Using Semantic Information for Improving Efficiency of Robot Task Planning

C. Galindo, J.-A. Fernández-Madrigal, and J. González
System Engineering and Automation Department
University of Málaga, (Spain)
{cipriano,jafma,gonzalez}@ctima.uma.es

A. Saffiotti
Center for Applied Autonomous Sensor Systems
Dept. of Technology, Örebro University
S-70182 Örebro, Sweden
{alessandro.saffiotti}@aass.oru.se

Abstract— **The use of semantic information in robotics is an emergent field of research. As a supplement to other types of information, like geometrical or topological, semantics can improve mobile robot reasoning or knowledge inference, and can also facilitate human-robot communication. These abilities are particularly relevant for robots intended to operate in everyday environments populated by people, which typically involve a great number of objects, places, and possible actions. In this paper, we explore a novel usage of semantic information: as an improvement for task planning in complex scenarios like the mentioned ones, where other planners easily find intractable situations. More specifically, we propose to first construct a "semantic" plan composed of categories of objects, places, etc. that solves a "generalized" version of the requested task, and then to use that plan for discarding irrelevant information in the definitive planning carried out on the symbolic instances of those elements (that correspond to physical elements of the world with which the robot can operate). Our results using this approach are promising, and have been compared to other existing approaches.**

Keywords- Semantic information, Abstraction, Generalization, Task Planning.

I. INTRODUCTION

The construction and study of representations of the real world is essential for the proper performance of mobile robots. The majority of approaches represent space and objects by only considering geometric information, for example building spatial maps (flat or hierarchical) with free and occupied areas [21],[17]. Although *geometrical* data is sufficient for solving a variety of robot tasks, other types of information are also useful: *topological* (to deal efficiently with large-scale maps), *hierarchical* (to deal efficiently with large amounts of information), and *semantic* (for performing more intelligently – from the point of view of a human being). In particular, semantic information could enable the robot to reason about the functionalities and characteristics of objects and environments [12], while topological symbols permit to communicate with humans using a proper set of terms and concepts [10].

The need to include semantic information in robot maps has been recognized for a long time [16],[5], but the integration of such information within spatial representations is still an emergent trend. Actually, most robots that incorporate provisions for task planning and/or for communicating with humans store implicitly some semantic information in their maps (e.g. [1],[20]), for example human-inspired classification of spaces (rooms, corridors, halls) or names of places and objects and the relations among them. However, these implicit approaches depend on the ability of the designers to capture the suitable set of semantic constraints and mechanisms.

In our previous work we have started to explore several ways in which the explicit use of semantic information may extend the robot capabilities:

- Semantic information enables a robot to infer new knowledge from its environment (e.g., to infer the type of a room according to the objects which are inside [12]).
- The use of both semantic and spatial information enhances human-robot communication by using concepts, terms, and reasoning understandable by people [10],.
- A robot can avoid misleading sensor readings by using semantic information. For instance, if the robot certainly knows that it is at a kitchen but its vision system detects a bathtub, it could discard this information since it knows that kitchens do not contain that kind of objects [12].

In this paper, we take one further step on using semantic information by exploring its benefits on improving symbolic task planning. We claim, as commented further on, that semantic information, when related appropriately to spatial information, can be exploited to help a robot to plan efficiently within large and/or complex worlds. In fact, some intractable problems under other planning approaches can become more tractable through the semantic support we proposed in this work.

In the robotics arena, though a great attention has been paid to path planning, task planning efficiency has been usually pushed to the background, most probably because of the simplicity of the scenarios considered so far in the field, or because of the limited forms of interaction between the robot and its environment.

This work has been supported by the Spanish Government under research contract DPI2005-01319, and by the Swedish KK Foundation.

ICRA-07 Workshop on Semantic Information in Robotics
IEEE Int Conf on Robotics and Automation. Rome, Italy, April 2007

However in the AI community, task planning efficiency has been largely studied [19],[22],[23],[24]. Some of their approaches have been borrowed in robotic applications, but up to our knowledge, only a few works have explicitly addressed the problem of planning efficiency for mobile robots [11],[9]. They usually rely on hierarchical structures upon topological representations to rid out irrelevant information to the task to be planned, achieving a significant speedup of the planning process.

Here we propose a novel hierarchical planning approach (called *SHPWA*) that exploits the semantic information managed by a mobile robot to improve its spatial-based task planning. Our semantic approach can be classified as a description logic system that organizes particular instances or world elements, say "my cup of coffee", into general categories, i.e. "Cup". This type of knowledge generalization produces a *taxonomic hierarchy* [14] (that we called here the *semantic hierarchy*).

The semantic hierarchy can provide the planning process (which is intended to construct a plan to solve the requested task using only spatial information) with information about the categories of world elements involved by the task at hand. Thus, we are able to discard irrelevant information for the spatial planning before it executes, reducing the computational effort. For example, if a servant robot is at a kitchen and it is commanded to "take my favorite fork" (let say *fork-1*), a possible solution could be "approach *shelf-1* and take *fork-1*". However, assuming a realistic environment, a kitchen may contain hundreds of objects and tens of distinctive places for navigation, which could prevent the planning process from achieving a successful result. In this paper, we claim that semantic information can help the planning process to reduce the search space: in our example, a fridge should not be considered for the task if the robot knows that forks are always in drawers; similarly, any spoon or knife could also be ignored although they can be found in a drawer. Our experiments show that using semantic information in this manner can help the robot to scale-up to environments containing thousands of objects.

In the rest of the paper, we describe our multi-hierarchical and semantic representation of the world (section II) and its use for improving the robot task planning process (section III). Section IV gives a comparison between our semantic-based hierarchical planner (SHPWA) and other non-semantic approaches (the flat planner Metric-FF [15], and the hierarchical planner HPWA [11]). Finally, some conclusions and future works are outlined.

II. SYMBOLIC REPRESENTATION OF THE WORLD

In our approach, we consider a mobile robot with a multi-hierarchical symbolic representation of its workspace [12]. This representation entails two hierarchies that represent the robot environment from two different perspectives: (i) a *spatial perspective*, that enables the robot to reliably plan and execute its tasks on existing objects, places, etc. (e.g., navigation, manipulation); and (ii) a *semantic perspective*, that provides it with inference capabilities (e.g., a *bedroom* is a *room* that contains a *bed*).

The spatial hierarchy contains symbols that represent particular elements of the environment, either perceived by the robot sensors or not. The ones that cannot be perceived are groupings of more detailed symbols and are useful for improving computational efficiency. The ones that can be sensed are created by means of the *anchoring technique* [11] that connects sensor data that refer to physical elements of the world, e.g. an image of my favorite fork, to particular symbols in the model, i.e. *fork-1*. This connection is represented by a data structure called *anchor* that includes a set of properties useful to re-identify the object, e.g., its color and position.

Figure 1 depicts our multi-hierarchical world representation, which includes the spatial and the semantic hierarchies. The *Spatial Hierarchy* arranges symbols in different hierarchical levels through abstraction of detail: (i) simple objects and distinctive places for navigation, (ii) the topology of the robot environment, and (iii) the whole environment represented by an abstract node. Additional intermediate levels could also be included.

The *Semantic Hierarchy* contains categories of the spatial symbols that represent particular elements of the world, i.e. *cup-1* in the spatial hierarchy is an instance of the category *Cup* in the semantic hierarchy. This categorization, represented in figure 1 as dotted lines, can be constructed by identifying particular properties of the correspondent anchors of the instances (i.e. a symbol is categorized as a cup if its anchor contains a perceptual image with a given size and shape), and/or through semantic inference, as presented in [12]. The semantic hierarchy may also model relations between categories, representing semantic knowledge like for instance that *Cups* "are usually on" *Tables*.

We manage this multi-hierarchical structure by using a mathematical model based on graphs called *MAH-graph* model [8], which has proved its suitability in reducing the computational effort of robot operations such as path-search [9] or symbolic task planning [11]. However, the semantic hierarchy can also be modeled by employing standard AI languages, like the *NeoClassic* language [18], in order to provide the robot with inference capabilities.

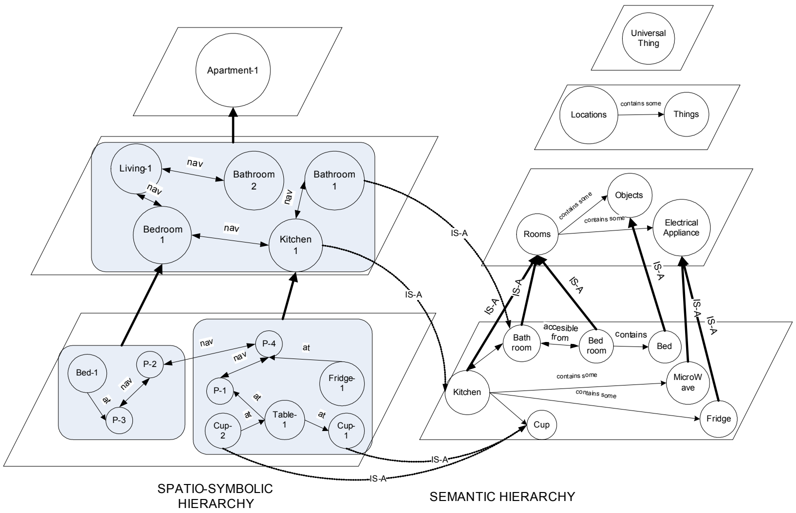

Figure 1. An example of the spatial and semantic hierarchies. On the left, spatial information represents the robot environment at different levels of detail. Shadowed regions indicate the set of symbols abstracted to the next upper hierarchical level. On the right, semantic information that models categories of spatial symbols and relations between them. This categorization starts in the links from symbols of the spatial hierarchy to categories in the semantic hierarchy (dotted lines). For clarity sake not all links and connections are shown in the figure, nor is anchoring.

A. The Spatial Hierarchy

The Spatial-Symbolic Hierarchy contains spatial and metric information from the robot environment. This model is based on *abstraction of detail*, that helps to minimize the information required to plan tasks by grouping/abstracting symbols in complex and high-detailed environments.

Symbols of this hierarchy are created by anchoring. On the one hand, laser-based gridmaps are anchored to symbols that represent distinctive places for navigation and open spaces [7], [3], on the other hand a visual pattern recognition system is employed to include information about the objects perceived by a camera mounted on the robot. When an object is detected that matches certain properties, i.e. a particular colour or shape, a new symbol is added to the model[1].

Symbols in both hierarchies are represented by vertexes which are interconnected through edges that represent relations between them. For instance, in the spatial hierarchy, two symbols that represent locations can be connected through a *navigability* edge, while a *located-at* edge models the relation between objects and a location (see figure 1).

Vertexes from the spatial hierarchy are abstracted into upper levels of the hierarchy to represent the topology of the environment, that is, a more general and less detailed representation of it. Different criteria can be adopted in order to construct these upper levels. We consider here grouping symbols according to their geometrical position and the normal distribution of human-like environments (places-rooms-apartments-buildings, etc.), though other techniques can be employed to construct hierarchies meeting other requirements, for instance for constructing hierarchies that improve the task planning process [9],[13].

B. The Semantic Hierarchy

The Semantic Hierarchy models semantic knowledge about the robot environment. All categories in this hierarchy are "refinements" of a common ancestor called *Universal Thing*, at the top level. Different sub-categorizations can be developed until reaching the lowest level that contains the most specific categories (kitchen, bedroom, cup, fridge, fork, etc.).

In our work, we incorporate semantic information within the multi-hierarchical model through semantic networks, as the one depicted in figure 2, although other mechanisms can be considered, like the *NeoClassic* system for knowledge representation and reasoning [18]. Regardless of the considered representation, the model should permit us to represent constraints or relations between categories, like for example the fact that a kitchen must have at least one fridge. The following is an example of how the category "kitchen" can be defined in the NeoClassic language. Intuitively, a kitchen is a room that has a stove, a fridge and a dish-washer, but does not have a bed, a bathtub, a sofa or TV set:

```
(createConcept Kitchen
(and Room
(atLeast 1 stove)(atLeast 1 fridge)
(atLeast 1 dish-washer)(atLeast 1 kitchen-furniture)
(and (atMost 0 bathtub) (atMost 0 sofa)
(atMost 0 bed) (atMost 0 tvset))))
```

The semantic network depicted in figure 2 shows the lowest level of the semantic hierarchy considered in our experiments.

[1] Since image recognition is out of the scope of our work we focus on recognizing simple objects, i.e., boxes, based on a particular shape and colour [12].

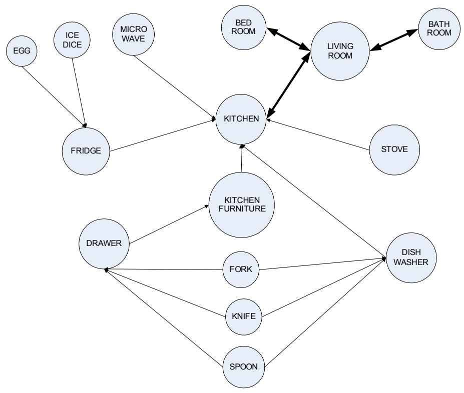

Figure 2. Example of the lowest level of a semantic hierarchy for an apartment scenario. Thick arcs indicate “can be accessed from”, while thin ones indicate “can be found in”.

As it will shown in the next section, a planning process able to cope with this semantic information could decide in advance the categories whose instances may be necessary for solving a given task, discarding the rest. Apart from the use of this information to guide task planning, it could also allow a mobile robot to perform some other types of inference, as described in [12]. For instance, if we know that “room-1” is a room that contains “obj-1”, which is a fridge, then we can infer that “room-1” is a kitchen.

III. THE PLANNING PROCESS

Our semantic planning process is based on a previous work on hierarchical planning called HPWA (*Hierarchical Planning through World Abstraction* [11]). HPWA runs a given planner (the so-called *embedded planner*), like for instance Metric-FF [15], using the information stored at different levels of the spatial hierarchy to improve planning efficiency in large and complex domains – including those that are not large-scale but contain a high number of objects and/or relations).

Broadly speaking, HPWA receives as input a specification of the goal to achieve specified using symbols of the spatial hierarchy that belong to the same hierarchical level (typically the ground level). First, HPWA abstracts the goal to the upper levels, until the goal loses its meaning or becomes trivial (please consult the details in [11]). Then, it solves the requested task (the goal) specified at the highest level of abstraction of detail that has been reached. That abstract plan is used to discard irrelevant spatial information at the next lower level of the hierarchy (more detailed) by discarding the elements that are not involved in the abstract plan plus all the elements of lower hierarchical levels that abstract to them. This process is repeated until the level of the hierarchy where the requested task was originally specified is reached, providing a plan made of simple actions that the robot can carry out if all the elements of the task are appropriately anchored. In a typical real-world environment, which may contain hundreds of objects and distinctive places for navigation, the reduction in computational cost achieved by HPWA can be very important, as demonstrated in [11].

In this paper, we improve the computational efficiency of HPWA by using semantic information. The new approach is called SHPWA (for *Semantic HPWA*) and requires the use of a multi-hierarchical model like the one described in section II: the spatial hierarchy will serve for planning tasks as described before (tasks that contain real, spatial elements of the world on which the robot can operate), while the semantic hierarchy will provide support for planning categorical tasks (those that contain semantic symbols, that is, categories of objects, places, etc.). Broadly, the SHPWA process consists of two executions of the HPWA, as follows:

1) A task is requested for planning, for instance “take my favorite fork”. This task is specified as a goal to achieve (a state of the world that have to be reached) using only symbols, i.e. *fork-1*, from some hierarchical level of the spatial hierarchy. Typically, that level is the ground level of the hierarchy and therefore all the symbols are anchored.

2) All the symbols in the goal are translated through the semantic links (“is-a”) that connect both hierarchies to categories in the semantic hierarchy. Thus, the task “take my favorite fork” is translated to “take FORK”. For the sake of simplicity we assume that all of the spatial symbols can be translated in this way; if this operation cannot be done, we pass directly to step 5 and thus semantics has not helped planning.

3) HPWA is used in the semantic hierarchy for constructing a categorical plan (one that only contains categories of objects, places, etc.) that solves the requested task. For the “take FORK” goal, considering the semantic network depicted in figure 2 and assuming that the robot is at the living-room, the categorical plan would be: “Go from LIVING-ROOM to KITCHEN”, “Go to KITCHEN-FURNITURE”, “Open DRAWER” and “Take FORK”. We assume that if a plan exists that solves the task, there will exist a corresponding categorical plan in the semantic hierarchy. If that is not true, we go to step 5 becoming a non-semantic task planning approach.

4) The set of all the semantic categories involved in the categorical plan indicates the set of particular spatial symbols (places, objects, etc.) that will be involved in the final plan that is to be constructed on the spatial hierarchy. Therefore, all the symbols in that hierarchy that do not correspond to categories involved in the categorical plan are discarded for the subsequent planning process. In our example, only distinctive places of the kitchen, and the living-room, drawers of the kitchen’s furniture, and forks are considered discarding the rest of elements of the domain.

5) HPWA is used in the spatial hierarchy for constructing the final plan considering only instances of the relevant categories provided by the semantic plan. In our example, the final plan would be: "go from living-1 to living-2", "go from living-2 to living-3", "go from "living-3 to kitchen-1", "go from kitchen-1 to furniture-1", "open drawer-3", "unstack fork-3", "unstack fork-2", "take fork-1", which can be executed by the robot.

Step 1 assumes that all the symbols in a requested task have "is-a" links to the semantic hierarchy. This implies that the semantic hierarchy must be complete for all the elements of the world on which the robot will need to plan tasks, a reasonable assumption specially when the anchoring process, that is in charge of maintaining connections between real world elements and their spatial symbols usually provides information for the generation of their semantic categories [6], [12].

Note that step 3 does not involve a high computational effort, since although spatial information can grow rapidly through the exploration of the world, semantic one (the categories) remains generally bounded. Also notice that to carry out planning on the semantic hierarchy we need that relations between categories different from "is-a" (that is, those that do not serve to generalize) correspond to operational needs of the task. For example, if the task requested is "Put the *fork-1* in *drawer-3*", the corresponding goal will be translated to the semantic hierarchy into "A fork is in a drawer" and then planning will be performed on that goal. Thus, for carrying out that planning, the semantic hierarchy must contain relations between categories such as "forks are usually in drawers" (to find the fork), "drawers belongs to a piece of furniture" (to find the drawer), "kitchens and dining-rooms contains pieces of furniture", "kitchens are accessable from dining-rooms – and vice versa" (to take the fork to the kitchen if it is at the dining-room), etc. Further research must be conducted on the appropriate construction of these semantic hierarchies, which can be done manually (by a human programmer), automatically (by analyzing the spatial relations between objects and places to produce semantic information), or through a combination of both. In the experiments of this paper we have chosen the manual construction for simplicity.

IV. EXPERIMENTS AND RESULTS

We have performed a number of experiments intended to verify the following hypothesis: using semantic information for task planning can reduce the complexity of the process, and therefore it allows us to scale up to larger domains than non-semantic planners. To test this hypothesis, we have used the SHPWA planner described in the previous section, and compared it to (non-semantic) HPWA. We also used a non-hierarchical planner (Metric-FF) for a baseline comparison. We have used a representation of a scenario with four rooms, containing many objects (grouped on piles) and distinctive places for robot navigation (see figure 3a). We have found that planning in such a scenario without semantic information may turn intractable even simple tasks.

We have run several experiments in which the number of objects and places have been gradually increased from 100 to 5000. For simplicity, in these experiments we have considered hierarchies with only two levels, though using additional hierarchical levels could improve even more planning efficiency [11]. It should be emphasized that the planned tasks were not actually executed on a robot: the goal of our experiments was to show the efficiency gained in the planning process by using a semantic-based planner, so executing the task was out of scope.

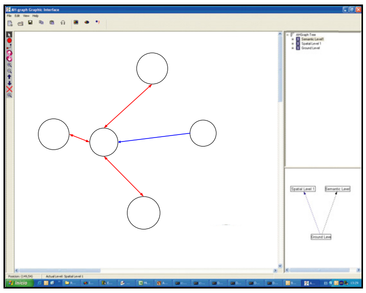

a)

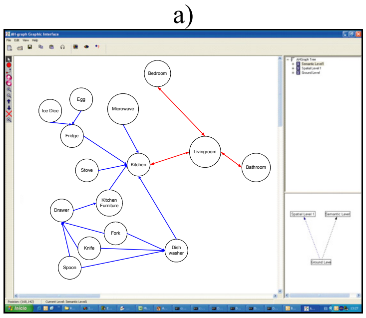

b)

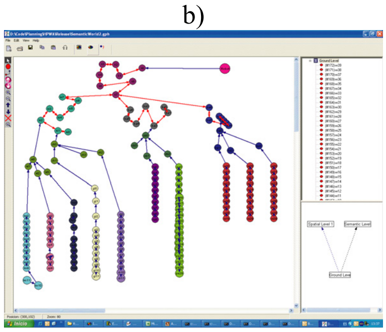

c)

Figure 3 Multi-hierarchical representation of a typical apartment scenario. a) First level of the spatial hierarchy that represents the different rooms and their connections. b) First level of the semantic hierarchy in which all the categories considered for this experiment and the relations captured by relations of particular instances are shown ("is accessable from" and "can be found at"). c) The spatial ground level for the simulated environment, with 100 world elements.

Figure 3 depicts the hierarchical model used in our experiments that represents an apartment. Figure 3a shows the first level of the *Spatial Hierarchy* in which topological information is grouped into rooms, and figure 3b shows the first level of the *Semantic Hierarchy* that categorizes the information of the ground level, which is shown in figure 3c.

In spite of the simplicity of the considered scenario, planning results using our semantic planning approach are promising. Several experiments have been conducted considering a random number of objects, places, and five different tasks (also chosen at random) consisting of taking an object (possibly manipulating other objects to take the desired one) and carrying it to a given location, i.e., "take *fork-1* to the *living- room-1*".

Figure 4 shows the average planning time for the set of random tasks varying the complexity of the environment (number of elements) with Metric-FF (planning only at the ground spatial level), HPWA (using all the levels of the spatial hierarchy, no semantics), and SHPWA. Although the behavior of each of the three planning approaches follows an exponential trend, the chart of figure 4 clearly demonstrates the benefits of using semantic information for planning. Also notice that in all cases the time of SHPWA is the shortest one, which proves that it actually alleviates the combinatorial explosion of the search involved in planning by discarding unnecessary objects for the task.

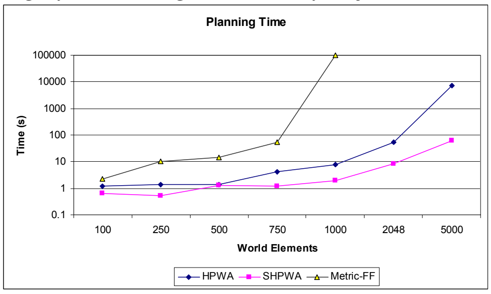

Figure 4. Task Planning comparison. Average time for planning five random tasks within random variations of the complexity of the considered simulated environment. Note the remarkable improvement achieved by SHPWA with respect to other non-semantic planners.

V. Conclusions and Future Work

This paper has explored a novel way of using semantic information in mobile robot applications: improving robot task planning through semantics. The proposed approach envisages a multi-hierarchical model that coherently entails different sources of environmental information for robot operation. Within that model, one hierarchy (the spatial one) represents spatial symbols, some of them anchored to geometric information, while another hierarchy (the semantic one) represents semantic information arising from relations between particular instances of objects, places, etc. The information provided by semantic information is exploited for constructing a categorical plan to satisfy a given goal, which will serve to discard particular instances not involved in the solution, improving thus the overall task planning process. Planning experiences have demonstrated the benefits of using semantic information for planning tasks within complex scenarios in which a relative low number of objects makes other planners fail.

In the future, we plan to explore the human participation to help the robot to acquire more complex semantic information, the use of automatic procedures for that purpose, and to perform tests in which semantic-based planning is used in a physical robotic platform.

References

[1] B.Bakker, Z.Zivkovic, and B.Krose. Hierarchical Dynamic Programming for Robot Path Planning. IROS 2005, pp.2756-2761.

[2] M. Beetz, T. Arbuckle, T. Belker, A.B. Cremers, D. Schulz, M. Bennewitz, W. Burgard, D. H¨ahnel, D. Fox, and H. Grosskreutz. Integrated Plan-Based Control of Autonomous Robots in Human Environments. *IEEE Intelligent Systems* 16(5):56–65, 2001.

[3] P. Buschka and A. Saffiotti. A Virtual Sensor for Room Detection. *IROS 2002*, Lausanne, CH, pp. 637-642.

[4] P. Buschka and A. Saffiotti. Some Notes on the Use of Hybrid Maps for Mobile Robots. *8th Int. Conf. on Intelligent Autonomous Systems* (IAS) 547–556.Amsterdam, 2004.

[5] R. Chatila and J.P. Laumond. Position referencing and consistent world modeling for mobile robots. In *Proc of the IEEE Int Conf on Robotics and Automation*, pages 138–145, 1985.

[6] S. Coradeschi and A. Saffiotti. An Introduction to the Anchoring Problem. *Robotics and Autonomous Systems* 43(2):85-96, 2003.

[7] E. Fabrizi and A. Saffiotti. Extracting Topology-Based Maps from Gridmaps. In *Proc. of IEEE Int. Conf. on Robotics and Automation(ICRA)*, San Francisco, CA, pp. 2972–2978, 2000.

[8] J.A. Fernández-Madrigal and J. González, Multi-Hierarchical Representation of Large-Scale Space. Series on Micro-processor-based and Intell. Systems Eng. 24, Kluwer Academic Pub., 2001.

[9] J.A. Fernandez and J. Gonzalez. Multihierarchical Graph Search. *IEEE T. on PAMI* 24(1):103-113, 2002.

[10] C. Galindo and J. González and J.A. Fernández-Madrigal. A Control Architecture for Human-Robot Integration: Application to a Robotic Wheelchair. IEEE Transactions on Systems, Man, and Cybernetics Part B, vol 36, nº 5, pp.1053-1068, 2006.

[11] C. Galindo, J.A. Fernandez-Madrigal, and J. Gonzalez, Hierarchical task planning through world abstraction, IEEE Trans. on Robotics, 20 (2004), pp. 667–690.

[12] C. Galindo, A. Saffiotti, S. Coradeschi, P. Buschka, J.A. Fernandez-Madrigal, and J. Gonzalez. Multi-Hierarchical Semantic Maps for Mobile Robotics. *IROS 2005,* pp. 3492-3497.

[13] C. Galindo. A Multi-Hierarchical Symbolic Model of the Environment for Improving Mobile Robot Operation. PhD Thesis, University of Malaga (Spain), 2006.

[14] Y. Gil. Description Logics and Planning. AI Magazine, 2005.

[15] Hoffmann J. and Nebel B., *The FF Planning System: Fast Plan Generation through Heuristic Search*. Journal of Artificial Intelligence Research, **14**, 2001 pp. 253-302.

[16] B. Kuipers.Modeling spatial knowledge.*Cognitive Science*,2, 1978.

[17] B. Lisiend, D. Morales, D. Silver, G. Cantor, I. Rekleitis and H. Choset. The Hierarchical Atlas. *IEEE Trans. On Robotics*, vol. 21, no.3, pp. 473-481, 2005.

[18] P.F. Patel-Schneider, M. Abrtahams, L. Alperin, D. McGuinness and A. Borgida. NeoClassic Reference Manual: Version 1.0. AT&T Labs Research, Art. Intell. Principles Res. Dept., 1996.

[19] E.D. Sacerdoti. Planning in a Hierarchy of Abstraction Spaces. Artificial Intelligence, **5**(2), 1974, pp. 115-135.

[20] S. Thrun, M. Bennewitz, W. Burgard, A.B. Cremers, F. Dellaert, D. Fox, D. Hähnel, C. Rosenberg, N. Roy, J. Schulte, and D. Schulz. MINERVA: A 2nd generation mobile tour-guide robot. ICRA 1999, pp. 1999-2005.

[21] S. Thrun. Robotic Mapping: *A Survey. Exploring Artificial Intelligence in the New Millennium, Morgan Kaufmann*, 2002.

[22] M. Ghallab, D. Nau, and P. Traverso. Automated Planning. Morgan Kaufmann Publishers, May, 2004.

[23] J. Rintanen, J. Hoffmann. An Overview of Recent Algorithms for AI Planning. Künstliche Intelligenz, **2**(1), 2001, pp. 5-11.

[24] A. Botea, M. Enzenberger, M. Müller, and J. Schaeffer. Macro-FF: Improving AI Planning with Automatically Learned Macro-Operators. J. Artf. Intell. Research 24 (2005). pp. 581-621.