be placed beneficially against a broader ontological background, adopting for this purpose the Descriptive Ontology for Linguistic and Cognitive Engineering (DOLCE) [4] together with some proposals that we are currently developing for extensions in the area of spatial ontology.

We structure the paper as follows. In Section 2, we set out the two selected models of robot navigation. In Section 3, we briefly introduce the account of ontology that we wish to draw upon, concentrating specifically on issues of space. In Section 4, we show how the two individual navigation models can be inter-related, placing them against the ontological background of our framework. Finally, in Section 5, we conclude with an explicit discussion of the benefit of incorporating ontology in the modelling of cognitive systems.

2 Two Models of Robot Navigation

Despite the limited application of ontologies within cognitive robots and navigation, there are now approaches that explicitly draw on ontology in their formalization. Both the Spatial Semantic Hierarchy and the Route Graph are of this kind. Our discussion of each will follow a similar pattern. We first identify the major ontological domains adopted or assumed by each model and then set out briefly the place and nature of relationships *between* those domains that the model specifies.

2.1 The Spatial Semantic Hierarchy

Kuipers' Spatial Semantic Hierarchy (SSH: [5, 6, 1]) is an approach to robot navigation which decomposes a robot's knowledge of its environment across several distinct layers in an *ontology hierarchy.* These layers allow distinct kinds of representations to co-describe the robot's experience and plans. This co-description serves to abstract an agent's spatial knowledge away from the details of its environment and its sensorimotor apparatus [6, p2]. Spatial knowledge can then be used that is derived, or abducted, from the sensory level for purposes of navigation, rather than simply having the agent rely on a reactive/sensory level. This approach finds its motivation in cognitive robotics and draws on research from cognitive science and psychology in an attempt to solve problems via high-level reasoning.

The SSH layers involve the following levels of abstraction; the symbolic levels are given in bold-face type: the sensory level, control level, **causal level**, **topological level**, and **metrical level**. Each level has, as Kuipers describes it, "its own *ontology* (the set of objects and relations it uses for describing the world) and its own set of inference and problem-solving methods" [6, p2].

The first two levels, *sensory* and *control*, concern the continuous output of sensors such as vision, laser or sonar range-sensing. This numerical data is represented by a sensory input vector $s(t) = [s_1(t), ..., s_n(t)]$ where $s(t)$ is the state of the agent at time t and $s_1(t), ..., s_n(t)$ are individual sensor outputs. These data are then abstracted via a set of *control laws* at the control level to discrete states encompassing position and orientation. This is the primary mechanism by which continuous descriptions and discrete, symbolic descriptions of behavior are related: the continuous numerical entities at the sensory and control levels are abstracted to a discrete symbolic representation for use at the causal and topological levels. We will have little more to say about the sensory and control levels in this paper, however. We will also not address the *metrical* level particularly. This consists simply of a global 2-D geometric map of the environment in a single frame of reference, a so-called "Map in the Head" [1, p195]. Our central concern will be on the central two levels, the causal and the topological.

The ontology of the *causal level* defines views, actions, events and the causal relations among them. Intuitively, a view is some state of affairs as perceived by the robot at some given moment while actions and events describe the motions that a robot may initiate. Views are thus defined as symbolic abstractions over the sensory input vector obtained at a locally distinctive state [1, p205]. A 'locally distinctive state' is another vector $s = (x, y, \theta)$ indicating the agent's position (x, y) and orientation θ within the environment. Views may be both complete and partial, i.e., attending to all or only to some subset of the available sensory inputs.[1] Changes in view are brought about by actions and the SSH defines two types: turns and travels. Turns leave the agent in the same place, while travels change the agent's location [1, p206]. All actions are caused by the agent applying one or more control laws in some distinctive state. The purpose of this abstraction in terms of actions is to lose "the details of how views are defined or how actions are implemented in particular circumstances" [1, p195]. Events are

---

[1] The sensor values are also considered to be functions of the agent's state [1, p199] but we will not consider this complication further here.

then used to describe the complex of a change in view via an action. They are represented by schemas of the form $\langle V, A, V' \rangle$, where an action A causes view V to change to view V'. Finally, a routine is defined as a set of such schemas indexed by the initial view [1, p207].

Reasoning on entities at the causal level is performed within the SSH with the aid of McCarthy's Situation Calculus [7, 8]. Within the causal layer the fact that some view V holds and that some action A is carried out at the current moment, *now*, is additionally associated with a particular situation, s_o. The state of affairs can then be represented by the Situation Calculus statements: $holds(V, s_o)$ and $do(A, now)$. Accordingly, the state of the world, some situation, is said to change when an action is applied, thus producing a new situation.

The next ontological layer is the *topological* level. This includes the categories of places, paths and regions, with their associated connectivity and containment relations. Places are defined simply as zero-dimensional entities which may lie on a path. A path imposes an order on the places lying on it by virtue of its direction: a path is thus a one-dimensional *subspace* leading from one place to another and having one of two possible directions (dir=±1). A topological, or place, graph can then be constructed as a map of the environment consisting of sets of places and their connecting paths. Paths also serve as the boundaries for regions. A region is defined as a two-dimensional subset of the environment, i.e., a set of places. Path directedness also allows a reference system to be determined. Each directed path divides the world into two regions: one on the right and one on the left. A bounded region is then defined by a directed path with the region on this path's right or on its left.

The SSH also uses regions to define a hierarchical view of space, as now amply motivated psychologically. This allows maps to be pitched at various levels of granularity with sets of places within a map of greater detail being represented within a map of lesser detail by single 'abstraction regions'. There are therefore two *levels of abstraction* within the topological layer, one for place and one for region. A place is still a zero-dimensional point but may also function as an abstraction of a region. To relate these levels Kuipers describes both upward and downward mappings. An *upward mapping* holds when "a place at a lower level is mapped to the place corresponding to the abstraction region that contains it" [1, p212]; a *downward mapping* holds when "a $\langle place, path, dir \rangle$ tuple at the higher level is mapped to a corresponding $\langle place, path, dir \rangle$ tuple at the lower level" [1, p212]. These levels of abstraction are not to be confused with the particular ontological levels defined by the ontology hierarchy. Places and regions occupy two different levels of abstraction within the same ontological level, while the causal, topological, and metric levels are considered ontologically distinct.

We can summarize those relations defined within the topological level by drawing on Kuipers' listing [1, p210] as follows:

- $on(place, path)$: *place* is on *path*;
- $order(path, place_1, place_2, dir)$: the order on *path* from $place_1$ to $place_2$ is *dir*;
- $right\text{-}of(path, dir, region)$: *path*, facing direction *dir*, has *region* on its right;

- $\mathit{left\text{-}of}(path, dir, region)$: $path$, facing direction dir, has $region$ on its left;
- $in(place, region)$: $place$ is in $region$.

In general, transforming from consecutive ontological layers in the SSH is carried out via a process of *abduction*. Thus, the places, paths and regions of the topological level are created by deducing some minimal description that is sufficient to explain the regularities found among the observed views and actions of the causal level [1, p209]. An example of an abduction rule from the causal level to the topological level is given as:

$$\forall view \exists place\ (at(view, place))$$

which means that an association will be established between views and particular locations. The following relations defined by the model combine topological and causal level constructs and so can be considered to be 'inter-ontology' relationships. We will bring out the special nature of these kinds of relationships further below.

- $at(view, place)$: $view$ is seen at $place$;
- $along(view, place, dir)$: $view$ is seen along $path$ in direction dir;

Entities from the control and topological levels may also be mixed in what the SSH terms 'axioms of commonsense'. Consider the gloss of one such axiom: "If the agent travels along a certain directed path, turns right, then travels again to reach a certain place, then that place lies within the region right-of that directed path" [1, p211]. This relates a routine consisting of three events to a region via several places; it also clearly requires several further assumptions to be made concerning the 'shape' of paths and the non-identity of the places mentioned in order to be accurate. Another example of an axiom is the following: $\langle V, (turn\ \alpha), V' \rangle \rightarrow \exists place[at(V, place) \wedge at(V', place)]$, glossed as "A turn action leaves the traveler at the same place" [1, p210]. This combines levels similarly, relating an event to a place. Note that in order to support sophisticated behavior, a considerable number of these kinds of axioms are necessary; providing such characterizations of the 'world' is precisely one of the tasks of ontology.

2.2 The Route Graph

The second area of research we discuss is the Route Graph (RG: [9, 2, 10, 3]). The RG was originally developed for practical robot navigation in real application contexts and so is also faced with the problems of mediating direct sensory input with abstract path control. The essence of the RG was that information concerning different routes can be integrated within a single network-like structure, combining a variety of data sources [2, p297]. Route graphs have accordingly been characterized in a number of different ways and several alternative descriptions of the RG that are closer to actual robot control structures have also been given in the literature [10, 11]. In our discussion here, however, we will draw particularly on the characterization of Krieg-Brückner *et al.* in this volume [3]. This brings a number of advantages for us—in particular that the RG is already defined there in terms of several distinct ontological areas, similar to the approach seen for the SSH. We will assume for present purposes, therefore, that the varied RG descriptions are all broady compatible with this latter ontologically-inspired account.

The starting point for the definition is basic graph theory, yielding nodes and edges. The edges are also directed, and so each has a source and a target. Edges may be combined into a sequence of edges, which may in turn be specialized as a path (possibly containing cycles), which may be specialized further as a route (containing no repeated edges). The Route Graph then refines the generic graph notions in a number of ways suited for concrete robot navigation and motion control involving a real physical robot, with dimensions and sensor capabilities.

First, nodes are refined to places. Places are anywhere that a robot can 'be'. They include in their specification a width (provided by a Voronoi representation of the free space constituting a place: cf. [3, 11]), and an origin that is used for defining the relative (angular) positions, or 'bearings' of all the nodes' incoming and outgoing edges. The origin constitutes a *reference system* for a place [2, p307]; orientation is therefore strictly local. Second, edges are refined to route segments leading from one place to another. Each route segment includes the angular displacements of its edge with respect to both the origins of its source and its target. This means that it is possible to specify, and also physically to rotate, a robot positioned with respect to the origin of the starting node, so that it is ready to follow a given route segment and, after having followed the segment, again to rotate the robot so that it faces in the direction of the origin of the target. Each segment specifies in addition its own length (so that the robotic agent knows how far to travel before expecting to be at its target) and its width (so that the robotic agent can know whether or not it will fit through the segment).

The RG also includes a notion of abstraction similar to that of the SSH. For example, at a particular level of abstraction an entry and an exit ramp of a highway can be considered as two different places (nodes in the RG), whereas, at a higher level of abstraction, the two nodes could be considered one place corresponding to the complex notion of a 'road junction'. Similarly, the complex possibilities for navigation within a train station might quite appropriately, at a higher level of abstraction, be collapsed to a single place: 'the station'. This is modelled formally in terms of an AbstractsTo relation, defined in terms of graph morphisms between RGs of differing granularities.

This abstract definition of a route graph is intended to be neutral across a wide variety of possible route planning tasks. Some intuitive examples of routes discussed include: "CommuterTrainLine, ShipRoute, FootPassage, CityRoad, Highway or Labyrinth" [2, p307]. To accomodate this, the RG incorporates a notion of *layers* where "each Layer represents a Route Graph of a particular Kind" [2, p310]. *Kind* refers to the nature of the places found within the RG. That is, the places in a railway system, e.g., the stations and various transfer points, are of the same kind, but are different from the places in a RG for of-

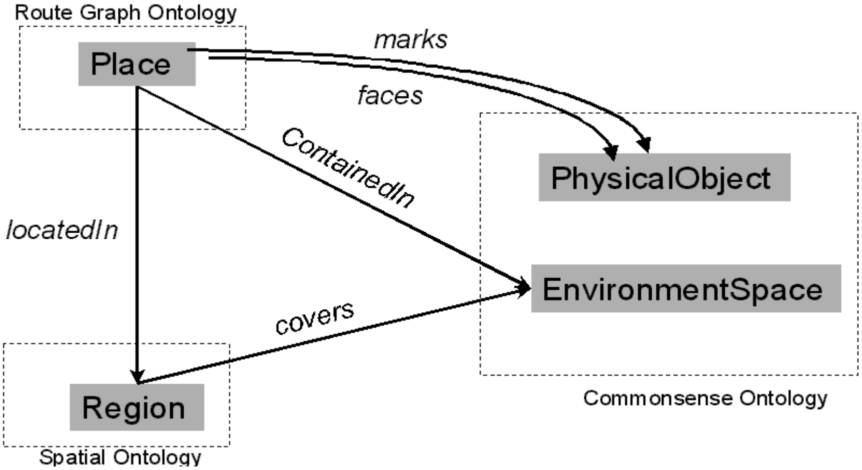

Fig. 1. Relations between RG, spatial, and commonsense ontologies: compiled from Krieg-Brückner *et al.* (2004)

fice navigation, e.g., doorways, corridors, and rooms. When the source and target nodes of a route segment are of the same kind, such RGs are called *homogeneous*. Heterogeneous route graphs do not allow direct transitions between them. Instead, a special type of route segment, whose source and target nodes can be of different kinds is introduced, called a *Transfer* [2, p311]. Krieg-Brückner *et al.* suggest that each of the RG-kind specializations may bring to bear additional kinds of information relevant for that specific application domain. For example, route segments may bring with them certain conditions that must be met before proceeding along the course of the segment.

This latest definition of route graphs also sets out some of the relations that are required between the world of route graphs (consisting of an ontology of places, route segments, paths, etc.) and other ontological domains more clearly than has hitherto been the case. Two such domains are explicitly identified: a spatial ontology, that is expected to provide *spatial regions*, and a 'commonsense ontology', providing everyday objects such as rooms, offices, corridors, and so on. The relationships provided are intended to allow inferences back and forth between places in a route graph, the spatial regions that such places occupy, and the everyday objects that those regions 'cover'. Thus, a place in a route graph can be defined as being locatedIn some given spatial region, and that spatial region might itself cover some physical 'environment space', such as a room; alternatively, the route graph place might be specified directly as being containedIn that environment space. The containedIn relationship is thus a composition of locatedIn and covers.

Two further inter-ontology relations are defined: marks and faces. Both relate a place in a route graph to a commonsense object. For example, a particular place in the route graph can be defined as a *routemark* which marks a commonsense object such as a window. With this in place, a route instruction such as "Go to the window" can be directly related to a route graph description which could drive a robot's motion to the required place. Similarly, a place can be said to face a particular object, so that the robot can be told to "face the elevator" with an appropriate realization of this at the level of the route graph description. A *landmark* is similar but marks points that lie not on a route but in the distance; this is simply modelled by incorporating special kind of edge (a vector) that leads to the landmark in question. Such marks are provided to support localization by triangulation.

We shall argue below that these inter-relationships are in fact crucial for combining the navigation and control-oriented conception of a route graph with further components, such as dedicated spatial reasoning systems or natural language interaction modules. Figure 1 collects together the current set of inter-ontology relationships proposed by Krieg-Brückner *et al.*, summarizing the ontological domains to which they refer.

3 Formal Ontology and Space

In this section we present the necessary background to orient the reader concerning formal ontology, and in particular the place of space within formal ontology. The reader is also referred to Bateman and Farrar [12] for an almost exhaustive account of the state-of-the-art concerning current approaches to space within ontology.

3.1 Formal Ontology: Background

The kind of formal and computational modelling that is our concern here has been described by Nicola Guarino as belonging to the *ontological level* [13]. Such descriptions are intended to provide a 'meaningful' structuring of some domain of concern that goes beyond purely logical adequacy, conforming with the necessary regularities of the world of our experience. As an example of such a description, Guarino considers the following logical expression:

$$\exists x \; (ball(x) \wedge red(x))$$

According to such a representation there is no ontological (or any other) difference between 'ball' and 'red'. Logically they are both unary predicates but, in terms of our experience of the world, this misses several important distinctions. For an adequate knowledge representation, we would rather state that actually a 'ball' is a concept and 'redness' is some property that such concepts can carry—e.g., 'red' is the value of a 'hasColor' relation. This begins to impose far more *structure* on the information being represented than is evident in the initial logical expression alone.

The ontological questions that arise then revolve around just what kind of categories can enter into 'hasColor' relationships, what kind of relation is 'hasColor', and what is the full set of such categories and relationships. An ontology is a way of making explicit those commitments and structural necessities that follow from the fact that we are modelling not knowledge in the abstract, but concrete objects, qualities, relations and events of the known world subject to a rich web of non-arbitrary constraints. As a corollary, an ontology is also a way of specifying explicitly just what follows from particular kinds of modelling decisions: were we to state that ‘ballness’ is a quality that inheres in certain colors, this would be a strong ontological statement and many consequences would follow from it. This means that the choice of logical representation is no longer arbitrary. The ontology therefore establishes a methodology and a set of principles for deciding in what way entities, relations, activities and so on are to be captured in a formal representation. Since the resulting ontologies are motivated by their anchoring in the world, it is generally hoped that representations that respect those ontologies will provide a more robust basis for inter-operability and knowledge sharing.

The specification of an ontology starts with a modelling language, which is used to represent the elements in the intended domain of discourse. Depending on the approach, there may be another language called an *ontology meta-language* used to describe the modelling language itself. The constants of the meta-language are used as predicates in the modelling language. The relationship between an ontology modelling language and a meta-language can be made to do some useful work by setting out clear methodological criteria for how ontology construction may proceed in terms of properties that need to hold at the meta-level. A successful example of this can be found in the *OntoClean* methodology [14, 15] which uses the notion of a meta-language extensively in its definition of *meta-properties.* Meta-properties are properties of properties, not of objects in the world and are used to constrain ontology development and to evaluate particular proposed ontological organizations. The meta-properties particularly important for OntoClean are: *rigidity*, *identity*, *unity*, and *dependence.* Rigidity refers to essential properties, i.e., properties that an entity cannot loose without ceasing to be itself; identity refers to criteria for discriminating entities from each other or for recognizing when one has a particular kind of entity; unity refers to the ‘wholeness’ of an entity, whether it has parts, boundaries and so on; and dependence reflects whether an entity can exist independently or whether it needs to be ‘carried’ by another (e.g., the color of an object is dependent for its existence on that object, the hole of a doughnut is dependent for its existence on that doughnut).

In the construction of formal ontologies, several issues immediately confront the knowledge (or, rather, ontological) engineer. The first is to consider the basic categories and their interrelationships so as to build up an organizational backbone for further specialization. We argue that representing such very basic, foundational features of the world in general, and of spatial objects in particular, is a prerequisite for constructing intelligent spatially-aware agents. Such systems can then operate in terms of situations or settings that are very much more like the kinds of settings that humans take for granted without the need for more *ad hoc* axiomatizations: this is the traditional link that is made to naive physics and modelling situations for intelligent behavior [16]. The foundational ontological properties are 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 be effected.

3.2 Generic Upper Ontologies

We are currently constructing a particular view of ontology which builds upon several state of the art formal ontology specifications. We assume a collection of abstract generic ontological modules that are used for the definition of more specific subontologies. Although the relation between ontological modules can be complex and requires more extensive discussion, we will not foreground this aspect here. We will simply assume for present purposes that relations between ontological modules take the form of structured mappings between the classes and relations of the modules involved. In the simplest case, an ontology submodule may straightforwardly extend a more generic ontology by the subsumption relation.

For reasons we have set out at length elsewhere [17, 12], we select the Descriptive Ontology for Linguistic and Cognitive Engineering (DOLCE) [4] as our main organizing framework. DOLCE was originally part of the WonderWeb project[2] whose aim was the development of foundational ontology libraries for the Semantic Web. DOLCE's upper level, a portion of which relevant for space is shown in Figure 2, provides a generically re-usable high-level characterization of the entities of the world that is particularly appropriate as a basis for further development. It was constructed applying the principles set out in the OntoClean methodology [15, 14] mentioned above and is supported by detailed axiomatization. In the following, we pick out particularly those aspects of DOLCE that are necessary and useful for our consideration of space in the context of robot navigation.

The most fundamental divisions made in DOLCE assert that the world can be divided into four classes of entities: first, there is a fundamental division between entities that unfold in time, called PERDURANTS, and entities which are present 'all-at-once' in time, called ENDURANTS, and second, there are QUALITIES, which inhere in other entities, and ABSTRACT entities. The physical objects generally of most concern to robots are a particular subclass of endurant (physical endurant: PED). Physical endurants are distinguished from non-physical endurants primarily by their relation to space: they are necessarily located spatially.

3.3 Spatial Ontology Within DOLCE

A very basic 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

---

[2] IST Project 2001-33052 WonderWeb: Ontology Infrastructure for the Semantic Web.

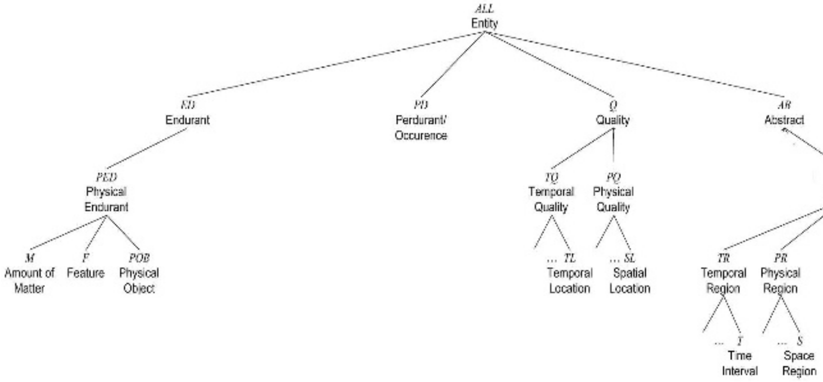

Fig. 2. DOLCE taxonomy relevant for space: extracted from Masolo *et al.* (2003)

objects [18, p2]. The first view is termed the *Newtonian*, *Galilean* or *absolutist* view of space, and the second the *Leibnizian* or *relationist* view. This distinction has important 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 accordingly. In contrast, this need not be the case in a Leibnizian ontology, where space is only present indirectly as relations between objects themselves.

The treatment of space and location in DOLCE is to consider location analogously to other ‘qualities’ that a physical endurant may ‘possess’, such as color or weight. Qualities are bound very closely to their bearing entities: thus, the color of some particular rose is uniquely *that* rose’s color and no other’s. Exactly parallel, a rose’s location is uniquely the location of that rose. The color and the location of the rose cannot be separated from that rose. Comparison between entities in terms of their qualities is only possible by a further step of relating the quality to ABSTRACT regions that give them values. These regions are ‘quality spaces’ in the sense of Gärdenfors’ *conceptual spaces* [19] and are thus not to be confused with the spatial notion of ‘region’ used in formalizations of space. Instances of QUALITY are said to “inhere” in their associated hosts, but their values are defined as points (elements) in a corresponding abstract quality space. This supports an ontologically sound understanding of situations described, for example, as “the color of the rose changed from red to brown” where it is clearly not some color, e.g., ‘red’, that changes but only the intrinsic color that inheres in the rose. That is to say that the position of that intrinsic color with respect to the abstract physical region, or quality space, of color changes. When applied to location, therefore, DOLCE maintains that each physical endurant necessarily has a quality location and this in turn receives a ‘value’ in terms of the SPACE REGION.

More formally, using the DOLCE axiomatization, particular spatial locations (SL: cf. Figure 2) are themselves unstructured and are kinds of PHYSICAL

QUALITY (PQ). These are related by a QUALITY relation ('qt') to entities organized within the SPACE REGION (S), a subcategory of ABSTRACT (AB) PHYSICAL REGIONS (PR). Significantly, the specification of particular possibilities for structuring the SPACE REGION is not refined further. Entities that are spatially present, e.g., physical endurants, are then bound into space by virtue of what DOLCE terms a *spatial mutual specific dependency* relationship (MSD_S). 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 restrict further how that might be specified. The generic quality axiom $MSD_S(PQ, PED)$ of DOLCE then states that physical qualities and physical endurants are mutually dependent and the qualities and their objects will necessarily co-locate.

One further spatially relevant subcategory of physical endurant within DOLCE is FEATURE. 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 (a DOLCE *one-sided generic dependence*: OGD). They are not then distinguishable as 'parts' of an object in the sense that they could be isolated (even potentially) from their wholes and include traditional problem cases such as 'holes' (cf. [20])—as in the holes of donuts—gulfs, openings, boundaries and so on. 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.

We consider 'relevant parts' to be related to spatial qualities proper and will pick out 'places' for separate treatment below.

3.4 Extensions to the DOLCE View of Spatial Ontology

We now focus in on the characterization of space and what precisely it may mean for physical objects (or events) to be *located at* particular locations in space. This raises questions about both how the objects concerned and the locations are to be identified. Our goal is to reach a specification that is adequate for an ontological account of robot navigation models and which will support inter-operability with other components.

First, we consider critically along the preliminary lines that we began setting out in [21] the treatment of locations within DOLCE as spatial qualities. It appears to us beneficial to separate clearly between notions of spatial extent and a broader notion of location. Whereas the former is indisputably an inherent quality of a physical object, any physical object necessarily takes up the space that it does and in the particular morphological form that it does, the latter is more complex. A reasonable specification of location, in the sense of where a physical object is located, requires reference to a broader scheme of possible positions. Only when such a scheme has been specified can we talk of the position of an entity at all. In this sense, then, there appears to be a subtle difference between the quality and qualia of color and the notions of location/position. Whereas it is not necessary to fix on a color space in order to acknowledge the existence of a rose's color: the rose would have the color quality regardless, it *is* necessary to fix on some decomposition of space in order to see that rose as 'positioned' at all. We will term such a spatial decomposition a LOCATION SCHEME and it is only through such location schemes that spatial access to entities is provided. Thus, although analogous to a quality space, a location scheme is more like a quale supporting qualities of an entire physical setting rather than for any particular entity.

Whereas particular colors (i.e., color qualities) that may exist (such as the particular red of this particular rose) depend on their bearing objects for their existence, particular positions depend on their locational scheme. A color quality is given a 'value' by a selected color quale, whereas a location/position *only exists* given a selected location scheme. We take this to be the case for both Newtonian/Galilean and Leibnitzian views of space. Any physical object will always be placed within an entire framework of spatial relationships. These relationships may make reference only to other objects (the Leibnizian view) or may rely on pre-structured space as such (the Newtonian/Gallilean view). The precise characterization of the placement depends on the location scheme adopted. Moreover, the position may exist regardless of any object that happens to occupy that position: the existence of the location scheme is sufficient.

Candidates for location schemes within a foundational ontology then include all of the formal accounts of space, of topology, of regions and so on that have been developed in the qualitative spatial reasoning and representation community [22, 12]. When such an account is also axiomatized in a manner compatible with the axiomatization provided for the rest of the foundational ontology, it can then directly constitute a parameterized ontology module in the sense of a submodule mentioned above. The provision of such ontology modules constitutes one of the ongoing research tasks of the Collaborative Research Center on Spatial Cognition within which our own research is situated [23].

To distinguish locations, positions and location schemes from the quality-quale axiomatization in DOLCE, we propose here to adopt the spatial primitives set out by Grenon and Smith in their Basic Formal Ontology (BFO: [24]). This also requires an adjustment to the categories found under physical endurants in DOLCE. Grenon and Smith posit both a purely spatial ontological category, SPATIAL REGION, including points, lines, surfaces, volumes, etc. as necessary (and is therefore Newtonian and absolutist), and a class of endurants that are intrinsically spatial, called SITES. Sites are just like other physical endurants as far as necessarily having a location in space, but also function themselves further as 'spaces' within which other physical objects (BFO: substances, DOLCE: physical endurants) can be. Grenon and Smith list holes, cavities and both real and fiat enclosed spaces as possible sites drawing on the formalization of such entities worked out by Casati and Varzi [20].

Sites are important for robot navigation because most of the objects with which we are concerned there, rooms, corridors and the like, are of this kind rather than simple physical objects. It should be clear that to talk of a 'room' ontologically as a physical object in the same way as a 'chair or a 'table' would be problematic. A room is a 'hole' of a particular kind that is defined by the physical object of its surrounding walls. In our conceptualizations of rooms, however, it is not the walls that are usually prominent—even though it is probably the walls that are more prominent to a robot's sensors. Bridging between these levels of reality is one of the tasks taken up by the modelling we have seen in the SSH and the RG, but it is also fundamentally involved in the specification of an adequate formal ontology. The most recent formalization of the RG then talks quite properly of a relationship between places and ENVIRONMENT SPACES: this maps quite directly to the notion of site as employed here.

The ontological relationships for location posited within the BFO are:

- SpatialLocation (or located at in Smith and Varzi [25]): a functional notion that associates an entity with a unique spatial region (a part of pure space);
- OccupiedBy: a relationship between a site and (one of) its occupants that ensures that the associated spatial regions align properly;
- place-of: a further functional notion associating occupants with their minimal places (functionally identified sites).

Although there are several further subtleties that also belong in a full account, we simply adopt these for present purposes without further discussion.

The notion of pure space we assume is taken to come with the structuring of a location scheme as suggested above. In addition to standard qualitative treatments of space, we also adopt here the notions of variable granularity and *qualitative coordinates* introduced by Bittner and Smith [26, 27]. This gives us the power necessary to talk of varying levels of abstraction such as a 'room' being a point in a network of routes, which can be used to position other rooms, or itself as a place in which movement can be pursued. The 'hole' or 'cavity' that is the room has itself a (parasitic) spatial extent and so can of necessity also serve as a component of a scheme for structuring location that can be utilized for qualitative coordinates. All such objects are then further subject to granularity selections in that a description can pick out positions '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' for color and 'in', 'in the corner', etc. for position). The particular location descriptions are similarly drawn from the make-up of a space region as defined by its location scheme. We can also at this point align the DOLCE dependent category of PLACES (a type of FEATURE: cf. Figure 2) with the BFO notion of site. As Grenon and Smith [24] write:

> "A room in a house is a site, as also is a landing strip, a meadow, the interior of your car or of your airline cockpit. The corner of a room is a site, as also is your alimentary tract or the interior of an oil pipeline."

We discuss this further elsewhere [21], where we also offer an illustrative informal example of the combination of location schemes with qualitative coordinates.

4 Towards an Ontological Account of Robot Navigation

In this section we take the final step of relating elements of the two models of robotic navigation described in Section 2 to the ontology framework sketched in the previous section. Our goal is to show that once this is done, we will have placed the individual formalizations within an ontologically broader context that can be used to shape subsequent design decisions more effectively. The deconstruction will also allow us to clarify some of the differences between the models and to consider how they may best be related to other components of a complete system.

4.1 Ontology and the SSH

As we saw above, the design of the SSH already makes reference to various levels whose elements are described, at least superficially, as originating from an ontology. 'Superficially' here means that the ontological levels described make little contact with generic ontological frameworks of the kind introduced above and so have remained isolated.

The first level of representation in the SSH that can be sensibly related to the terms of a spatial ontology is the causal level. First of all in the SSH implementation presented in [1], the entities at the causal level are tied to the ontology assumed by the Situation Calculus. As we saw above, this ontology is rather neutral, consisting of only first-order entities (*view*, *situation*, and *action*), plus second-order predicates for operating over them (*holds*, *result*, *do*). Here, we consider only the first order terms and not the second-order predicates of the Situation Calculus, since the SSH is not strictly dependent on them.

According to the model, views are *observed* in particular situations. Views, then, are meant to correspond to configurations of real-world states, such as, 'the chair being next to the desk' or 'the cat being on the mat'. These are intended as representational abstractions over the direct sensory input. The precise terms in which such richly structured representations can be formed is a difficult issue. Here a specification in terms of the everyday objects and relations of a 'commonsense' ontology would provide a sensible target for such abstractions however they are constructed. This also fills in the possible definitions of situations to include more than just the agent's state captured by position and orientation. Actions would need to be assimilated under the DOLCE class of perdurants, as they unfold in time. Clearly, any specific ontological characterization of a domain in which the robotic agent is to act, such as an office environment, or a car navigation scenario, should be able to feed directly into the specification of views, actions, and situations. This needs to be managed as a case of importing generic ontology modules and not as a specific piece of additional axiomatization specific to the SSH.

Perhaps more revealing is the consideration of the topological ontological level of the SSH. Here we can see that the definitions provided by the SSH are directly spatial: places are zero-dimensional points, sets of points define regions, paths are boundaries of regions and so on. In other words, the SSH topological level already commits to a particular location scheme for decomposing space. Moreover, it takes on the task of providing reasoning about space within that location scheme. This location scheme gives rise to a SPACE REGION which has points, (directed) lines, and regions, plus a granularity operation by which regions can be collapsed to places. There are, however, other alternatives explored within the qualitative spatial reasoning community[3] and it might be considered beneficial to have a more parameterizable approach whereby these alternatives can also be selected, under appropriate conditions, for driving reasoning. Such modularity is also an intended goal of ontological specification.

This is also significant for communication with such robotic agents. It is now known that mismatching conceptualizations between users and robots can seriously impede effective communication [29, 30]. Fixing a particular location scheme means that space is conceptualized only in the way that that location scheme provides. Concretely for the SSH, this appears to involve just the relations: *on*, *order*, *in*, *right-of* and *left-of*. As Tenbrink in this volume shows [31], the necessary variation in natural communication is considerably higher and would be difficult to support by this scheme alone. When a user is conceptualizing the navigational task in other terms, a re-organizational overhead will be inevitable.

4.2 Ontology and the Route Graph

Whereas the original Route Graph showed more of its origins in a concrete structure for guiding mobile robot navigation by maintaining several rather distinct kinds of information, ranging from perceptual inputs to linguistic node-labels, the latest specification described above (Section 2.2 and [3]), presents a far more ontologically rigorous account. Here, we see distinct information types being separated out, each allocated to its own appropriate ontological subdomain. This is precisely what an adherence to ontological engineering principles requires and allows reuse of as much formal organizational structure in those domains as is available, leveraging off more detailed and established accounts of various aspects of spatial (and other) knowledge.

Much of the specific spatial information that we see incorporated within the SSH model is therefore in the RG case ‘contracted out’. The RG specification commits to certain ontological classes being present in order to function and these serve additionally to position the RG account within a more general lattice of ontological modules. We illustrate this briefly with respect to the basic connectivity of the RG and then turn to more interesting aspects concerned with the explicit inter-ontology relations that are defined.

As we saw above, the basic organizational structure of Route Graphs is defined as a refinement of generic graph theory with graph morphisms for supporting abstraction. Several current ontologies include such graph theory modules and Krieg-Brückner *et al.* [3] also provide such a formalization drawing on the

---

[3] Some of these are also applied within the situation calculus approach: cf. Dylla and Moratz in this volume [28].

algebraic specification language CASL (cf. [23]). Nodes and edges of generic graphs are accordingly mapped to places and route segments within the ontology of route graphs. This kind of inter-module import is crucial to ontology design and several schemes have been proposed (cf. [32, 3]).

The connectivity of a RG is therefore achieved in a very different way to that of the SSH. In the latter, we see a direct modelling of places and connections in terms of spatial relationships; in the former, we see an abstract model of connectivity that is quite distinct from a concrete spatial instantiation. This allows for the possibility of providing more detail concerning the route graph places and edges than would be coherent with a strictly spatial interpretation. We can also see that 'places' as such need to be seen as very different in the RG and SSH, despite their superficially similar functional roles. Places in the RG need to define reference systems: that is, they have an internal spatial structure that is not compatible with the basic zero-dimensional spatial notion of a place within the SSH.

Proceeding further, we can now state that the most generic way of allocating a spatial structure to places is to relate reference systems directly to location schemes in the sense introduced above. At present, the only location scheme that appears to be envisaged in the RG specification is that of angular displacement, or bearing—although even here there are a number of possibilities; for example, schemes can vary according to their granularity (e.g., 360 degrees *vs.* first-quarter, second-quarter, . . .) or according to their orientation (absolute: north, south, east, west *vs.* intrinsic: forwards-backwards-left-right). Since the route graph is also intended to drive robot motion, the more refined and exact reference systems will probably offer more effective choices here. This is quite different when one considers communication, however, where again it is useful to explore the appropriateness of a variety of location schemes. The RG class place must also therefore import properties of an 'intrinsically oriented region'; a kind of entity that has internal spatial parts that lie in some specifiable spatial arrangement that can be determined appealing to the spatial relations provided by some location scheme, i.e., a SPACE REGION.

This means that there is actually little ontological difference between a place in a route graph and the route graph as a whole: both may be related to some spatial region. And this should not come as a surprise since this is precisely what the RG abstraction relation abstractsTo enforces. A place may stand for an entire route graph: the selection is one of granularity.

We see an analogous situation with the ontological class of SITE that we adopted above from BFO. Sites have an occupant, and that is the agent that is situated within the route graph. At one degree of resolution, then, a route graph might most naturally be related to a site. The route graph defines those places that the occupant of the site can be. The site thus defines a certain functional potential for action (movement/navigation), the precise possibilities for which are set out by the connectivity of the route graph and the properties of the route segments. Sites, within our generic ontology, are physical endurants and are therefore positioned in space (via a location scheme). But the occupants of sites are also physical endurants and so the possibility of recursion is built in.

This alignment may be carried further by considering more closely the relations defined between the Route Graph ontology and other ontological domains that Krieg-Brückner *et al.* introduce and which we summarized in Figure 1 above. These 'glue' relations serve to anchor route graphs both to spatial representations and to everyday commonsense representations of the world. We now align these briefly with the resources provided by our generic upper ontology.

The locatedIn relation assumes that some spatial ontology will at least provide REGIONS. This provides a channel for importing any logical specification of the properties and behavior of regions that an adopted location scheme provides. Regions will group route graph places into spatial neighborhoods via their participation in locatedIn relationships. There is so far no guarantee, however, that the connectivity of the route graph is 'well-behaved' with respect to the hierarchy and connection relations over spatial regions. For example, nodes that are immediately connected in the route graph may be spatially positioned in disjoint spatial regions and more distant route graph nodes may be spatially positioned in the same region. Explicitly imposing 'good-behavior' constraints on the route graph during its construction via the properties of spatial regions may be a way of improving the reliability of the construction process—at least for naturalistic route graphs. Route graph neighborhoods would then align with spatial inclusion relationships among regions.

The ContainedIn relationship assumes that some commonsense ontology will provide ENVIRONMENTSPACES. Such a class is already provided by the generic ontology class SITE. A RG node may then be containedIn a commonsense SITE. As we suggested above, a route graph node is then picking out the possible (from the route graph perspective) positions that a site's occupant may occupy. Note, however, that this is precisely what a location scheme would offer for a site in any case: a way of picking out positions within the spatial region that is the site's location. Under this interpretation, the route graph as such might even be accomodated as a structuring of space alongside other such possible structurings—e.g., traditional spatial calculi such as the Region Connection Calculus (RCC: [33]) or, more directly related to current work with route graphs, dipoles [28]. We might then hypothesize that a route graph is a location scheme that decomposes the space of a site according to the possibilities for movement within that site rather than, for example, according to connection and overlap of regions and subregions. A move that may well also be more in line with embedded cognition approaches to space.

This then raises a precisely analogous situation to the 'good-behavior' constraints for spatial inclusion mentioned for the locatedIn relation. When building up a hierarchical nesting of route graphs, there is no *a priori* guarantee that that nesting will reflect commonsense categories. The fact that some nodes are to be containedIn in a room and some others in a connecting corridor is not a necessary consequence of the connectivity of the route graph.[4] Moreover, as noted in Krieg-Brückner *et al.*, there may even be *differences* between the 'commonsense' decompositions provided by robot and user and, even, between different users. For a cognitive agent that is following a route graph to be able to communicate with its users, therefore, we need both to impose sensible ContainedIn relationships and to negotiate these with the particular commonsense SITE categories that some particular user is adopting. The potential for confusion that this raises for communication within a RG-based model is discussed at greater length in Ross *et al.* in this volume [34].

The covers relationship assumed by the RG specification between the spatial and commonsense ontologies stands in for an entire complex of issues, several of which we have discussed above. We have already seen both the spatial region and environment space entities that this relation relies upon. The relationship itself is then simply that of SpatialLocation between a site and its position. How that positioning is achieved is defined by an adopted location scheme.

We can now state some simple correspondences between the RG specification and the classes and relations available in the generic ontology. We write the SpatialLocation relation as l_{LS}, indicating that we consider location to be always relative to a specified location scheme (LS), and the spatial 'part of' relation as $\subset$. The parameters of the relations are also informally typed, drawing on RG or generic ontology constructs as appropriate.

- $covers(Region : R,\ Site : S) \longleftrightarrow l_{LS}(S) = R$
- $containedIn(Place : P,\ Site : S) \longleftrightarrow l_{LS}(P) \subset l_{LS}(S)$
- $locatedIn(Place : P,\ Region : R) \longleftrightarrow l_{LS}(P) \subset R$
- $marks(Place : P,\ PED : x) \longrightarrow l_{LS}(P) = l_{LS}(x)$
- $faces(Place : P,\ PED : x) \longrightarrow l_{LS}(\textit{in-front-of}\, P) = l_{LS}(x)$[5]

In the case that the adopted location scheme is itself that of a route graph, the left-hand location operators can then be ommitted from the generic ontology statements. This is because the route graph specification is then naturally itself already the 'location'. The correspondence for the marks relation, for example, becomes:

$$marks(Place : P, PED : x) \longrightarrow P = l_{RG}(x)$$

That is, the location of the physical endurant x, with respect to a location scheme that is a route graph, is the place in that route graph that marks x. In general, however, the location scheme will not be a route graph and other schemes will mediate attributions of location.

---

[4] Although the free-space geometry naturally generates hypotheses.

[5] Whereas marks is essentially a simple 'naming' relationship based on identity of position, the faces relation includes instrinsic relative-orientation information, in particular, the 'functional relationship' [35] of being 'in front of'. We include this here without further discussion, although its relationship to the possibilities offered by the location scheme needs to be spelled out in more detail.

Finally, we note that these inter-ontology relations themselves provide a backbone for relationships that may be specialized to express more or less detailed correlations between the two levels of ontological abstraction.

4.3 An Illustration of the Benefits of an Ontological Foundation

We have suggested that it is beneficial to embed robotic navigation models within a more generic framework grounded in a broad area of commonsense and spatial information. One area where this is particularly evident is in the relation between navigation models and higher cognitive functionalities, such as communication. We believe that the ontological placement of the various navigational components provides a much improved foundation for building in sophisticated communication functionality. We illustrate this by example, showing how a simple linguistic utterance concerning navigation requires activation of all of the distinct components of the model. Any reduction in the range of ontological modules employed brings with it an automatic restriction in the range of communicative functionality that can be supported.

The following utterance is a realistic directive given to a robotic system in an office environment such as we are working with:

> "Go to the window and follow the corridor until the last room on the right."

We will assume that speech recognition and grammatical and semantic analysis have been carried out so that what remains is a shallow semantic representation unresolved against context (cf. [36, 34]). There still remain significant problems for a navigation system, however; particularly we will consider the process of mapping between such a shallow semantic representation and concrete actions that can be carried out by a robotic system on the basis of a navigational representation such as the route graph. Note that we envision a situation here where there is an ongoing interaction between user and robot—the environment may not have been completely mapped out and labelled. And even there, it is still possible, perhaps likely, that a user may deviate from that labelling.

First "go to the window": the shallow semantics makes it clear that a movement action is being called for and that the destination of that movement is being labelled *by the user* as belonging appropriately to a semantic type window. The use of the definite article also sets a reference resolution problem, the robotic system with dialogue component must be able to locate a real-world entity (PED) that is considered describable by the semantic type. The particular entity will be revealed either by the discourse context (i.e., the window we have been discussing) or by perception. For the latter, we make use of the additional fact that any such utterance must in general be seen as selecting a certain 'ontological granularity'. In this case, the use of the semantic type 'window' selects a subdomain of a commonsense ontology concerning SITES with bona fide boundaries, such as walls, doors, windows. The task of the navigation system using a RG is then to locate a suitable node marked as corresponding to a PED 'window' or at least facing such an entity.

If the nodes of the RG have already been partitioned according to containedIn relationships, then the search for a node can be restricted to a neighborhood defined by an appropriate site. Note that no assumption of the kind that the node corresponding to the window (or facing or marking the window) is an immediate neighbor of the current position holds. The site may itself be a complex arrangement with its own internal route graph structure. If the nodes of the RG have not been so partitioned, then the RG-reasoner will need to follow edges until candidate nodes have been found.

There is also, however, the very real possibility that the user and the robot *disagree* about how exactly a specified physical object is to be described. What for the user may be a window could for the robot be labelled as a 'glass door'. For the robot to resolve this problem gracefully, it is necessary to invoke the commonsense ontological information that both windows and glass doors can form (parts of) boundaries of certain sites and that there is a certain confusion likelihood because of the similarity in the material of the entities. This kind of 'flexible' reference is only possible when substantial real-world knowledge is available; and a link to a commonsense ontology provides just this. Note that this problem can arise *whenever* a semantic type is used: the assumption that such types can be unproblematically resolved by appropriate linguistic labelling of a navigation graph is unfounded. Given sufficient uncertainty, the robot can also engage in precisely focused clarificatory dialogue, for example "do you mean the glass door over there?"

Assuming that the navigation system has located a node in the route graph window that both the robot and the user agree is the 'window', we come to the next component of the directive: "follow the corridor". As before, this may involve confusions concerning 'what is a corridor' while the utterance itself selects a certain ontological granularity. Here we are dealing with SITES such as rooms, corridors, lifts, and so on. The RG nodes containedIn the corridor define the search space for the subsequence search for the 'last room on the right'. In general, however, the precise set of nodes belonging to this corridor might not yet be clear; this can then lead to targetted clarification dialogues such as "are we still in the corridor?"

We can also employ the connection described above between, on the one hand, the potentially recursive structure of SITES and, on the other, route graphs and nodes related by the RG abstractsTo relationship. If we define a relation $\mathsf{containedIn}_{collective}$ that relates the collection of all the RG nodes containedIn a given SITE to that SITE, then this is equivalent to a composition of abstractsTo and marks. Thus, when the linguistic utterance selects a granularity of SITE appropriate for corridors, this also corresponds within the RG to stating that there is some set of nodes that stand in an abstractsTo relation to a more 'abstract' node that may be marked by the site 'corridor'. A corridor will also bring with it from the commonsense ontology attributes that can be used to constrain appropriate collections of RG nodes that are to be grouped: for example, that it is essentially a path with exits.

Finally, assuming that the navigation system has found a collection of RG nodes that user and robot agree can be called a ‘corridor’ (which automatically allows a set of `containedIn` relations, an `abstractsTo` relation, as well as a `marks` relation to be recorded if not already present), the RG reasoner can be given the task of locating a sequence of nodes of some type that are all on the ‘right’ hand side of the path through the corridor. Here there needs to be explicit communication with the spatial ontology and some particular location scheme that supports interpretations of ‘left’/’right’ and such *intrinsic* references. Moreover, the problem of potential confusion occurs here as always: the semantic type ‘room’ used in the last component of the directive may refer for the user to a very different set of potential places than it does for the robot; more discussion of this particular problem is given in the description of our dialogue system presented by Ross *et al.* [34].

Without the provision of the commonsense ontology, suggesting possible candidates for confusion and fixing appropriate granularities, the spatial ontology, for determining spatial relationships, and the route graph ontology itself for handling navigation, flexible communication and, above all, natural resolutions of communicative problems during that communication, would not be conceivable.

5 Conclusion

We have seen that most of the ‘knowledge-level’ issues that are involved in robotic modelling and which have been considered within the SSH and RG models also have their correlates in a thorough formal ontological modelling of space and the possibilities for movement that space entails. Explicitly building into a model possibilities for importing and exporting the information necessary should provide for significantly improved development. The graph-like nature of a navigation graph, for example, can be modelled directly by importing a generic ontology of graphs and graph morphisms. Also, the explicit link between navigation graphs and spatial regions provides access to calculi for reasoning about the various connection relations, etc.

A further benefit is that with such ontologies in place, the relation of further components, for example, those of natural language, to the various ontological levels of robot navigation is clarified. Kuipers suggests that verbal route directions correspond naturally to the causal level of the SSH, i.e., as sequences of imperative corresponding to actions [1, p228]; our discussion of the previous section should have made clear how far this is away from the flexibility required for genuinely natural interaction and its possibilities for misunderstanding and self-correction. The RG account, while not itself providing a model of natural language interaction, comes closer in that the interfaces between the ontological modules are identified and defined.

We can therefore place approaches to robot navigation along a continuum ranging across: (i) no use of an ontological foundation, where an account has to provide its own knowledge modelling and reasoning capabilities from scratch; through (ii) the adoption of ontological modules as a design principle, although the contents of these modules are also provided from scratch (cf. the SSH); to (iii) an ontologically aware design that decomposes the problem across distinct modules, only some of which need to be developed for navigation alone (cf. the RG). We have motivated the utility of adopting the third option and are currently developing our own account of human-robot interaction on this basis [30, 34].

References

1. Kuipers, B.: The spatial semantic hierarchy. Artificial Intelligence **19** (2000) 191–233
2. Werner, S., Krieg-Brückner, B., Herrmann, T.: Modelling navigational knowledge by route graphs. In Freksa, C., Brauer, W., Habel, C., Wender, K., eds.: Spatial Cognition II - Integrating Abstract Theories, Empirical Studies, Formal Methods, and Practical Applications. Springer, Berlin (2000) 295–316
3. Krieg-Brückner, B., Frese, U., Lüttich, K., Mandel, C., Mossakowski, T., Ross, R.: Specification of an Ontology for Route Graphs. In: this volume. (2004)
4. Masolo, C., Borgo, S., Gangemi, A., Guarino, N., Oltramari, A.: Ontologies library (final). WonderWeb Deliverable D18, ISTC-CNR, Padova, Italy (2003)
5. Kuipers, B., Levitt, T.: Navigation and mapping in a large-scale space. AI Magazine **9** (1988) 25–43 Reprinted in *Advances in Spatial Reasoning*, Volume 2, Sushing Chen (Ed.), Norwood NJ: Ablex Publishing, 1990.
6. Kuipers, B.: A hierarchy of qualitative representations for space. In: Proceedings of the 10th International Workshop on Qualitative Reasoning about Physical Systems, Menlo Park, CA, AAAI Press (1996) 113–120 AAAI Technical Report WS-96-01.
7. McCarthy, J.: Situations, actions and causal laws. In Minsky, M., ed.: Semantic Information Processing. MIT Press (1968) 410–417
8. Reiter, R.: Knowledge in Action: Logical Foundations for Describing and Implementing Dynamical Systems. MIT Press, Bradford Books, Cambridge, MA (2001)
9. Krieg-Brückner, B., Röfer, T., Carmesin, H.O., Müller, R.: A taxonomy of spatial knowledge for navigation and its application to the bremen autonomous wheelchair. In Freksa, C., Habel, C., Wender, K., eds.: Spatial Cognition I - An interdisciplinary approach to representing and processing spatial knowledge. Springer, Berlin (1998) 373–397
10. Röfer, T., Lankenau, A.: Route-based robot navigation. Knstliche Intelligenz (2002) Themenheft Spatial Cognition.
11. Wallgrün, J.O.: Autonomous construction of hierarchical Voronoi-based route graph representations. In: this volume. (2004)
12. Bateman, J., Farrar, S.: Spatial ontology baseline. SFB/TR8 internal report I1-[OntoSpace]: D2, Collaborative Research Center for Spatial Cognition, University of Bremen, Germany (2004)
13. Guarino, N.: Formal ontology, conceptual analysis and knowledge representation. International Journal of Human-Computer Studies **43** (1995) 625–640
14. Guarino, N., Welty, C.: An overview of OntoClean. In Staab, S., Studer, R., eds.: Handbook on Ontologies. Springer-Verlag, Heidelberg and Berlin (2004)
15. Guarino, N., Welty, C.: Evaluating ontological decisions with OntoClean. Communications of the ACM **45** (2002) 61–65
16. Hayes, P.J.: The naive physics manifesto. In Michie, D., ed.: Expert systems in the microelectronic age. Edinburgh University Press, Edinburgh, Scotland (1979)

17. Farrar, S., Bateman, J.: General ontology baseline. SFB/TR8 internal report I1-[OntoSpace]: D2, Collaborative Research Center for Spatial Cognition, University of Bremen, Germany (2004)
18. Borgo, S., Guarino, N., Masolo, C.: A pointless theory of space beased on strong congruence and connection. In Aiello, L., Doyle, J., eds.: Principles of knowledge representation and reasoning, Proceedings of the 5th International Conference KR96, San Mateo, CA, Morgan Kaufmann (1996)
19. Gärdenfors, P.: Conceptual spaces: the geometry of thought. MIT Press, Cambridge, MA. (2000)
20. Casati, R., Varzi, A.C.: Holes and other superficialities. MIT Press (Bradford Books), Cambridge, MA and London (1994)
21. Bateman, J.A., Farrar, S.: Towards a generic foundation for spatial ontology. In: Proceedings of the International Conference on Formal Ontology in Information Systems (FOIS-2004), Trento, Italy (2004)
22. Cohn, A., Hazariki, S.: Qualitative spatial representation and reasoning: an overview. Fundamenta Informaticae **43** (2001) 2–32
23. Lüttich, K., Mossakowski, T.: Specification of ontologies in CASL. In: Proceedings of the International Conference on Formal Ontology in Information Systems (FOIS-2004), Trento, Italy (2004)
24. Grenon, P., Smith, B.: SNAP and SPAN: Prolegomenon to geodynamic ontology. In: Spatial Cognition and Computation. (forthcoming)
25. Smith, B., Varzi, A.C.: The Niche. Noûs **33** (1999) 214–238
26. Bittner, T., Smith, B.: 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 (2001) 28–43
27. Bittner, T.: Granularity in reference to spatio-temporal location and relations. In: Proceedings of the 15th Florida Artificial Intelligence Research Symposium (FLAIRS). (2002)
28. Dylla, F., Moratz, R.: Exploiting qualitative spatial neighborhoods in the situation calculus. In: this volume. (2004)
29. Moratz, R., Fischer, K., Tenbrink, T.: Cognitive modelling of spatial reference for human-robot interaction. International Journal On Artificial Intelligence Tools **10** (2001)
30. Moratz, R., Tenbrink, T., Bateman, J., Fischer, K.: Spatial knowledge representation for human-robot interaction. In Freksa, C., Brauer, W., Habel, C., Wender, K., eds.: Spatial Cognition III. Springer, Berlin, Heidelberg (2002)
31. Tenbrink, T.: Identifying objects on the basis of spatial contrast: an empirical study. In: this volume. (2004)
32. Farquhar, A., Fikes, R., Rice, J.: The Ontolingua server: a tool for collaborative ontology construction. In: Proceedings of the 10th Banff Knowledge Acquisition for Knowledge Based Systems Workshop (KAW95), Banff, Canada (1996)
33. Randell, D., Cui, Z., Cohn, A.: A spatial logic based on regions and connection. In: Proceedings of the 3rd. International Conference on Knowledge Representation and Reasoning, San Mateo, Morgan Kaufmann (1992) 165–176
34. Ross, R., Shi, H., Vierhuff, T., Krieg-Brückner, B., Bateman, J.: Towards dialogue based shared control of navigating robots. In: this volume. (2004)
35. Coventry, K.R., Garrod, S.C.: Saying, seeing and acting. The psychological semantics of spatial prepositions. Essays in Cognitive Psychology series. Psychology Press (2004)
36. Alshawi, H., ed.: The Core Language Engine. MIT Press, Cambridge, Massachusetts (1992)