Semantic Mapping with a Probabilistic Description Logic

Rodrigo Polastro, Fabiano Corrêa, Fabio Cozman, and Jun Okamoto Jr.

Escola Politécnica da Universidade de São Paulo

Abstract. Semantic mapping employs explicit labels to deal with sensor data in robotic mapping processes. In this paper we present a method for boosting performance of spatial mapping, through the use of a probabilistic ontology, expressed with a probabilistic description logic. Reasoning with this ontology allows segmentation and tagging of sensor data acquired by a robot during navigation; hence a robot can construct metric maps topologically. 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

Until a few years ago, robotic mapping tasks were restricted to 2D reconstruction of environments [1]. With recent improvements in sensors and algorithms, researches are now aiming at 3D maps of large environments with ever increased detail, using both laser sensors and images of vision systems [2].

There is current interest in annotating measured structures with semantically meaningful labels such as "building" or "tree" [3]. The classification of data in such categories allows dimensionality reduction in mapping related problems [4], as well as the construction of maps that are useful beyond navigation (for instance, useful for task planning [5]). Work on this topic is grouped under the name of semantic mapping. As stated by Hertzberg and Saffiotti [6], two points must be present in applications to fully use semantic knowledge in robotics:

1. an explicit representation of knowledge available to the robot; that is, an ontology for the domain of interest;
2. the need for grounding the symbols used in this representation in 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, if those categories are not related to each other, there is no semantic reasoning involved.

Galindo et al [5] have presented one of the few proposals that really explore semantics in depth. They provide an ontology for indoor environments and use it to reason during task planning. Two ways of exploiting semantics are investigated: deduction of new information, and improvement in task planning efficiency. They

A.C. da Rocha Costa, R.M. Vicari, F. Tonidandel (Eds.): SBIA 2010, LNAI 6404, pp. 62–71, 2010.
© Springer-Verlag Berlin Heidelberg 2010

combine a spatial hierarchy with a conceptual hierarchy expressed with a description logic. One hierarchy is related to the another through anchoring [7].

Limketkai et al [8] propose a more integrated mapping process, although they only explore semantics to construct a map from laser data. They use a first-order probabilistic model based on Relational Markov Networks [9] to classify lines processed from laser sensor data, so as to identify doors and corridors in the environment. Spatial and relational information between the objects in scene are used to facilitate the discrimination. Wang and Domingos [10] apply Markov Logic networks (with extensions to handle continuous variables) to that same domain.

Regarding only sensor data classification, some authors propose different methods to cluster sensor data with the objective of creating a topological map of the environment. Posner et al [11] consider as a scene whatever data a robot acquires with its sensors, and shows how to classify outdoor scenes (they use odometry information to provide continuity in the classification). Zivkovic et al [12] use omnidirectional images of indoor environments and propose a method that identifies images that composes different rooms, constructing thus a topological map.

Vasudevan et al [13] use identification of certain objects and their spatial position in the environment to create a map of objects. They detect doors with a laser sensor and use SIFT [14] to detect other objects from images.

In this paper, we propose 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 relate datasets with different areas of the environment. This way we can split sensor data in small clusters, map them separately, and them assemble them together on a topologic semantic map. Thus moving towards the goal of adding high-level abilities to robotic navigation.

2 Representing Environments in CR$\mathcal{ALC}$

This paper introduces a method for boosting spatial mapping through reasoning over objects contained in the environment. To achieve that, relationships between the environment and the objects in it should be first modeled. We resort to a *description logic* for such modeling [15]. Description logics are largely used to build ontologies, as they usually contain a fragment of first-order logic and can organize concepts into hierarchies. Such a logic seems to be the right tool in the present context, as we need a language to describe high-level labels; for instance, an "office" must 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 contexts in robotics; in particular, there is always uncertainty associated with sensor data. Besides, in practice two instances of the same label are often distinct; that is, two different offices contain chairs and tables, but the quantity of this objects is unique to each office; moreover one of them may have vending machines for its employees while the other may not. As another example, two parks may have completely different vegetation but both still need to be labeled by a single concept "park". Standard description logics cannot model alone these matters. We thus resort to a *probabilistic description logic*; that is, a description logic that allows probabilities to be attached to its formulae.

2.1 Credal CR$\mathcal{ALC}$

Cozman and Polastro have recently proposed a probabilistic description logic, called CR$\mathcal{ALC}$ [16,17], 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 other probabilistic description logics [18,19,20,21,22,23,24] as it is based in the popular $\mathcal{ALC}$ logic and attends to our needs.

The vocabulary of CR$\mathcal{ALC}$ contains *individuals*, *concepts*, and *roles* [15]. Concepts and roles are combined to form new concepts using a set of *constructors* from $\mathcal{ALC}$ [25]. These constructors are *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.

CR$\mathcal{ALC}$ also supports probabilistic inclusions. A probability inclusion reads $P(C|D) = \alpha$, where D is a concept and C is a concept name. The semantics of such a probabilistic inclusion is:

$$\forall x : P(C(x)|D(x)) = \alpha, \tag{1}$$

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 "true" concept $\top$, then we simply write $P(C) = \alpha$. 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).

We also allow assessments such as $P(r) = \beta$ to be made for roles, with semantics

$$\forall x, y : P(r(x, y)) = \beta, \tag{2}$$

where again the probabilities are over the set of all interpretation mappings.

We assume that every terminology is acyclic; that is, a concept does not use itself. Under some additional restrictions (unique-name assumption, precision of probability assessments, known and finite domain), any terminology in CR$\mathcal{ALC}$can be grounded into a Bayesian network [16,17]. For better understanding of the grounding, consider the 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 1 we have: a) ontology graph; b) the grounded network for 2 individuals.

2.2 An Ontology for the Domain of Spatial Mapping

To characterize the status of a robot's location with the information of its sensors, we contribute proposing a probabilistic ontology as follows. We start with two primitive concepts $\mathsf{Object}(x)$ and $\mathsf{Environment}(x)$. As CR$\mathcal{ALC}$ requires that *a*

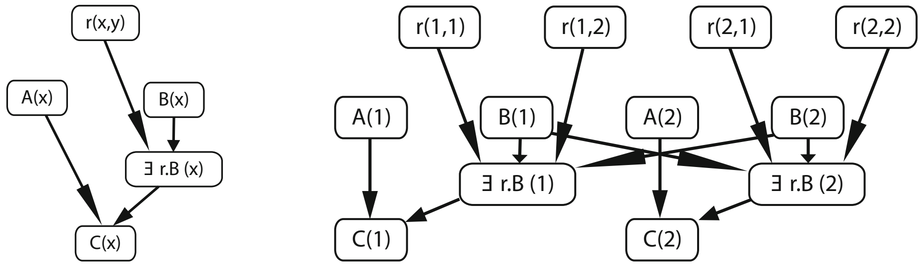

(a) Ontology graph. (b) Grounding.

Fig. 1. Graph for a simple ontology and its grounding for a domain with 2 individuals

priori probabilities must be specified for primitive concepts; as there is no prior information on objects, we assign relatively neutral probabilities as follows:

$$P(\mathsf{Object}) = 0.5,$$
$$P(\mathsf{Environment}) = 0.5.$$

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 neutral, using the same probabilities:

$$P(\mathsf{contains}) = 0.5,$$
$$P(\mathsf{near}) = 0.5.$$

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

$$\begin{aligned}
\mathsf{InteriorObject} &\sqsubseteq \mathsf{Object},\\
\mathsf{ExteriorObject} &\sqsubseteq \mathsf{Object},\\
\mathsf{OfficeObject} &\sqsubseteq \mathsf{Object},\\
\\
\mathsf{Table} &\sqsubseteq \mathsf{InteriorObject},\\
\mathsf{Chair} &\sqsubseteq \mathsf{InteriorObject},\\
\\
\mathsf{Cabinet} &\sqsubseteq \mathsf{InteriorObject} \sqcap \mathsf{OfficeObject},\\
\mathsf{Monitor} &\sqsubseteq \mathsf{InteriorObject} \sqcap \mathsf{OfficeObject},\\
\mathsf{Printer} &\sqsubseteq \mathsf{InteriorObject} \sqcap \mathsf{OfficeObject},\\
\\
\mathsf{Sign} &\sqsubseteq \mathsf{ExteriorObject},\\
\mathsf{Extinguisher} &\sqsubseteq \mathsf{ExteriorObject},\\
\mathsf{Switchbox} &\sqsubseteq \mathsf{ExteriorObject},
\end{aligned}$$

$$\begin{aligned} \mathsf{Door} &\sqsubseteq \mathsf{Object}, \\ \mathsf{Board} &\sqsubseteq \mathsf{Object}. \end{aligned}$$

Note that $B \sqsubseteq A$ implies $P(B|\neg A) = 0$ but it implies nothing about $P(B|A)$ remaining $P(B|A) \in [0,1]$; whenever this happens we adopt a uniform distribution over possible objects.

Composite objects can now be described:

$$\begin{aligned} \mathsf{Desk} &\equiv \mathsf{Table} \sqcap \exists \mathsf{near.Chair}, \\ \mathsf{Entrance} &\equiv \mathsf{Door} \sqcap \exists \mathsf{near.Sign}. \end{aligned}$$

Finally, the whole environment can be described as:

$$\begin{aligned} \mathsf{Room} \equiv\ & \mathsf{Environment} \\ & \sqcap\ \exists \mathsf{contains.Door} \\ & \sqcap\ \exists \mathsf{contains.Table} \\ & \sqcap\ \exists \mathsf{contains.Chair} \\ & \sqcap\ \neg\exists \mathsf{contains.ExteriorObjects} \\ \mathsf{Office} \equiv\ & \mathsf{Room} \\ & \sqcap\ \exists \mathsf{contains.Desk} \\ & \sqcap\ \exists \mathsf{contains.Cabinet} \\ & \sqcap\ \exists \mathsf{contains.Monitor} \\ \mathsf{Classroom} \equiv\ & \mathsf{Room} \\ & \sqcap\ \exists \mathsf{contains.Board} \\ & \sqcap\ \neg\exists \mathsf{contains.OfficeObjects} \\ \mathsf{Hallway} \equiv\ & \mathsf{Environment} \\ & \exists \mathsf{contains.Entrance} \\ & \sqcap\ \exists \mathsf{contains.Extinguisher} \\ & \sqcap\ \exists \mathsf{contains.Switchbox} \\ & \sqcap\ \neg\exists \mathsf{contains.InteriorObjects} \end{aligned}$$

The graph representing this ontology is presented in Figure 2.

The network generated by this ontology is relatively dense due to the presence of quantifiers, and the computational effort needed to calculate probabilities is not trivial. Although exact inference is theoretically possible, there are reliable algorithms for approximate inference that seem more adequate in practice. In particular, Loopy Propagation "breaks" the network connections, replacing local probability distribution with more convenient functions.

3 Semantic Mapping

Mapping larger environments requires dealing with huge amounts of data. Another difficulty in mapping, particularly in 3D mapping, is data association; that

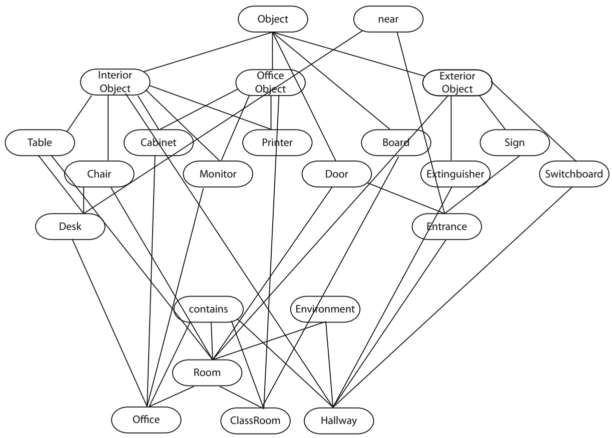

Fig. 2. Graph for the proposed ontology

is, sensor readings from two different positions must be registered into a common coordinate system. It is desirable to separate all the data in various groups, labeled as ceiling, wall, or perhaps kitchen or office.

Our proposed scheme for semantic mapping is an offline process. The robot navigates through the environment, collecting images, laser and odometry data in several positions, and then the data are processed at once. The idea is to identify objects in the images collected by a robot, and to classify the locations in which those images were taken. Then, we use the laser and odometry data related to those images to map smaller areas of the environment. For instance, data of an office let us reconstruct spatially only the office; using odometry information it becomes possible to determine a topology of the environment, and then to unify all metric maps in a single one through the detected doors.

The classification of objects in images is done using a previous set of features, robustly extracted using the SIFT algorithm. The objects present in the ontology have a SIFT descriptor previously computed. Figure 3 shows two templates of objects in the matching with scenes acquired by the robot, and the robot Pioneer 3-AT mounted with a SICK laser to obtain 3d data and a camera to recognize objects in images, used in the experiments.

As the objects in the images are identified, the probabilistic description logic model is used to label sensor data. Inferences are conducted in the probabilistic graph instantiated by the CR$\mathcal{ALC}$ model, thus producing probabilities for the types of environments. Not all sensor are labeled as some images may contain objects not matched against SIFT descriptors. But using the principle of continuity, sensor data not labelled between two or more with identical labels received the same label.

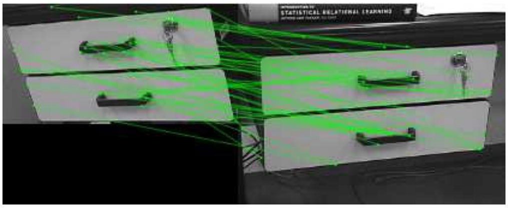

(a) Desk.

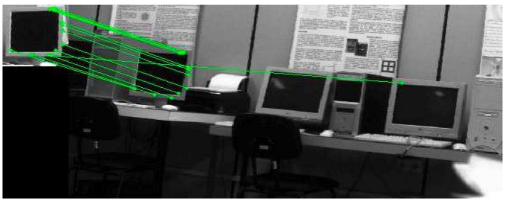

(b) Monitor.

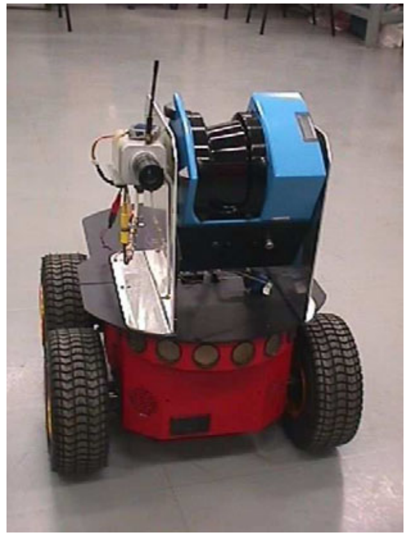

(c) Pioneer AT-3.

Fig. 3. Correspondence between template objects and scenes acquired by a mobile robot

4 Experiments

Our primary goal in doing these experiments is to evaluate the ontology and the reasoning procedure with CR$\mathcal{ALC}$. Both must be appropriate to deal with real robotic problems and sensor data.

We provide some results of an experiment consisting in the robot navigating through three different areas of an indoor environment: two labs and a hallway connecting both. In this experiment, we gathered some images, 3D points from a laser sensor, and the estimated pose in each data gathering location given by odometry and a gyroscope. We determine a priori a set of objects that could be easily identified through SIFT features and used them to characterize different areas.

Following the path of the robot, that goes from a laboratory to the other, passing by the hallway, we picked sequentially 6 points to gather data from the laser and camera, two in each area. Each point was then classified as Office, Classroom or Hallway, accordingly to the identified objects in the images using SIFT; the result of the reasoning can be seen in Table 1. Whenever two sequential points have different labels, the data are split into a new area. In this case, we found 3 different areas. Table 1 shows the inferred values for each point. 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.

Table 1. Identified ambients

Datapoints						
Location	1	2	3	4	5	6
Observations						
Objects	3 chairs 1 table 1 monitor 1 cabinet 1 door	2 chairs 2 tables 1 monitor 1 cabinet 1 door	3 doors 2 signs 1 extinguisher 1 switchbox	2 doors 2 signs 1 extinguisher 1 switchbox 1 board	2 chairs 1 table 2 monitors 1 cabinet 1 door	2 chairs 1 table 1 monitor 2 cabinets 1 door
Inferred Probability for Each Label						
P(Office)	**0.4278**	**0.4258**	0.0000	0.0000	**0.4280**	**0.4280**
P(Classroom)	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
P(Hallway)	0.0000	0.0000	**0.1578**	**0.1619**	0.0000	0.0000
Area	1		2		3	

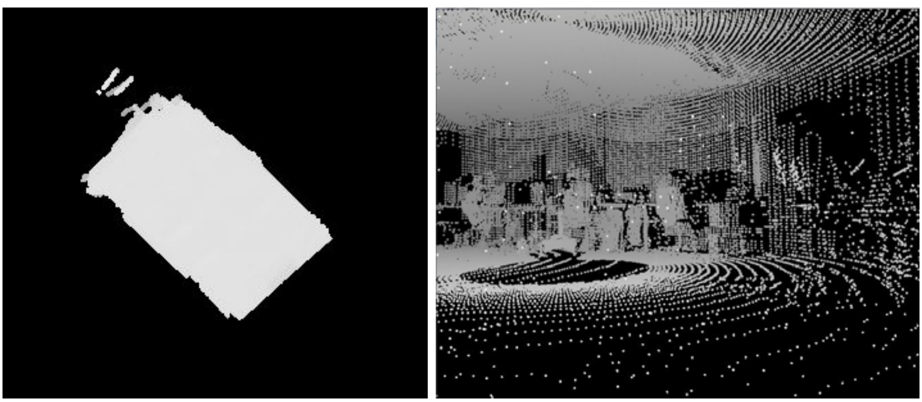

(a) Top-view. (b) View from inside

Fig. 4. An example of a 3d metric map of a given area

Mapping each identified area alone, we could map the 3 environments. Figure 4 shows the 3d map from one of the areas, a laboratory. Figure 4(a) is a top-view from the laboratory and Figure 4(b) is a view from inside the laboratory.

5 Conclusion and Future Work

Semantic knowledge can have a significant impact in robotics. In this paper, we show how to model a robotic mapping task with a probabilistic description logic, and we show how to process sensor data so as to reason about the identity of specific areas in an indoor environment. The main objective of this process is to automatically segment the data acquired by a robot, so that the map could be constructed topologically, thus providing a way of scale the mapping process to larger environments. This objective is said to be achieved, once we managed to automatically split the data in convenient smaller and tagged sets, each one being mapped alone.

Instead of using a SIFT algorithm to provide labels, we might have used an SVM classifier, or perhaps a Conditional Random Field, with better accuracy and generalization. We plan to integrate the image labeling process with the probabilistic description logic in the near future.

Acknowledgments

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

References

1. Thrun, S.: Exploring Artificial Intelligence in the New Millennium. In: Robotic mapping: a survey, pp. 1–36. Morgan Kaufmann Publishers, San Francisco (2003)
2. Agarwal, S., Snavely, N., Simon, I., Seitz, S.M., Szeliski, R.: Building rome in a day. In: IEEE 12th International Conference on Computer Vision, pp. 72–79 (2009)
3. Anguelov, D., Taskar, B., Chatalbashev, V., Koller, D., Gupta, D., Heitz, G., Ng, A.: Discriminative learning of Markov random fields for segmentation of 3d scan data. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, vol. 2, pp. 169–176 (2005)
4. Nuechter, A.: 3D Robotic Mapping. Springer Tracts in Advanced Robotics (2009)
5. Galindo, C., Fernndez-Madrigal, J.-A., Gonzlez, J., Saffiotti, A.: Robot task planning using semantic maps. Robotics and Autonomous Systems 11, 955–966 (2008)
6. Hertzberg, J., Saffiotti, A.: Using semantic knowledge in robotics. Robotics and Autonomous Systems 56, 875–877 (2008)
7. Coradeschi, S., Saffiotti, A.: An introduction to the anchoring problem. Robotics and Autonomous Systems 43, 85–96 (2003)
8. Limketkai, B., Liao, L., Fox, D.: Relational object maps for mobile robots. In: Proceedings of the International Joint Conference on Artificial Intelligence, vol. 1, pp. 1471–1476 (2005)
9. Taskar, B., Abbeel, P., Wong, M.-F., Koller, D.: Relation Markov Networks. In: Introduction to Statistical Relational Learning, pp. 175–199. MIT Press, Cambridge (2007)
10. Wang, J., Domingos, P.: Hybrid Markov logic networks. In: Proceedings of the 23rd National Conference on Artificial Intelligence, vol. 2, pp. 1106–1111 (2008)
11. Posner, I., Schroeter, D., Newman, P.: Using scene similarity for place labeling. In: Proceedings of the 10th International Symposium on Experimental Robotics (2006)
12. Zivkovic, Z., Booij, O., Krse, B.: From images to rooms. Robotics and Autonomous Systems 55, 411–418 (2007)

13. Vasudevan, S., Gchter, S., Nguyen, V., Siegwart, R.: Cognitive maps for mobile robots - an object based approach. Robotics and Autonomous Systems 55, 359–371 (2007)
14. Lowe, D.G.: Distinctive image features from scale-invariant key-points. International Journal of Computer Vision 60, 91–110 (2004)
15. Baader, F., Calvanese, D., McGuinness, D.L., Nardi, D., Patel-Schneider, P.F.: Description Logic Handbook. Cambridge University Press, Cambridge (2002)
16. Cozman, F.G., Polastro, R.B.: Loopy propagation in a probabilistic description logic. In: Greco, S., Lukasiewicz, T. (eds.) SUM 2008. LNCS (LNAI), vol. 5291, pp. 120–133. Springer, Heidelberg (2008)
17. Polastro, R.B., Cozman, F.G.: 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 (2008)
18. Ding, Z., Peng, Y., Pan, R.: BayesOWL: Uncertainty modeling in semantic web ontologies. In: Soft Computing in Ontologies and Semantic Web. Studies in Fuzziness and Soft Computing, vol. 204, pp. 3–29. Springer, Heidelberg (2006)
19. Heinsohn, J.: Probabilistic description logics. In: Conference on Uncertainty in Artificial Intelligence, pp. 311–318 (1994)
20. Jaeger, M.: Probabilistic reasoning in terminological logics. In: Principles of Knowledge Representation (KR), pp. 461–472 (1994)
21. Koller, D., Levy, A.Y., Pfeffer, A.: P-CLASSIC: A tractable probablistic description logic. In: AAAI, pp. 390–397 (1997)
22. Lukasiewicz, T.: Expressive probabilistic description logics. Artificial Intelligence 172, 852–883 (2008)
23. Costa, P.C.G., Laskey, K.B.: PR-OWL: A framework for probabilistic ontologies. In: Conference on Formal Ontology in Information Systems (2006)
24. Sebastiani, F.: A probabilistic terminological logic for modelling information retrieval. In: Croft, W., Rijsbergen, C.V. (eds.) 17th Annual International ACM Conference on Research and Development in Information Retrieval (SIGIR), pp. 122–130. Springer, Dublin (1994)
25. Schmidt-Schauss, M., Smolka, G.: Attributive concept descriptions with complements. Artificial Intelligence 48, 1–26 (1991)