Robotics and Autonomous Systems 56 (2008) 875–877

Contents lists available at ScienceDirect

Robotics and Autonomous Systems

journal homepage: www.elsevier.com/locate/robot

Editorial

Using semantic knowledge in robotics

There is a growing tendency to introduce high-level semantic knowledge into robotic systems and beyond. This tendency is visible in different forms within several areas of robotics. Recent work in mapping and localization tries to extract semantically meaningful structures from sensor data during map building, or to use semantic knowledge in the map building process, or both. A similar trend characterizes the cognitive vision approach to scene understanding. Recent efforts in human–robot interaction try to endow the robot with some understanding of the human meaning of words, gestures and expressions. Ontological knowledge is increasingly being used in distributed systems in order to allow automatic re-configuration in the areas of flexible automation and of ubiquitous robotics. Ontological knowledge was also used recently to improve the inter-operability of robotic components developed for different systems.

While these trends have many questions and issues in common, work on each one of them is often pursued in isolation within a specific area, without being aware of the related achievements in other areas. The aim of this special issue is to collect in a single place a set of advanced, high-quality papers that tackle the problem of using semantic knowledge in robotics in many of its different forms.

The submissions to this special issue made it clear that there are many ways in which semantic knowledge may play a role in robotics. Interestingly, they also revealed that there are many ways in which the term *semantic knowledge* is being interpreted. Before turning to the technical papers, then, it is worth spending a few words on this matter.

1. Semantic knowledge in robotics

A recurring question by the reviewers, and debate among reviewers, was triggered by the request in the Call for Papers that submissions should

> "report novel contributions related to the creation, representation and use of semantic knowledge in autonomous robots".

The question in many cases was "What in this paper is supposed to be semantic knowledge?".

The question of what should count as semantic knowledge has been asked repeatedly in the history of philosophy. In more recent times, it has been the subject of much debate in the Artificial Intelligence (AI) community, especially in the area of knowledge representation and reasoning (KR), and it is now central to the development of the semantic web. As general sources for these topics, we refer to the textbooks by Russell and Norvig [1], by Brachman and Levesque [2], and by Antoniou and van Harmelen [3].

Interestingly, many of the views on semantic knowledge that have been adopted (explicitly or implicitly) by the authors and by the reviewers of submitted papers had appeared before in this debate. To make a long story short, let us emphasize two important points about semantic knowledge, which have long been discussed in KR, and which are central in the present context: (1) the need for an *explicit representation* of knowledge inside the robot; and (2) the need for *grounding the symbols* used in this representation in real physical objects, parameters, and events.

1.1. Explicit representation

In knowledge-based systems, the term *semantic knowledge* usually denotes descriptions of the concepts and relations used to define the domain of interest. The main hypothesis in the field of KR is that these descriptions are represented *explicitly* inside the system, that is, in a way that allows the system to access and manipulate them. Employing pieces of knowledge in a robot, then, requires that each and every concept be linked to other concepts by declarative statements in terms of the domain theory. The domain theory, together with some inference engine, can be used for inferring new facts about the domain from known facts. The archetype of a knowledge base in KR is a theory in first-order logic, but many other formalisms can be used, including probabilistic representations (e.g., Bayes networks), constraint networks, Description Logics, or planning-related representations. In fact, what formalism to use usually depends on the specific domain: in KR terminology, a formalism has to be representationally and inferentially efficient to be useful for a task.

Irrespective of the specific formalism used to represent semantic knowledge, the point to keep in mind is: we have no semantic stance yet, if we simply attach labels (like `chair`, or `kitchen`, or an event like `scoring a goal`) to a set of sensor data, like a region in a camera image. A robot that is able to represent and employ semantic knowledge must certainly be able to do object recognition or classification in some sense, i.e., to attach labels to compact subsets of its sensor data, in order to get a foothold for some type of interpretation of its environment. However, these labels have to be embedded in a domain theory in order to allow the robot to do reasoning. To become meaningful categories, they need to co-exist with other categories in some form of ontology, be it represented as it may. "Bare" labels unconnected to an ontology may induce meaning to a human observer, but this meaning is owed to this observer's understanding of the domain, not the robot's: the robot cannot use these labels as long as they are purely syntactic. Ancient AI folklore had the nickname of *wishful mnemonics* for such labels that bore meaning for human observers, and only for them.

0921-8890/$ – see front matter
doi:10.1016/j.robot.2008.08.002

1.2. Symbol grounding

Using a semantic stance *in robotics* involves a specific challenge: All elements used in the knowledge representation have to be effectively *grounded* in the robot's sensor and motor signals. So, KR for robots has to add to "regular KR" the means to recognize objects and/or assess the truth or falsity of propositions from the sensor data stream of the given sensor configuration in the given environment. In most practical cases, the robot should be able to do this in real time.

This is, of course, an enormous requirement — some may even say it is intimidating. It is closely related to the symbol grounding problem, [4] which is generally perceived as fundamental. Exactly how fundamental it is, is still being debated; in any case, robotics should not wait until the problem is solved, before turning to using semantic knowledge. The good news is: grounding facts and concepts in sensor data on-line and on-board the robot appears to be doable in some cases — this Special Issue contains a number of examples. The bad news is: Although KR has a healthy number of representation and reasoning formalisms on stock that are ready to be employed on systems that can rely on their human users for having their symbols grounded, matters are different for knowledge-based robots. So, in choosing a KR formalism for a robot supposed to use semantic knowledge, one must also consider the need to provide the grounding for all represented concepts. There may be representation formalisms that are better suited than others in this context.

2. In this issue

Finally, an answer needs to be given to the obvious question: If using semantic knowledge on a robot is difficult, why should we be working on it in the first place? The collection of papers contained in this Special Issue provides a constructive answer to this question by example. The papers tackle several facets of the problem of endowing a robot with the capability of acquiring and using semantic knowledge. They can be seen along three aspects, which typically co-exist in a knowledge-based robot.

The first aspect, namely, *acquiring semantic knowledge from sensor data* on board a mobile robot, is prominent in the first five papers in this issue. Environment sensor data considered here consist of 2D or 3D laser data, camera images, or both; categories for recognition are handcrafted, or learned as part of the process; these categories include objects of limited sizes and fixed forms, like a cupboard, as well as environment structures of unbounded sizes and arbitrary forms, like a pathway.

To start with the papers featuring learning, Modayil & Kuipers take the most basic approach by employing unsupervised learning on 2D laser scans for generating object representations. D'Este & Sammut describe a method for ILP-based learning of concepts from camera data in interaction with a human teacher. Posner et al. apply supervised learning for classifying objects and structures in large scenes available in 2D/3D laser and camera data.

Yet, data recorded by a robot interacting with its environment contain more than just a set of distal readings of environment texture or geometry, which hit the robot like rainfall — the interaction itself is part of the sensing, and of the learning in the case of the three papers just mentioned. So, spatio-temporal information about robot pose, and teacher interaction (D'Este & Sammut) is in the background of the learning process here. Therefore, the objects or structures that get recognized by the trained classifiers are naturally located in space. Posner & al. use that for generating geometrical or spatial scene descriptions, or "semantic robot maps", that involve the detected categories.

The remaining two papers under the first aspect focus exactly on this: using sensor data interpretation routines, which may be trained or canned, for building semantic maps. Nüchter & Hertzberg elaborate on the processing pipeline for getting from individual 3D laser scans to a consistent geometry-plus-categorical environment representation; they also use some feedback from semantic map contents to lower-level sensor data processing. Rusu et al. start from 3D laser scan data using similar low-level methods, but later use for processing them the prior knowledge that their robot operates in a world of kitchen objects (cupboard, dishwasher, etc.); part of that knowledge is available in the explicit form of a Description Logic ontology.

The second aspect concerns the ability to use semantic knowledge for *improving planning and control aspects in the robot itself*. It is present in most papers in this issue, but it is central to the following five.

Bouguerra et al. use semantic knowledge, which comes in a Description Logic and a probabilistic domain model, for monitoring the execution of explicitly represented plan actions. Both explicit and inferred expectations about action results are used for verifying by observation successful action execution. The paper by Galindo et al. emphasizes the plan generation aspect prior to execution, by basing it on a semantic map as mentioned above. This allows spatial and propositional knowledge about the environment to be used efficiently in combination for planning. Stulp & Beetz focus on plan optimization in the context of plan generation and execution. Plan generation is based on a declarative planning domain description language (PDDL, to be precise) and an off-the-shelf planner; optimization applies learned prediction of effects for generated actions, based on an execution model.

Ferrein & Lakemeyer apply planning and reasoning by means of a different representation framework than the other papers, namely, the GOLOG dialect READYLOG. This includes plan generation, execution and monitoring, where the paper emphasizes handling dynamic domains in this context. Calisi et al. take a representation rather than a plan generation approach for specifying robot control and improving performance, centering around the notion of context. Contexts are explicitly represented and manipulated, resulting in a context-based architecture that features rule-based control.

The third aspect is the *usage of semantic knowledge in robotics motivated in particular tasks*. Most of the papers in this collection focus on methods and concepts around semantic knowledge in robotics rather than on applications, so this aspect is not prominent here. However, it is important in the last one by Holzapfel et al., which deals with the problem of learning semantic knowledge in interaction with a human. The approach assumes that a robot (a humanoid one in this case) that is supposed to cooperate with humans needs to describe its environment in the same terms like the human does, so effective human–robot interaction (HRI) requires parallel or similar categories for perceiving and describing the common world. The paper by D'Este & Sammut in this collection shares this view; the one by Rusu et al. with its aim of modeling human-designed environments in human categories also has a strong flavor of it. A number of projects aiming at HRI reported in the literature would agree.

Based on the state of the art in using semantic knowledge in robotics, as presented in this special issue, it is still quite a way to go, until a robot ontology can possibly resemble a human one, to be in tune with human communication based on speech or gestures. Yet the first steps are being taken.

Acknowledgements

The collection of papers issued here has a precursor. The IEEE International Conference Robotics and Automation (ICRA'07) in April 2007 included a workshop on *Semantic Information in Robotics* with largely the same aim in mind. The response to the workshop made it obvious that it was time to put this topic on the robotic agenda.

The call for the present Special Issue in Fall 2007 yielded 22 submissions, of which the present 11 were accepted after a two-stage reviewing process. This process used the services of a guest editorial board, whose members are:

Michael Beetz	Amy Loutfi	Andreas Nüchter
Wolfram Burgard	Patric Jensfelt	Claude Sammut
Tom Duckett	Geert-Jan Kruijff	Mary-Anne Williams
Cipriano Galindo	Daniele Nardi	Jeremy L. Wyatt

We received help from additional reviewers; namely, Luca Iocchi, Lars Karlsson, Kai Lingemann, and Danijel Skocaj. The dedicated work of these colleagues was pivotal in selecting the best submissions and in interacting with the authors to improve their submissions even further. Our sincere thanks to all of them.

References

[1] S. Russell, P. Norvig, Artificial Intelligence: A Modern Approach, 2nd edition, Prentice Hall, Englewood Cliffs, NJ, 2003.
[2] R. Brachman, H. Levesque, Knowledge Representation and Reasoning, Morgan Kaufmann, San Francisco, CA, 2004.
[3] G. Antoniou, F. van Harmelen, A Semantic Web Primer, MIT Press, Cambridge, MA, 2004.
[4] S. Harnad, The symbol grounding problem, Physica D 42 (1990) 335–346.

Joachim Hertzberg*
University of Osnabrück,
Inst. of Computer Science,
Knowledge-Based Systems Research Group,
D-49069 Osnabrück, Germany
E-mail address: joachim.hertzberg@uos.de.

Alessandro Saffiotti
Örebro University,
Department of Technology,
AASS Mobile Robotics Lab,
S-70182 Örebro, Sweden

* Corresponding editor.