Robotics and Autonomous Systems 61 (2013) 1131–1143

Contents lists available at ScienceDirect

Robotics and Autonomous Systems

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

Inferring robot goals from violations of semantic knowledge

Cipriano Galindo [a,*], Alessandro Saffiotti [b]

[a] *Department of System Engineering and Automation, University of Málaga, Campus de Teatinos, Málaga, Spain*

[b] *AASS Cognitive Robotic Systems Lab, Örebro University, Sweden*

ARTICLE INFO

Article history:
Available online 16 January 2013

Keywords:
Semantic maps
Mobile robotics
Goal autonomy
Planning
Norms
Ontology
Description logics
Fault detection and isolation
Environment stabilization

ABSTRACT

A growing body of literature shows that endowing a mobile robot with semantic knowledge and with the ability to reason from this knowledge can greatly increase its capabilities. In this paper, we present a novel use of semantic knowledge, to encode information about how things should be, i.e. *norms*, and to enable the robot to infer deviations from these norms in order to generate goals to correct these deviations. For instance, if a robot has semantic knowledge that perishable items must be kept in a refrigerator, and it observes a bottle of milk on a table, this robot will generate the goal to bring that bottle into a refrigerator. The key move is to properly encode norms in an ontology so that each norm violation results in a detectable *inconsistency*. A goal is then generated to bring the world back in a consistent state, and a planner is used to transform this goal into actions. Our approach provides a mobile robot with a limited form of *goal autonomy*: the ability to derive its own goals to pursue generic aims. We illustrate our approach in a full mobile robot system that integrates a semantic map, a knowledge representation and reasoning system, a task planner, and standard perception and navigation routines.

1. Introduction

Mobile robots intended for service and personal use are being increasingly endowed with the ability to represent and use semantic knowledge about the environment where they operate [1,2]. This knowledge encodes general information about the entities in the world and their relations, for instance, that a kitchen is a type of room which is used for cooking and which typically contains a refrigerator, a stove, and a sink; that milk is a perishable food; and that perishable food is stored in a refrigerator. Once this knowledge is available to a robot, it can be exploited to better understand the environment or plan actions [3–7], assuming of course that this knowledge is a faithful representation of the properties of the environment. There is, however, an interesting issue which has received less attention so far: what happens if this knowledge turns out to be in conflict with the robot's observations?

Suppose for concreteness that the robot observes a bottle of milk lying on a table. This observation conflicts with the semantic knowledge that milk, a perishable item, should be stored in a refrigerator. The robot has three options to resolve this contradiction: (a) to verify its perceptions, e.g., by looking for clues which may indicate that the observed object is not a milk bottle; (b) to update its semantic knowledge base, e.g., by adding a subclass of milk that is not perishable; or (c) to modify the environment, e.g., by putting the bottle in the refrigerator. While some works have addressed the first two options [6,8–10], the last one has not received much attention so far. Interestingly, the last option leverages the distinctive ability of robots to modify their physical environment. The goal of this paper is to investigate this option.

Our investigation proceeds in four steps. First, we address the problem of how to *encode* normative knowledge in a robot, that is, semantic knowledge on how things should be. For this we use a hybrid semantic map [8], which combines traditional robot maps with description logics [11], and enrich it with the notion of "normative" concepts. Second, we study how the robot can automatically *detect* violations of its normative knowledge, and *isolate* the causes of these violations. For this we rely on our encoding of norms to transform norm violations into logical inconsistencies. This allows us to use the mechanisms of description logics to detect an inconsistency and to identify the objects and relations which are involved in it. Third, we discuss how to go from the detection of a violation to a *recovery* strategy. We define a mechanism to automatically generate a goal, which represents the intention to achieve a specific state of the world that satisfies the violated norm. If this goal is fed to a standard task planner, it will result in a plan to execute the actions needed to bring the world back to a consistent state — provided of course that the robot has the right action repertoire.

Troubles rarely come alone, so our fourth and last step is to extend the above mechanism to the case of *multiple violations* of norms. This extension is not straightforward because of several reasons: (i) standard inference systems based on tableau methods

* Corresponding author.
E-mail address: cipriano@ctima.uma.es (C. Galindo).

0921-8890/$ – see front matter
doi:10.1016/j.robot.2012.12.007

do not behave well with multiple, simultaneous inconsistencies; (ii) violations may be inter-dependent, and solving one violation may produce another one; and (iii) some violations may be more important than others. We propose an algorithm that alternates violation detection, goal generation, and simulated recovery until a feasible sequence of recovery plans is found, also taking into account user defined priorities. This algorithm enables a mobile robot to generate a "to do list" in order to keep its workspace, as it perceives it, consistent with respect to a set of given norms.

In the rest of this paper we describe the above steps in more detail. We complement the formal descriptions with algorithms and examples to allow other researchers to reproduce our results. We also report a proof-of-concept experiment that shows the concrete applicability of our approach to real robotic systems. It should be emphasized that in this work we focus on the detection of norms violations and on the goal inference mechanism: the development of perception and action capabilities and the possible use of semantic knowledge in that context are beyond the scope of this paper.

In the next section we review some related work. Section 3 introduces our semantic map. In Section 4 we present the above first three steps, while Section 5 deals with the extension to the case of multiple concurrent norm violations. Section 6 reports the proof-of-concept experiment. We then discuss our results in Section 7 and conclude.[1]

2. Related work

The robotics community increasingly recognizes that future robots will have to be endowed with semantic knowledge [1,13,14]. Most current approaches rely on a shallow interpretation of semantic knowledge: the data used by the robot are simply augmented with labels, like "door" or "kitchen", which carry a semantic meaning to humans, but this meaning is not explicitly represented into the robot. Often these semantic labels are used for the human–robot interaction [15,16]. Many proposals have been made to allow the robot to acquire these labels automatically [17,3,5,18,19], even in a life-long perspective [20].

A few proposals exist, however, that take a deep semantic stance, in that semantic labels are embedded in a domain theory and are put in relation with other categories in some form of ontology. A robot can then effectively use this deep semantic knowledge for reasoning. For example, a robot may include an ontology that represents the relation that a kitchen is a type of room which contains a stove. This robot could use the fact that a room is labeled as "kitchen" to form the expectation that there is a stove in it; conversely, if the robot detects a stove in a room it will classify that room as a kitchen [6]. Proposals that adhere to a deep semantic stance include [6,7,4,8,9,21,22,10]. The European project RoboEarth goes one step further and use ontologies not only to allow a robot to perform new inferences, but also to enable meaningful communication among heterogeneous robots [23].

The above approaches have shown that endowing a robot with an explicit representation of semantic knowledge can increase the robot's behavioral autonomy, by improving their basic skills (planning, navigation, localization, etc.) with deduction abilities. Some researchers have also addressed the use of semantic knowledge to increase the robot's *goal* autonomy, that is, the robot's ability to pro-actively generate its own goals given generic motivations [24,25]. Two examples of this are the Curious George project [4] and the CogX European project [26]. In both cases, the authors explore the ability of the robot to generate its own perceptual goals, based on some innate "curiosity" that pushes it to increase its knowledge. Our work can also be seen as contributing to the robot's autonomy, but in a different vein. In our approach, the robot generates its own *action goals*, based on some innate "sense of order" that pushes it to maintain its environment in good order with respect to a given set of norms, encoded in a declarative way in its internal semantic representation.

The idea that a robot might modify its environment to better suit its needs has been proposed before [27,28]. The proponents of this idea see the modification of the environment as an acceptable option to improve the fitness of a robot to its environment, somehow complementary to the usual avenue of modifying the robot. In this paper we take a different point of view: we let the robot modify its environment to make it conformant to a set of norms. These norms may have many different purposes, but typically they are not meant to simplify the operation of the robot.

Our approach to goal autonomy can be seen as a case of normative goals applied to agents which act based on beliefs and intentions [29,30]. However, normative goals are often considered as simple *if-then* rules triggered when particular stimuli are given in the environment [31,32]. Other works have used the term *maintenance goals* to represent innate goals that are aimed to satisfy a particular state of the world over time, e.g., the battery level should always be over a certain value [33,34]. Our approach substantially diverges from those works, since it is not based on procedural rules, i.e., motivation-action pairs, nor on if-then rules. Instead, we rely on a declarative representation of the domain, from which the robot infers what should be done according to the current factual information in order to maintain the consistency between the environment and its representation. A declarative representation is easier to extend and to maintain. It also offers the possibility to be used in several directions: the same item of information that milk is stored in a fridge can be used to plan where to look for the milk and to detect that the milk is out of place.

The work presented in this paper is related to the general problem of fault detection, isolation, and recovery (FDIR) [35,36]. FDIR is the problem of detecting faults in a system, accurately isolating the causes, and bringing the system back to its normal behavior. Detection is typically based on the comparison between the expected behavior of the system, computed through a predictive model, and its observed behavior. Our approach can be interpreted as a special case of FDIR in which the observed system is the environment in which the robot operates, and the reference model is the knowledge base that describes the system at the semantic level. The use of an ontology-based model enables the use of knowledge-based reasoning methods for fault isolation and fault recovery. Approaches to FDIR based on ontology are not yet common in the literature, but a few examples exist [37,38].

Finally, we should mention that the research community in description logics is actively working on techniques to deal with inconsistency and incoherence [39,40]. This is especially important when different ontologies must be combined [41]. Those works are orthogonal to the one presented in this paper, since they handle inconsistency within the system, i.e., through a revision of the knowledge base; in contrast, the tenet of our approach is to handle the inconsistency outside of the system, i.e., through a modification of the physical environment.

3. A semantic map for mobile robot operation

The semantic map used in this work, derived from [6], comprises two different but tightly interconnected parts: a *spatial box*, or S-Box, and a *terminological box*, or T-Box. Roughly speaking,

[1] The work reported here is a major extension of earlier work presented at ECMR-2011 [12].