Functional topological relations for qualitative spatial representation

Kristoffer Sjöö, Andrzej Pronobis and Patric Jensfelt

***Abstract*— In this paper, a framework is proposed for representing knowledge about 3-D space in terms of the functional *support* and *containment* relationships, corresponding approximately to the prepositions "on" and "in". A perceptual model is presented which allows for appraising these qualitative relations given the geometries of objects; also, an axiomatic system for reasoning with the relations is put forward.**

We implement the system on a mobile robot and show how it can use uncertain visual input to infer a coherent qualitative evaluation of a scene, in terms of these functional relations

I. INTRODUCTION

Having already made great inroads into industrial settings, robotics is now making an effort to enter into environments such as homes, offices or hospitals. These kinds of spaces are, more than anything, *human-oriented,* that is constructed by and for people, used and modified by people, and occupied by people.

As a result, nearly every aspect of those spaces is shaped by the propensities, preferences and mental habits of human beings. From this association, they take on human *semantics* [1], [2], semantics that must be internalized by any robot that is to have a chance of interacting meaningfully with such environments and their occupants.

An important part of this semantics is *spatial relations*. Spatial relations are abstract, functional relationships between entities in space; they show themselves in the way humans speak about space [3], [4], albeit in a limited fashion. Inspired by these psycholinguistic clues, this work aims to imbue a robot with the ability to understand space in terms of two of the most important spatial relations in the human repertoire – "on" and "in". It proposes computational models as well as a first-order logic axiomatic system for the spatial abstractions that underlie these ubiquitous expressions. We demonstrate by experiment that the approach is suitable for automatic extraction of scene descriptions from undertain visual perception.

A. Functional relations

We humans speak of, and think of, reality in certain terms because those terms are useful to us. Abstractions permit us to make sense of the endless variability of the world, allowing for structured learning, planning and communication. Spatial relations are no exception. They represent some aspect of the environment that has functional relevance – if there was none, they would not be used and thus not learned [3], [5]. Related is the notion of object *persistence*, meaning that objects are expected to remain in the same qualitative relation over time, even if the exact geometrical positions change [6].

Functions may be things such as transporting ("groceries in a bag"), protecting ("trophies in a display case"), allowing to dry ("clothes on a line"), communicating a location ("the door on your right") – or any number of others. The variation among these functionalities is infinite; nevertheless, studies of different languages [7], [8] have indicated that there are recurring patterns, clusters of abstract functionality that are instantiated and extended in different ways for different languages and situations.

This work centers on two such clusters: mechanical support and containment, corresponding – although not perfectly – to the English prepositions "on" and "in" respectively. The importance of these concepts, evident from language, as well as their topological nature provide a hint to their potential for organising the world in a manner that can be shared between robots and people.

B. Related work

There has been research into quantifying spatial relations previously. [9] uses results from brain research to isolate geometrical factors that are important to some relations. [10] introduces a computational model in the *Attention Vector Sum*, verifying it against actual human responses. Another model is suggested by [11] in the form of spatial templates, prototypes which can provide a more or less accurate match to a situation. [12] and [13] both present graphical systems in which spatial relations are used for interaction with a user.

None of the above investigate the functionally important topological spatial relations nor are their approaches based on a functional conceptualization, something we believe important as explained above.

Topological relations are surveyed in [14]. One well-known approach is *Region connection calculus* and its variants, which provide a language for expressing qualitative relationships between regions – such as containment, tangential contact etc. Although there is some overlap with the qualitative axioms introduced below, RCC is purely geometrical and does deal with functional relationships.

This paper builds upon initial work published in [15], [16]. These earlier efforts concentrate on only one of the topological relations (ON), whereas the present work introduces the IN relation and proposes an axiomatic system detailing the relationship between IN and ON, providing the means for qualitative high-level reasoning to incorporate topological information.

The authors are with the Centre for Autonomous Systems at the Royal Institute of Technology (KTH), Stockholm, Sweden. This work was supported by the SSF through its Centre for Autonomous Systems (CAS), and by the EU FP7 project CogX and the Swedish Research Council, contract 621-2006-4520 (K. Sjöö)

C. Organisation of this paper

In Section II the concept of topological spatial relations is explained and the specific instances ON and IN are introduced. Section III details a set of first-order logic axioms structuring the relations, and Section IV shows how those axioms can be included in a probabilistic reasoning framework. Section V describes the system as implemented on a mobile robot and verifies its function experimentally. Finally, conclusions and ideas for future work are contained in Section VI.

II. TOPOLOGICAL SPATIAL RELATIONS

Spatial relations represent the configuration of a focus object, or *trajector*, relative to one or more other objects termed *landmarks*. In language, spatial relations are typically divided into different groups based on the salient geometric relationship: *Projective* spatial relations constrain the trajector's location within an essentially *directed* region relative to the landmark. The direction may depend on many factors, such as intrinsic properties of either object, or the frame of reference of an onlooker. Examples in English include "to the left of", "behind" and "past".

Topological relations, in contrast, locate the trajector in some manner that is independent of direction and the location of an observer. Typical examples are "on", "at" and "inside". Topological relations seem to be among the first to be learned in humans [17]. Topology is very useful for structuring space in a systematic, hierarchical way, allowing us to put together sentences such as "my keys are in a briefcase on the desk in my office on the second floor at our branch in New York". This hierarchical property makes for efficient storage and inference. For this reason, this paper focuses on the arguably most significant topological relations, "on" and "in".

A. ON

1) Ideal schema: The word "on" in English carries a central functional meaning: *support* against gravity. This encompasses for example "the book on the table", "the fly on the wall", "the ring on the finger". Other languages extend the concept differently [18], but the support criterion remains central.

Support goes together with other functional aspects, such as *location control*. Location control imposed by one object on another means that the latter moves together with the former, such as is the case with e.g. trays, plates, buses and trains. Other connotations such as attachment or "weighing down" also overlap with the central "support" notion.

2) Computational model: Although mechanical support provides an objective criterion for defining a spatial relation, it is not typically possible for a robot to ascertain that one object is in fact supporting another. Even humans use perceptual models to estimate this, and those models may sometimes fail (see Fig. 1). We have previously suggested a computational model for a robot to be able to make such an estimate from vision [16] – briefly, three numerical criteria are weighed together to produce a quantitative function ON estimating how well one object supports another:

- *Distance*: Since the objects must touch in order for one to support the other, apparent separation between objects (as well as apparent interpenetration) penalizes the function.
- *Stability*: As the (apparent) center of mass of an object moves beyond the area of contact with its support (as in Fig. 1(b)), the function value is decreased.
- *Verticality*: When the contact surface between objects is horizontally oriented, the function is high, dropping off as the surface becomes more vertical.

Using this computational model, it is possible both to evaluate a perceived configuration of objects in terms of how well they correspond to the support relation, and to estimate the most likely configuration if the support relation is given.

This model is restricted to cases when an object is being supported "on top of" another, as opposed to hanging or adhesive support; the latter entail entirely different geometries and would need a separate perceptual model.

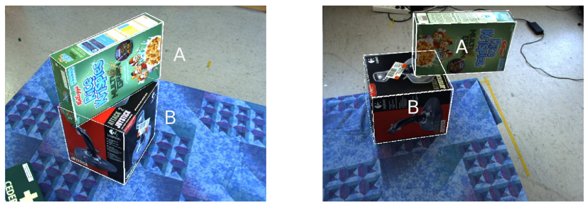

(a) Obvious support relation (b) Incongruent-seeming support relation

Fig. 1. Estimation of support through vision is imperfect

B. IN

1) Ideal schema: "In" as a word has a wider variety of connotations than "on" does. Besides location control and object persistence, "in" often entails aspects of concealment, protection, constraint among others. This variety of meanings is difficult to pin down precisely, but a robust approximation can be found in the idea of *containment*.

Containment signifies the inclusion of most or all of an object within the interior of another object or group of objects. "Interior" is not itself unambiguous, but even with a simple interpretation such as the convex hull, many if not most situations represented by "in" can be covered; Fig. 2 shows two examples of this.

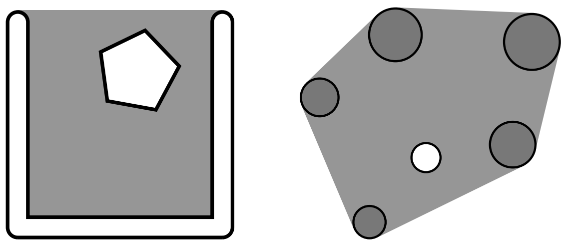

(a) Shape "in" box (b) Circle "in" group

Fig. 2. Convex hull defining "in"

2) Computational model: Containment is computed directly as the proportion of an object O that falls within the convex hull of the container object C (see Fig. 3(a)). This proportion is termed $\text{IN}_{con} \in [0, 1]$.

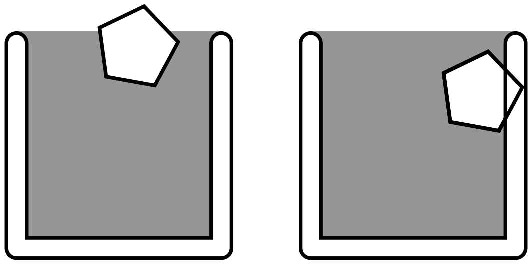

(a) Shape partly "in" box

(b) Shape and box interpenetrating

Fig. 3. Penalties on "in" estimate

However, if this were the only factor determining degree of containment, cases where O and C overlap in space – which is not physically plausible (Figure 3(b)) – would be evaluated the same as realistic configurations. Because such cases are bad examples of the relation, the model is supplemented with a penalty function on apparent object interpenetration:

$$\text{IN}_{pen} \triangleq \begin{cases} 1 & d \geq 0 \\ e^{d/k} & d < 0 \end{cases} \tag{1}$$

where d is the minimum distance between O and C (as defined in Sec. II-A.2) and k a falloff constant.

The total estimate function for the containment spatial relation is taken to be:

$$\text{IN} \triangleq \text{IN}_{con} \cdot \text{IN}_{pen} \tag{2}$$

Both "on" and "in" carry a plethora of additional, metaphorical and indirect meanings that transfer some of the concrete aspects mentioned above into other domains than space by analogy: "on my side", "in theory". Although these are illustrative of the thought processes that support spatial relations and interesting in their own right, the present work shall restrict itself to concrete, spatial usage.

III. AXIOMATIC SYSTEM

One of the main uses for a model that can translate a geometrical relationship between perceived objects into qualitative spatial relations (and back) is to perform high-level reasoning. In order to permit that, a set of rules, or axioms, for the relational predicates is required.

Here follows a suggestion for such an axiomatic system, involving the predicates $\text{On}(x, y)$ and $\text{In}(x, y)$, which are first-order symbols corresponding to the support and containment relations described in Section II. As is inevitable with abstract reasoning, the axioms represent an idealization that will not always apply to the real world. They are reasonable approximations, however, and may be included selectively depending on the application.

Support tends to be *transitive*: if z supports y and y supports x, then z supports x as well. This is obviously not covered by the computational model in Sec. II-A; therefore, a third relation symbol is introduced, termed On_t (for "transitive On"), the properties of which are derived from the axioms.

A. Basic axioms

$$\text{On}_t(x, y) \rightarrow \neg\text{On}_t(y, x) \tag{3}$$
$$\text{In}(x, y) \rightarrow \neg\text{In}(y, x) \tag{4}$$

- (3): Support is antisymmetric
- (4): Containment is antisymmetric

The above also entail irreflexivity ($\neg\text{On}_t(x, x)$, $\neg\text{In}(x, x)$)

B. Transitivity axioms

$$\text{On}(x, y) \rightarrow \text{On}_t(x, y) \tag{5}$$
$$\text{On}_t(x, y) \wedge \text{On}_t(y, z) \rightarrow \text{On}_t(x, z) \tag{6}$$
$$\text{In}(x, y) \wedge \text{In}(y, z) \rightarrow \text{In}(x, z) \tag{7}$$

- (5): Direct support implies transitive support.
- (6): Support is transitive – if y takes the weight of x, and z the weight of y, then that will include x as well.
- (7): Containment is transitive; this is a reasonable assumption given simple geometry and the definition of ON.

C. Interchangeability axioms

$$\text{On}_t(x, y) \wedge \text{In}(y, z) \rightarrow \text{In}(x, z) \tag{8}$$
$$\text{In}(x, y) \wedge \text{On}_t(y, z) \rightarrow \text{On}_t(x, z) \tag{9}$$

$$\begin{aligned} \text{On}_t(x, y) \rightarrow\ & \text{On}(x, y) \\ \vee\ & \exists z. ((\text{On}(x, z) \wedge \text{On}_t(z, y)) \\ \vee\ & (\text{In}(x, z) \wedge \text{On}_t(z, y))) \end{aligned} \tag{10}$$
$$\exists y. (\text{On}_t(x, y) \vee \text{In}(x, y)) \tag{11}$$

- (8): "Generous containment". Typically containment will physically prevent objects from sticking out. This means supported objects will also be contained.
  One consequence of this axiom is that geometrical containment may be violated for In in some cases. Figure 4(a) illustrates, however, that even in such cases functional aspects such as location control, confinement and so forth are often largely preserved and so we tend to extend the use of the word "in" to these cases as well. The axiom is thus intuitively justifiable.
- (9): "Containment provides support". When an object is contained by another, as a rule it is prevented from contact with outside objects and so must receive its supporting force directly or indirectly from the container, as illustrated in Figure 4(b).
- (10): "Support requirement". This is the necessary condition that corresponds to the sufficient conditions in Eqs. (5), (6) and (9), and asserts that an object must

be supported directly by *some* object in order to be indirectly supported.
- (11): “Base requirement”. Every object must be supported by some other object.

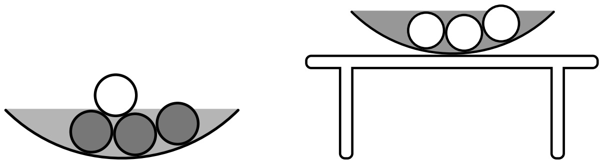

(a) The ball is “in” the bowl (b) The balls are “on” the table

Fig. 4. Effect of interchangeability axioms

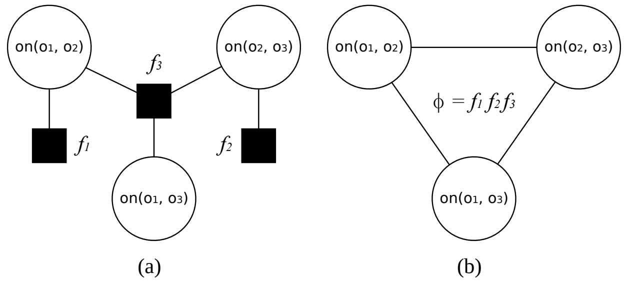

Fig. 5. Factor graph representing “on” object relations connected with a transitivity axiom (a) and a corresponding undirected graphical model (b).

D. Hierarchy axioms

$$\mathrm{On}(x, y) \wedge (y \neq z) \quad \rightarrow \quad \neg\mathrm{On}(x, z) \tag{12}$$
$$\mathrm{On}_t(x, y) \wedge \mathrm{On}_t(x, z) \quad \rightarrow \quad \mathrm{On}_t(y, z) \vee \mathrm{On}_t(z, y) \tag{13}$$
$$\mathrm{In}(x, y) \wedge \mathrm{In}(x, z) \quad \rightarrow \quad \mathrm{In}(y, z) \vee \mathrm{In}(z, y) \tag{14}$$

The hierarchy axioms ensure that the spatial relations form a tree-like structure, which is useful for representation and reasoning.

- (12): Asserts uniqueness of (direct) support. The intuitive justification for this assumption is that an object often is substantially supported by only one other object, and the *majority* of its support nearly always comes from one source.
- (13): Extends the unique-support assumption to the indirect support On_t.
- (14): Although situations can be constructred wherein two containers overlap such that each contains an object, while neither contains the other, such situations are uncommon in practice. Factors such as location control are also unlikely to be present in such cases[1].

E. Using the relational axioms

The axioms proposed in the preceding sections are valuable when processing spatial knowledge on a qualitative level.

A few examples:

- Transitivity and interchangeability axioms allow for deducing In and On_t relations even where not directly given by the computational models.
- Incomplete and qualitative knowledge can be used to guide active search for an object; for example, learning from different sources that “the bowl is on the table” and that “the apple is in the bowl”, the robot can search for the table in order to help find the apple.
- Concrete-support and hierarchy constraints provide the possibility of learning about spatial relations through the process of elimination, given a closed-world assumption.
- Hierarchy constraints furthermore ensure that relations form a tree-like structure and thus make for compact storage (only a few relations need be stored whereas the rest can be deduced), as well as the potential for effectivizing algorithms operating on this knowledge.

In a practical application, obviously a great deal of instance knowledge will apply in addition to the axioms. Many pairs of objects will be patently impossible in the context of On and In; a room cannot be “in” a desk, and that desk can probably not be “on” an apple. Such commmonsense knowledge can be added to the knowledge base to reduce the space of possibilities. Also, for practical applications some objects (such as the floor or the room) must be exempt from Eqns. 10 and 11, as an infinite number of objects would be required otherwise.

IV. Probabilistic inference

In real-world scenarios, the information about objects perceived by a robot is inherently uncertain. This makes it important to provide the ability to transform the axioms defining object relations into a form that permits probabilistic reasoning and integration with probabilistic models such as directed or undirected probabilistic graphs [19]. Here, we introduce a probabilistic representation of axioms and show that such representation can be automatically generated according to the uncertain perception of a scene.

A. Factor-based Representation of Axioms

There is no straightforward way of defining a probabilistic interpretation of the axiomatic system presented above. Except for the fact that configurations contradicting the axioms perforce must have probability 0, nothing is said about the relative likelihoods of permitted configurations. Expressing the axioms through conditional probabilities as in e.g. a Bayes Net [19] will be non-trivial and potentially inefficient, since the relationships expressed are not causal in nature and introduce a great deal of circular cross-dependencies.

One way of introducing probabilities is to use *factor graphs* [20]. Factor graphs are bipartite graphical models, where random variables are represented using variable nodes, connected to each other not directly but via *factor nodes* –

[1]Eqn. (14) implies that, in Fig. 4(a), $\mathrm{On}_t(Ball, Bowl)$ must hold. While this rings true as regards mechanical support, one would not likely say that “The ball is on the bowl”. Here, “in” takes linguistic precedence. However; while this paper gets inspiration from language, it is not primarily about modeling language *per se*.

see Fig. 5(a). Each factor node j defines a function f_j on its connected variables X_j; the joint probability is expressed as

$$p(x_1, \ldots, x_n) = \prod_j f_j(X_j)$$

This factorization makes it easy to encode the various constraints provided by the axioms. For example, if $\mathcal{O}$ is the set of objects, Eqn. (5) becomes

$$\forall \langle o_1, o_2 \rangle \in \{\mathcal{O} \times \mathcal{O}\} : \quad f_{(5)} = \begin{cases} 0, & \begin{array}{l}\mathrm{On}_t(o_1, o_2) \\ \wedge\neg\mathrm{On}(o_1, o_2)\end{array} \\ 1, & \text{otherwise} \end{cases}$$

Similarly, each axiom can be modeled as a factor on every applicable tuple with a value of 0 or 1. The tuples may prove intractable in some cases, such as Eqn. (10). Here, it may be necessary to reduce the set of tuples by heuristically excluding combinations that are impossible, depending on the domain. A typical example would be to divide objects into a group of base objects (e.g. a table) and mobile objects (e.g. a book) and exclude the cases when a base object is On or In any of the mobile objects.

Apart from these axiomatic factors, "probabilistic" factors can be introduced on relations and tuples of relations for which probability needs to be modeled.

An example:

$$\forall \langle o_1, o_2 \rangle \in \{\mathcal{O} \times \mathcal{O}\} : \quad f^{\star} = \begin{cases} \alpha_1, & \begin{array}{ll} & \mathrm{In}(o_1, o_2) \\ \wedge & \textsc{Book}(o_1) \\ \wedge & \textsc{Library}(o_2) \end{array} \\ \alpha_2, & \ldots \end{cases} \tag{15}$$

The above encodes the likelihood, all other things being equal, that objects of different categories are inside containers of different categories. Note that the α:s are not probabilities *per se*; rather they are parameters that, in conjunction with other factors, influence the probabilities of their associated tuples in a systematic way. These parameters are prime candidates for learning. They might also be influenced by other sources of knowledge such as commonsense knowledge about typical man-made environments.

B. *Automatic Generation of the Factor-based Representation*

We have shown that it is possible to establish a direct correspondence between object relations and factor graph variables as well as relation axioms and factor graph factors. This can be used to design an automatic procedure converting an uncertain perception of a visual scene into a probabilistic model performing scene understanding. In the sequel, we propose such a procedure.

Our method takes as input the set of objects, enumerates all object pairs and posits a relation for each pair and relation type. In order to make the reasoning more efficient, it is possible to additionally exclude certain relations which are *a priori* impossible, such as On(A,A). The algorithm subsequently incorporates observations of given object relations, obtained by analysing the visual input as outlined in Section II. Those observations are provided in the form of values in the range $[0, 1]$ quantifying each of the perceived

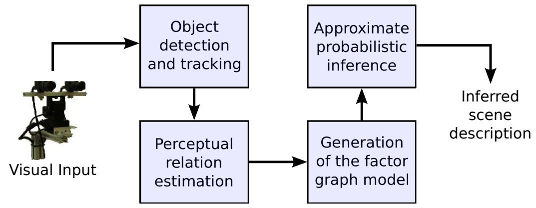

Fig. 6. Data flow through the scene description estimation system.

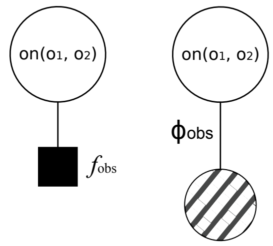

Fig. 7. An excerpt from an undirected graphical model and a corresponding factor graph illustrating the way the uncertain observations of object relations are included

relations. The data flow through the system is presented in Fig. 6.

The algorithm iterates over the possible relations and generates factor graph variables accordingly. Then, it analyses all relation sets matching any of the axioms specified in Section III and introduces an axiom factor for each of them. Finally, factors representing observations are generated for those relations for which the observations are available, as presented in Fig. 7. The following section show that the resulting representation may be successfully applied to the problem of understanding real-world scenes in the presence of uncertain perception.

V. Experiments

In order to show how the system proposed in the preceding sections could be used in robotics applications we have implemented it on a mobile robot. The platform used is a Pioneer III wheeled robot, equipped with a a camera mounted at 1.4 m above the floor. In this experiment the robot was controlled manually so as to place the objects within the view of the camera. We assume the geometries of the objects are known in advance, but not their positions nor the qualitative relations between them.

A. *Vision*

For detection and pose estimation of objects, we are using a system developed at Vienna University [21]. In it, objects are detected using SIFT features trained from a variety of view points; this also provides an initial pose estimate. The pose is refined and tracked using edges and textures.

Given the estimated poses and the known geometries of the detected objects, the perceived values for the functions On and In were computed as described in Section II. Because of noise in the pose estimates, the values obtained fell

(a) Example 1: "A on B on C"

(b) Example 2: "A on B in C"

Fig. 8. Examples of consistent scene evaluation

	Example 1		Example 2	
	Per	Inf	Per	Inf
On(A,B)	92.5%	TRUE	98.9%	TRUE
On_t(A,B)		TRUE[1]		TRUE[1]
In(A,B)	0%	FALSE	0%	FALSE
On(A,C)	4.4%	FALSE	95.2%	FALSE[4]
On_t(A,C)		TRUE[2]		FALSE
In(A,C)	0%	FALSE	16.2%	TRUE[3]
On(B,A)	0%	FALSE	2.1%	FALSE
On_t(B,A)		FALSE		FALSE
In(B,A)	0%	FALSE	0%	FALSE
On(B,C)	96.4%	TRUE	1.7%	FALSE
On_t(B,C)		TRUE[1]		FALSE
In(B,C)	0%	FALSE	99.9%	TRUE

TABLE I

EXAMPLE 1, 2 EVALUATION. "PER" STANDS FOR PERCEIVED VALUE, "INF" FOR INFERRED TRUTH VALUE.

[1]Using Eqn. 5. [2]Using Eqn. 6. [3]Using Eqn. 8. [4]Using Eqn. 13.

Fig. 9. Example 3: an ambiguous scene

within the continuous range $[0, 1]$. Figure 8 shows examples of scenes and Table I the extracted relation values.

B. Inference

Using the set of detected objects and their perceived relation values, the scene was instantiated as a factor graph (Section IV). Each possible relation pair In(x, y), On(x, y), On$_t(x, y)$ was instantiated as a node in the graph, as were the axioms – except that the box "C" was considered a "base object", exempting it from appearing as the first argument in any relation and from needing a support.

The observed values of ON and IN were included as well, as unary factors working on the corresponding nodes. Inference was then performed and the maximum *a posteriori* (MAP) estimate obtained.

C. Results

Figure 8 shows two examples of scenes for which visual processing and inference were performed. The wireframe boxes indicate the object tracker's estimated pose of each object. Table I shows the perceived as well as the inferred values for the relations.

Note that the resulting maximum-a-posteriori solutions obey the axioms. In Example 1, it can be seen that On_t is deduced in accordance with Eqs. 5, 6. Example 2 shows the effect of the interchangeability and hierarchy axioms. Here, both On(A,B) and On(A,C) are indicated by vision, but Eqn. 13 forbids them to be true simultaneously, unless On_t(B,C). Since A already has a support, B, On(A,C) is inferred to be false rather than setting On(B,C) to true. Note also that In(A,C) is made true by Eqn. 8.

In Example 3 (Figure 9), failure to recognize an object means that the object B is seemingly without a proper support. Nevertheless, Eqn. 11 causes On(B,C) to be inferred as the only consistent explanation.

It is seen that the proposed method does indeed produce

	Example 3	
	Per	Inf
On(B,C)	36.9%	TRUE
On_t(B,C)		TRUE
In(B,C)	0%	FALSE

TABLE II

EXAMPLE 3 EVALUATION

consistent qualitative descriptions of a scene, even in the presence of uncertainty, helping to bridge the gap between sensors and metric representations on the one hand and high-level reasoning on the other.

VI. CONCLUSIONS AND FUTURE WORK

In this work we have suggested the use of functional, topological relations based on the notions of support and containment in order to structure spatial knowledge for autonomous robots. An axiomatic system was suggested, consisting of rules that model the first-order logical properties of the abstract relations and that will aid in high-level cognitive activities concerning space. We have demonstrated an implementation of the theory on a real robot and shown that it yields consistent results.

Spatial relations have already been put to use in object search [16], [22], wherein the relations were assumed to be given in advance. A natural next step is to use the axioms to infer the relations likely to hold in a scene and thus create priors for unseen objects, or to aid in tracking.

Another avenue of inquiry is integrating the concepts with computational linguistics, which is appropriate since the this work draws inspiration from language. Spatial relations are important for giving instructions or asking questions about objects; this work should help a robot determine which questions to ask and how to incorporate the answers into its knowledge.

The use of factor graphs to represent the relations and axioms permits their integration with more complete and expressive models directly indicating the type of the modeled relationships between random variables and clearly representing conditional independence between them. As future work, we intend to integrate the factor graph representation of axioms with a complete conceptual spatial knowledge representation within a single chain graph model [23]. Chain graphs are probabilistic graphical models providing a generalization of directed (Bayesian Networks) and undirected (Markov Random Fields) graphical models. As such, chain graphs allow for modeling both "directed" causal as well as "undirected" symmetric or associative relationships, including circular dependencies. In the context of the chain graphs, the presented representation becomes a powerful tool for reasoning about object relations that can easily be incorporated into a more complete probabilistic environment models such as the one presented in [24].

Obviously, this paper has only scratched the surface of the rich repertoire of spatial relations that humans use. Though the schemata of support and containment are doubtless very important, many others as important remain unmodeled out there. It is our belief that the function-based treatment given the relations in this paper can successfully be applied to them as well, helping to build understanding of the world that surrounds us and our robots.

REFERENCES

[1] Y.-F. Tuan, *Space and Place*. University of Minnesota Press, 1977.
[2] T. Jordan, M. Raubal, G. B., and M. Egenhofer, "An affordance-based model of place in gis," in *8th Int. Symposium on Spatial Data Handling (SDH'98)*, 1998.
[3] S. Levinson, *Space in Language and Cognition: Explorations in cognitive diversity*. Cambridge University Press, 2003.
[4] K. Coventry and S. Garrod, *Saying, seeing and acting : the psychological semantics of spatial prepositions*. Hove, 2003.
[5] C. Vandeloise, *Spatial Prepositions: a case study from French*. The University of Chicago press, 1991.
[6] D. B. M. Haun, J. Call, G. Janzen, and S. C. Levinson, "Evolutionary psychology of spatial representations in the hominidae," *Biology*, 2006.
[7] S. Levinson and S. Meira, "natural concepts in the spatial topological domain – adpositional meanings in crosslinguistic perspective: An exercise in semantic typology," *Language*, vol. 79, no. 3, 2003.
[8] A. Herskovits, *Language and Spatial Cognition*. Cambridge University Press, 1986.
[9] J. O'Keefe, *The Spatial Prepositions*, ch. 7. The MIT Press, 1999.
[10] T. Regier and L. A. Carlson, "Grounding spatial language in perception: An empirical and computational investigation," *Journal of Experimental Psychology*, vol. 130, no. 2, pp. 273–2098, 2001.
[11] G. Logan and D. Sadler, *A Computational Analysis*, ch. 13. The MIT Press, 1999.
[12] K. Lockwood, K. Forbus, D. Halstead, and J. Usher, "Automatic categorization of spatial prepositions," in *Proceedings of the 28 th Annual Conference of the Cognitive Science Society.*, 2006.
[13] J. Kelleher, *A Perceptually Based Computational Framework for the Interpretation of Spatial Language*. PhD thesis, Dublin City University, 2003.
[14] A. Cohn and S. Hazarika, "Qualitative spatial representation and reasoning: An overview," *Fundamenta Informaticae*, 2001.
[15] K. Sjöö, A. Aydemir, T. Mörwald, K. Zhou, and P. Jensfelt, "Mechanical support as a spatial abstraction for mobile robots," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, October 2010.
[16] A. Aydemir, K. Sjöö, and P. Jensfelt, "Object search on a mobile robot using relational spatial information," in *Proceedings of the 11th International Conference on Intelligent Autonomous Systems (IAS'10)*, (Ottawa, Canada), September 2010.
[17] J. Piaget and B. Inhelder, *The child's conception of space*. Routledge, 1948.
[18] S. Levinson, *Encyclopedia of cognitive science*, ch. Spatial Language. Nature Publishing Group, 2003.
[19] C. M. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006.
[20] F. R. Ksischang, B. J. Frey, and H.-A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE TRANSACTIONS ON INFORMATION THEORY*, vol. 47, pp. 498–519, February 2001.
[21] T. Mörwald, J. Prankl, A. Richtsfeld, M. Zillich, and M. Vincze, "BLORT - The Blocks World Robotic Vision Toolbox," in *Best Practice in 3D Perception and Modeling for Mobile Manipulation (in conjunction with ICRA 2010)*, 2010.
[22] K. Sjöö, A. Aydemir, D. Schlyter, and P. Jensfelt, "Topological spatial relations for active visual search," Tech. Rep. TRITA-CSC-CV 2010:2 CVAP317, Centre for Autonomous Systems, KTH, Stockholm, July 2010. http://www.csc.kth.se/˜patric/publications/cvap317-avs.pdf.
[23] S. L. Lauritzen and T. S. Richardson, "Chain graph models and their causal interpretations," *Journal Of The Royal Statistical Society Series B*, vol. 64, no. 3, pp. 321–348, 2002.
[24] M. Hanheide, N. Hawes, C. Gretton, A. Aydemir, H. Zender, A. Pronobis, J. Wyatt, and M. Gbelbecker, "Exploiting probabilistic knowledge under uncertain sensing for efficient robot behaviour," in *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI'11)*, (Barcelona, Spain), 2011.