

Dealing with Semantic Knowledge in Robotics with a Probabilistic Description Logic

Fabiano Corrêa, Rodrigo Belizia Polastro, Fabio Gagliardi Cozman, Jun Okamoto Jr.

Escola Politécnica da Universidade de São Paulo, Brasil

Abstract. One has often to deal with large quantities of data in robotics, either coming from sensors or from background knowledge. Background knowledge, with attached semantics, are usually modeled logically, and sensor data, due to uncertainties concerning their nature, are modeled probabilistically. In this paper we present a scalable method for spatial mapping of indoor environments, through the use of a probabilistic ontology. Reasoning with this ontology allows segmentation and tagging of sensor data acquired by a robot during navigation. We report experiments with a real robot to validate our approach, thus moving closer to the goal of integrating mapping and semantic labeling processes.

1 Introduction


As robots have moved from static indoor environments to dynamic outdoor applications, displaying complex interaction patterns, the interest in semantic knowledge in robotics has grown [15]. For instance, semantic knowledge can offer substantial help in 3D reconstruction of environments [24] and in transferring knowledge learned from one environment to different ones [25]. Recent research has actually investigated the acquisition of semantic knowledge directly from robot sensors [11], moving beyond classic work on (basically) hand-coded semantic knowledge [3,19].

As stated by Hertzberg and Saffiotti [15], two points must be present in applications to fully use semantic knowledge in robotics. The first point is that an explicit representation of knowledge must be available to the robot (that is, an ontology for the domain of interest must be present). The second point is that symbols used in the representation must be grounded into physical objects that can be detected by robot sensors. Even though several proposals claim to be using semantics in robotics when automatically classifying sensor data in categories, most such proposals do not reason about objects in the domain. One of the difficulties in reasoning about objects is the presence of uncertainty in real world robotics. While most languages for encoding semantic knowledge are based on logic, in robotics one must deal with sensors, actuators and changing environments; dealing with the unavoidable uncertainties is essential. There is a growing interest in combining logic and probability [13] in the most different fields, with particular robotic applications in semantic mapping [10,20]. In this paper we benefit from this literature, focusing on a probabilistic description logic.


This article shows some improvements regarding a previous ongoing work, that proposes a combination of traditional methods in robotics to find known objects in images and to register 3D points, while using a probabilistic description logic to match datasets with areas of the environment. We can split sensor data into smaller clusters, map them separately, and assemble them together to build a topologic semantic map. This scheme allows us to handle large environments with distinct parts. We thus move towards the goal of adding high-level abilities to robotic navigation. We present the basic idea in Section 2, and the probabilistic ontology in Section 3. Section 4 describes experiments with real data and Section 5 concludes the paper.

2 Semantic mapping in large environments


Mapping environments requires dealing with huge amounts of data. To produce scalable solutions to this problem, it is necessary to explore similarities in patterns. For instance, indoor environments consist of rooms and hallways, and different objects in each one of these areas help in characterizing the location. An additional difficult in mapping, particularly 3D mapping, is data association; that is, registering sensor readings from two distinct positions. If we can label data with the location they come from, say ceiling, wall, kitchen, office, the possible combinations of data association are minimized.



(a) Monitor.



(b) Table.



(c) Pioneer 3-AT.

Fig. 1. Correspondence between template objects and scenes acquired by a mobile robot.

There is current interest in annotating observed structures with semantically meaningful labels such as “building” or “tree” [1]. The classification of data in such categories simplifies data association [23] and leads to maps that are useful beyond navigation; for instance, a map can be useful for task planning as well [12]. Such work is often called *semantic mapping*.

A quick literature review on the use of semantic knowledge in robotics is in order. Galindo et al. [12] present an ontology for indoor environments, and exploit the semantics for task planning. They investigate the deduction of new information and the use of this information to improve plans. Their spatial hierarchy is combined with a conceptual hierarchy expressed with a description logic, and both hierarchies are related through anchoring [4]. Another relevant previous work has been produced by Limketkai et al. [20], who exploit semantics to construct a map from laser data. They use Relational Markov Networks [31] to classify lines processed from laser sensor data. Doors and corridors are identified, and relations are used to improve discrimination. Yet another relevant work has been produced by Wang and Domingos [33], who apply Markov Logic networks to the same domain, developing extensions that handle continuous variables. With regard to sensor data classification, several authors aim at creating topological maps by clustering data points. For instance, Posner et al. [28] shows how to classify outdoor scenes, using odometry to provide continuity in classification. Zivkovic et al. [34] use omnidirectional images of indoor environments and propose a method that builds a topological map. Vasudevan et al. [32] use identification of certain objects (with SIFT [21]) to create a map of objects.

We propose in this paper an offline scheme for semantic mapping. The robot navigates through the environment, collecting images, laser and odometry data, and then all data are processed at once. The idea is to identify objects in the images and to determine the label of the location each image was taken. We then segment data based on location labels, thus creating a topological map. We run a SLAM algorithm to the individual locations and create maps of smaller areas of the environment. For instance, data from an office are used to reconstruct spatially only the office. Odometry allows us to determine the topology of the environment and to unify all metric maps into a single one.

Our implementation employed SIFT to obtain features and the logic *CRALC* to encode semantic knowledge about environments. The figure also shows the robot Pioneer 3-AT, mounted with a SICK laser to obtain 3D data and a camera to recognize objects in images, as used in the experiments. The classification of objects in images was done with features robustly extracted using the SIFT algorithm; for that, objects described in the ontology have a SIFT descriptor previously computed. Figure 1 shows two templates of objects. Not all sensor data are labeled as some images may contain objects not matched against SIFT descriptors. But using the principle of continuity, sensor data not labeled between two or more with identical labels receives the same label.

After objects are identified, the probabilistic description logic ontology in *CRALC* is used to label sensor data. Inferences are conducted in a Bayesian

network generated from the $\text{CR}\mathcal{ALC}$ model, thus producing probabilities for the various types of environments.

3 Representing environments in $\text{CR}\mathcal{ALC}$

To reason about objects contained in an environment, we propose in this section a probabilistic ontology, expressed in a probabilistic description logic. Description logics are largely used to build ontologies, as they usually contain a fragment of first-order logic and can organize concepts into hierarchies [2]. For instance, an “office” can be described as a set of “chairs”, “tables” and “computers” in some structured manner. The challenge is that there is usually uncertainty attached to descriptions of environments in robotics, and there is always uncertainty associated with sensor data. In practice two instances of the same label are often distinct; for instance, two parks may have completely different vegetation but both still need to be labeled as “park”. Standard description logics cannot model alone these matters. We thus resort to a description logic that allows probabilities to be attached to its formulae. Even though some of these probabilities can easily be estimated, most of the probabilities cannot, and so we must deal with imprecision in probability values. An extreme case would be to leave the probability of a particular relation to be in the interval $[0, 1]$, with no further constraints.

3.1 Credal \mathcal{ALC}

Cozman and Polastro have recently proposed a probabilistic description logic, called Credal \mathcal{ALC} ($\text{CR}\mathcal{ALC}$) [6,7,27], that adopts an interpretation-based semantics and resorts to graph-theoretical tools so as to allow judgements of stochastic independence to be expressed. This logic was chosen over several other probabilistic description logics in the literature [9,14,17,18,22,5,30] as it is based in the popular \mathcal{ALC} logic and due to its interpretation-based semantics and connection with Bayesian networks (hence leading to relatively efficient inference algorithms).

The vocabulary of $\text{CR}\mathcal{ALC}$ contains *individuals*, *concepts*, and *roles*. Concepts and roles are combined to form new concepts using a set of *constructors* from \mathcal{ALC} [29]: *conjunction* ($C \sqcap D$), *disjunction* ($C \sqcup D$), *negation* ($\neg C$), *existential restriction* ($\exists r.C$) and *value restriction* ($\forall r.C$). A *concept inclusion* is denoted by $C \sqsubseteq D$ and a *concept definition* is denoted by $C \equiv D$, where C and D are concepts; we assume in both cases that C is a concept name. We then say that C *directly uses* D ; the relation *uses* is the transitive closure of *directly uses*. Also, the concept \top denotes $C \sqcup (\neg C)$ for some concept C . As in \mathcal{ALC} , the semantics is given by a domain \mathcal{D} , a set of elements, and an interpretation mapping \mathcal{I} that assigns an element to an individual, a set of elements to a concept, and a binary relation to a role. An interpretation mapping must also comply with constructs of the language; for instance, the interpretation of concept $C \sqcap D$ is $\mathcal{I}(C) \cap \mathcal{I}(D)$, while the interpretation of concept $\forall r.C$ is

$\{x \in \mathcal{D} \mid \forall y : (x, y) \in \mathcal{I}(r) \rightarrow y \in \mathcal{I}(C)\}$. Additionally, $\text{CR}\mathcal{ALC}$ accepts *probabilistic inclusions* as follows. A probability inclusion reads

$$P(C|D) \in [\alpha_1, \alpha_2],$$

where D is a concept and C is a concept name. The semantics of such a probabilistic inclusion is, informally:

$$\forall x : P(C(x)|D(x)) \in [\alpha_1, \alpha_2], \quad (1)$$

where it is understood that probabilities are over the set of all interpretation mappings \mathcal{I} for a domain \mathcal{D} . If D is the concept \top then we write $P(C) \in [\alpha_1, \alpha_2]$. Probabilistic inclusions are required to only have concept names in their conditioned concept (that is, inclusions such as $P(\forall r.C|D)$ are not allowed). Yet another type of probabilistic assessment is possible in $\text{CR}\mathcal{ALC}$: for a role r , we can have $P(r) \in [\beta_1, \beta_2]$ to be made for roles, with semantics:

$$\forall x, y : P(r(x, y)) \in [\beta_1, \beta_2], \quad (2)$$

where again the probabilities are over the set of all interpretation mappings for a given domain.

Every ontology is assumed *acyclic*; that is, a concept does not use itself. If we write down an ontology as a directed graph where each node is a concept or role, and arcs go from concepts that are directly used to concepts that directly use them, we obtain that this graph must be acyclic. We refer to such a graph as an *ontology graph*. For instance, consider concepts A, B, C and the role r , and suppose $P(A) = 0.7$, $B \sqsubseteq A$, $P(B|A) = 0.4$, $P(r) = 0.5$ and $C \equiv A \sqcap \exists r.B$. In Figure 2.a we have the ontology graph.

Under some additional restrictions (unique-names assumption, known and finite domain), any ontology expressed in $\text{CR}\mathcal{ALC}$ can be grounded into a Bayesian network, possibly with attached probability intervals [6,7,27]. That is, grounding an ontology with a finite and known domain leads to a *credal network* [8]. In Figure 2.b we have the grounded network for the ontology described in the previous paragraph, for a domain with only 2 individuals. Note that each concept is instantiated with each one of the individuals, while each role is instantiated with each pair of individuals.

A change from previous uses of $\text{CR}\mathcal{ALC}$, is that we implement types in the language so that it is possible to instanciate only some parts of the graph to do inference.

3.2 An ontology for the domain of spatial mapping

We now present a probabilistic ontology suitable for dealing with robotic mapping. We take as base an ontology proposed in [26] and add more topological knowledge between the identified ambients. We start with two primitive concepts **Object** and **Environment**. As $\text{CR}\mathcal{ALC}$ requires that *a priori* probabilities must be




Fig. 2. A small ontology and its grounding for a domain with 2 individuals.

specified for primitive concepts; as there is no prior information on objects, we assign relatively large probability intervals as follows:

$$P(\text{Object}) \in [0.2, 0.8], \quad P(\text{Environment}) \in [0.2, 0.8].$$

We introduce two roles, one to express that an environment contains an object, and the other to express that two objects are near. We leave the probabilities for these roles rather free, using the same intervals:

$$P(\text{contains}) \in [0.2, 0.8], \quad P(\text{near}) \in [0.2, 0.8].$$

Along with Objects and Environments we introduce the concept **Connector** that will be used to describe elements such as doors and hallway junctions and the role **accesses** that tell us which environment can be accessed from each Connector. Again, we use the same intervals.

$$P(\text{Connector}) \in [0.2, 0.8], \quad P(\text{accesses}) \in [0.2, 0.8].$$

We then propose the following object hierarchy, using constructs in $\text{CR}\mathcal{ALC}$:

$$\begin{aligned} \text{InteriorObject} &\sqsubseteq \text{Object}, \\ \text{ExteriorObject} &\sqsubseteq \text{Object}, \\ \text{OfficeObject} &\sqsubseteq \text{Object}, \end{aligned}$$

$$\begin{aligned} \text{Table} &\sqsubseteq \text{InteriorObject}, \\ \text{Chair} &\sqsubseteq \text{InteriorObject}, \end{aligned}$$

$$\begin{aligned} \text{Cabinet} &\sqsubseteq \text{InteriorObject} \sqcap \text{OfficeObject}, \\ \text{Monitor} &\sqsubseteq \text{InteriorObject} \sqcap \text{OfficeObject}, \\ \text{Printer} &\sqsubseteq \text{InteriorObject} \sqcap \text{OfficeObject}, \end{aligned}$$

$$\text{Sign} \sqsubseteq \text{ExteriorObject},$$

Extinguisher \sqsubseteq ExteriorObject,
 Switchbox \sqsubseteq ExteriorObject,

Board \sqsubseteq Object,
 Door \sqsubseteq Object \sqcap Connector,
 Junction \sqsubseteq Connector.

As $B \sqsubseteq A$ implies $P(B|\neg A) = 0$ but implies nothing about $P(B|A)$, in several cases we would have $P(B|A) \in [0, 1]$. We then adopt $P(B|A) \in [0.2, 0.8]$.

Composite objects can now be described:

Desk \equiv Table $\sqcap \exists$ near.Chair, Entrance \equiv Door $\sqcap \exists$ near.Sign.

Finally, the whole environment can be described as:

Room \equiv Environment
 $\sqcap \exists$ contains.Door
 $\sqcap \exists$ contains.Table
 $\sqcap \exists$ contains.Chair
 $\sqcap \neg \exists$ contains.ExteriorObjects
 Office \equiv Room
 $\sqcap \exists$ contains.Desk
 $\sqcap \exists$ contains.Cabinet
 $\sqcap \exists$ contains.Monitor
 Classroom \equiv Room
 $\sqcap \exists$ contains.Board
 $\sqcap \neg \exists$ contains.OfficeObjects
 Hallway \equiv Environment
 \exists contains.Entrance
 $\sqcap \exists$ contains.Extinguisher
 $\sqcap \exists$ contains.Switchbox
 $\sqcap \neg \exists$ contains.InteriorObjects

The network generated by grounding this ontology is relatively dense and contains probability intervals. Inference is nontrivial from a computational point of view. In our experiments we resort to approximate algorithms based on Loopy Propagation [7,16].

4 Experiments

We provide results of an experiment consisting of a robot navigating through three different areas of an indoor environment: a laboratory, a professor's room,

Table 1. Identified environments for 6 datapoints.

Datapoints						
Location	1	5	7	10	13	17
Observations						
Objects	3 chairs	2 chairs	3 doors	2 doors	2 chairs	2 chairs
	1 table	1 table	2 signs	2 signs	2 table	1 table
	1 monitor	2 monitor	1 extinguisher	1 extinguisher	1 monitors	1 monitor
	1 cabinet	1 cabinet	1 switchboard	1 switchboard	1 cabinet	2 cabinets
	1 door	1 door		1 board	1 door	1 door
Inference result						
P(Office)	[0.000 , 0.502]	[0.000 , 0.602]	[0.000, 0.001]	[0.000, 0.001]	[0.000 , 0.554]	[0.000 , 0.554]
P(Classroom)	[0.000, 0.450]	[0.000, 0.532]	[0.000, 0.060]	[0.000, 0.060]	[0.000, 0.385]	[0.000, 0.385]
P(Hallway)	[0.000, 0.020]	[0.000, 0.020]	[0.000 , 0.767]	[0.000 , 0.767]	[0.000, 0.020]	[0.000, 0.020]
Area	1		2		3	

and a hallway connecting both. In this experiment, we gathered images and 3D points with a laser sensor, with the pose of each data gathering location given by odometry and a gyroscope. We picked sequentially 18 points to gather data from the laser and gather images.

Each point was then classified accordingly to the identified objects; the result of inferences for 6 of the points can be seen in Table 1. Whenever two consecutive points have different labels, the data are split into a new area. In this case, we found 3 distinct areas. Table 1 shows the inference values for 2 points for each area. Note that in the ontology the possible environments were not set as mutually exclusive; hence the probabilities are for individual objects and are not required to add to one across objects.

Note that the good result of the inferences are directly related to the use of probability intervals. In [26] exact probabilities where used instead of intervals. When choosing this probabilities, if probabilities such as 0.5 or higher where picked, the results get extremely binary, e.g. if a table where identified, the probability for hallway drops to zero. This is undesirable since the robot may be subjected to misidentified objects, or even identify an unrelated object through a door or window. On the other hand, if smaller probabilities where to be chosen, inferences would result in insignificant values, i.e. all the environments probabilities would be inferred as low in [0.00, 0.15]. And using the open intervals to specify the ontology gives much significant results once mean that we really does not know if an object is relevant to the environment. All environment start with a probability in [0.0, 1.0] and for each identified object, the upper limit of the interval drops slightly for the undesired environment.

With the sensor data set segmented in three different areas, it was possible to map each identified area alone, using only the pertinent sub-set of data to map the 3 areas. To produce the 3D maps, we used 6D SLAM software, available




Fig. 3. An example of a 3d metric map of area 1.

from <http://openslam.org> [23]. Figure 3 shows the top-view of the 3D map from the mapped areas. Figure 3(a) is a top-view of the laboratory, Figure 3(b) is a top-view of the hallway and Figure 3(c) is a top-view of the professor's room. Figure 4 shows a tridimensional view of the environment mapped.




Fig. 4. Three-dimensional view

In Table 2, we see that mapping the different areas separately is faster (even if we process one after the other) than processing all the data to produce one single map. But the approach is even more interesting if we calculate the different areas in parallel, in which the time consumed is the time of the area that consumes

the most, in this case, Cozman's room. The difference in time is due to the calculation of registration between all the sensor readings. With less readings, lesser combinations and faster processing. Experiments were executed in a Intel Core 2 Quad (2.4 GHz / 2.39 GHz).

Table 2. Time comparisons

Mapped areas	time (ms)
Cozman's room	52,186
Hallway	31,670
Jun's lab	7,267
Complete	114,703

5 Conclusion and Future Work

We have proposed a set of techniques that incorporate semantic knowledge into robotic mapping; the techniques are geared towards mapping large domains. The idea is to make the mapping process scalable by breaking it into small units that receive semantic labels specified using a probabilistic description logic. This objective has been achieved, as we have been able to automatically split data in convenient smaller and tagged sets, each one being mapped by itself. Semantic knowledge can have a significant impact in robotics, and we hope to have offered a viable path for substantial future development. Improvement from a previous version of this work includes a clearer ontology to the problem of robotic mapping, inference considering imprecise probabilities, implementation of types in the probabilistic logic language to instantiate smaller graphs for inference, and experiments in a larger environment to really demonstrating the benefits of the approach.

Further work with probabilistic description logic should deal directly with data coming from sensors; that is, the classification should consider more primitive objects as lines from laser or patches of images. We plan to investigate how to learn relations between concepts to extend the ontology as the robot executes its tasks. Future applications will use more deeply the potential of semantic knowledge and the power of probabilistic logic languages.

Acknowledgements

The first author is supported by FAPESP grant 2009/14396-8; the second author is supported by FAPESP grant 2008/53292-0; the third author is partially supported by CNPq. The project has been partially supported by FAPESP process 2008/03995-5.

References

1. D. Anguelov, Ben Taskar, V. Chatalbashev, D. Koller, D. Gupta, G. Heitz, and A. NG. Discriminative learning of markov random fields for segmentation of 3d scan data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2005.
2. F. Baader, D. Calvanese, D.L. McGuinness, D. Nardi, and P.F. Patel-Schneider. *Description Logic Handbook*. Cambridge University Press, 2002.
3. R. Chatila and J.P. Laumond. Position referencing and consistent world modeling for mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 1985.
4. S. Coradeschi and Alessandro Saffiotti. An introduction to the anchoring problem. *Robotics and Autonomous Systems*, 43:85–96, 2003.
5. Paulo C. G. Costa and Kathryn B. Laskey. PR-OWL: A framework for probabilistic ontologies. In *Conference on Formal Ontology in Information Systems*, 2006.
6. Fabio G. Cozman and Rodrigo B. Polastro. Loopy propagation in a probabilistic description logic. In Sergio Greco and Thomas Lukasiewicz, editors, *Second International Conference on Scalable Uncertainty Management*, Lecture Notes in Artificial Intelligence (LNAI 5291), pages 120–133. Springer, 2008.
7. Fabio G. Cozman and Rodrigo B. Polastro. Complexity analysis and variational inference for interpretation-based probabilistic description logics. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, 2009.
8. Cassio Polpo de Campos and Fabio G. Cozman. Computing lower and upper expectations under epistemic independence. *International Journal of Approximate Reasoning*, 44:244–260, 2007.
9. Zhongli Ding, Yun Peng, and Rong Pan. BayesOWL: Uncertainty modeling in semantic web ontologies. In *Soft Computing in Ontologies and Semantic Web*, volume 204 of *Studies in Fuzziness and Soft Computing*, pages 3–29. Springer, Berlin/Heidelberg, 2006.
10. Pedro Domingos and Matthew Richardson. *Introduction to Statistical Relational Learning*, chapter Markov Logic: a unifying framework for Statistical Relational Learning, pages 339–371. The MIT Press, 2007.
11. C. Galindo, A. Saffiotti, S. Coradeschi, P. Buschka, J.A. Fernandez-Madriral, and J. Gonzalez. Multi-hierarchical semantic maps for mobile robotics. In *International Conference on Intelligent Robots and Systems*, 2005.
12. Cipriano Galindo, Juan-Antonio Fernandez-Madriral, Javier Gonzalez, and Alessandro Saffiotti. Robot task planning using semantic maps. *Robotics and Autonomous Systems*, 11:955–966, 2008.
13. Lise Getoor and Ben Taskar, editors. *Introduction to Statistical Relational Learning*. MIT Press, 2007.
14. J. Heinsohn. Probabilistic description logics. In *Conference on Uncertainty in Artificial Intelligence*, pages 311–318, 1994.
15. Joachim Hertzberg and Alessandro Saffiotti. Using semantic knowledge in robotics. *Robotics and Autonomous Systems*, 56:875–877, 2008.
16. Jaime Shinsuke Ide and Fabio Gagliardi Cozman. Approximate algorithms for credal networks with binary variables. *Int. J. Approx. Reasoning*, 48(1):275–296, 2008.
17. Manfred Jaeger. Probabilistic reasoning in terminological logics. In *Principles of Knowledge Representation (KR)*, pages 461–472, 1994.

18. Daphne Koller, Alon Y. Levy, and Avi Pfeffer. P-CLASSIC: A tractable probabilistic description logic. In *AAAI*, pages 390–397, 1997.
19. B. Kuipers and Y.-T. Byun. A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Journal of Robotics and Autonomous Systems*, 8:47–63, 1991.
20. B. Limketkai, L. Liao, and Dieter Fox. Relational object maps for mobile robots. *Proceedings of the International Joint Conference on Artificial Intelligence*, 1:1471–1476, 2005.
21. D. G. Lowe. Distinctive image features from scale-invariant key-points. *International Journal of Computer Vision*, 60:91–110, 2004.
22. Thomas Lukasiewicz. Expressive probabilistic description logics. *Artificial Intelligence*, 172:852–883, 2008.
23. Andreas Nchter. *3D Robotic Mapping*. Springer Tracts in Advanced Robotics, 2009.
24. A. Nuchter, O. Wulf, K. Lingemann, J. Hertzberg, B. Wagner, and H. Surmann. 3d mapping with semantic knowledge. In *RoboCup International Symposium*, 2005.
25. Charles Petrie. The semantics of "semantics". *IEEE Internet Computing*, 13:94–96, 2009.
26. Rodrigo B. Polastro, Fabiano Corrêa, Fabio G. Cozman, and Jun Okamoto Jr. Semantic mapping with a probabilistic description logic. In *Simpósio Brasileiro de Inteligência Artificial (submitted)*., 2010.
27. Rodrigo B. Polastro and Fabio G. Cozman. Inference in probabilistic ontologies with attributive concept descriptions and nominals. In *4th International Workshop on Uncertainty Reasoning for the Semantic Web (URSW) at the 7th International Semantic Web Conference (ISWC)*, Karlsruhe, Germany, 2008.
28. I. Posner, D. Schroeter, and P. Newman. Using scene similarity for place labeling. In *Proceedings of the 10th International Symposium on Experimental Robotics*, 2006.
29. M. Schmidt-Schauss and G. Smolka. Attributive concept descriptions with complements. *Artificial Intelligence*, 48:1–26, 1991.
30. Fabrizio Sebastiani. A probabilistic terminological logic for modelling information retrieval. In W.B. Croft and C.J. van Rijsbergen, editors, *17th Annual International ACM Conference on Research and Development in Information Retrieval (SIGIR)*, pages 122–130, Dublin, Ireland, 1994. Springer-Verlag.
31. Ben Taskar, Pieter Abbeel, Ming-Fai Wong, and Daphne Koller. *Introduction to Statistical Relational Learning*, chapter Relation Markov Networks, pages 175–199. The MIT Press, 2007.
32. Shrihari Vasudevan, Stefan Gchter, Viet Nguyen, and Roland Siegwart. Cognitive maps for mobile robots - an object based approach. *Robotics and Autonomous Systems*, 55:359–371, 2007.
33. Jue Wang and Pedro Domingos. Hybrid markov logic networks. In *Proceedings of the 23rd National Conference on Artificial Intelligence*, volume 2, pages 1106–1111, 2008.
34. Zoran Zivkovic, Olaf Booij, and Ben Krse. From images to rooms. *Robotics and Autonomous Systems*, 55:411–418, 2007.