Biologically Inspired Cognitive Architectures (2013) 5, 1–9

Available at www.sciencedirect.com

SciVerse ScienceDirect

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

RESEARCH ARTICLE

A meta-cognitive architecture for planning in uncertain environments ☆

Vincenzo Cannella, Antonio Chella, Roberto Pirrone *

Dept. of Chemical, Management, Computer, and Mechanical Engineering (DICGIM), University of Palermo, Italy

Received 30 May 2013; received in revised form 1 June 2013; accepted 1 June 2013

KEYWORDS
Cognitive and meta-cognitive artificial agents;
Planning;
Uncertainty;
Markov Decision Processes

Abstract

The behavior of an artificial agent performing in a natural environment is influenced by many different pressures and needs coming from both external world and internal factors, which sometimes drive the agent to reach conflicting goals. At the same time, the interaction between an artificial agent and the environment is deeply affected by uncertainty due to the imprecision in the description of the world, and the unpredictability of the effects of the agent's actions. Such an agent needs meta-cognition in terms of both self-awareness and control. Self-awareness is related to the internal conditions that may possibly influence the completion of the task, while control is oriented to performing actions that maintain the internal model of the world and the perceptions aligned. In this work, a general meta-cognitive architecture is presented, which is aimed at overcoming these problems. The proposed architecture describes an artificial agent, which is capable to combine cognition and meta-cognition to solve problems in an uncertain world, while reconciling opposing requirements and goals. While executing a plan, such an agent reflects upon its actions and how they can be affected by its internal conditions, and starts a new goal setting process to cope with unforeseen changes. The work defines the concept of ''believability'' as a generic uncertain quantity, the operators to manage believability, and provides the reader with the u-MDP that is a novel MDP able to deal with uncertain quantities expressed as possibility, probability, and fuzziness. A couple u-MDPs are used to implement the agent's cognitive and meta-cognitive module. The last one is used to perceive both the physical resources of the agent's embodiment and the actions performed

* Corresponding author.
E-mail addresses: vincenzo.cannella26@unipa.it (V. Cannella), antonio.chella@unipa.it (A. Chella), roberto.pirrone@unipa.it (R. Pirrone).

2212-683X/$ - see front matter
http://dx.doi.org/10.1016/j.bica.2013.06.001

by the cognitive module in order to issue goal setting and re-planning actions.

1. Introduction

The effects of the actions performed by an artificial agent like an autonomous robot, interacting with a natural environment cannot be predicted exactly, due to both internal and external factors. Autonomous robots operating outdoor, and engaged in rescue or equipment repairing tasks are a good example of such a scenario.

Natural environments are prone to sudden changes of the operating conditions (i.e. weather conditions, earthquakes, accidents, and so on). Moreover, the robot may experience imprecise perceptions from the environment itself. Finally, the internal representation of the world may be not well defined; so the agent plans its actions erroneously.

On the other side of the agent's boundary, the internal status of the robot can influence its performance severely. The agent may experience partial failures or a decrease in power supply; so new goals arise that are aimed at repairing malfunctions, and can be in conflict with the task at hand.

A robot aimed at coping with all the issues mentioned above, has to be *self-aware* to plan its behavior properly. Moreover, the robot has *control* because it performs actions, which maintain its uncertain model of the world and its uncertain perceptions aligned; as a consequence the agent can *judge* if it is performing properly. In a few words, such a robot must exhibit meta-cognitive abilities while planning in uncertain environments.

Classical robotics deals with uncertainty in both perception and the model of the world, while neglecting the rest of the problem. A classical solution is offered by the Probability Theory. The concept of ''uncertainty'' has been investigated over the years, and new definitions have been formulated to capture other aspects of the meaning. Possibility Theory and Fuzzy Logic represent some of the result of such studies (Zadeh, 1999). On the contrary, meta-cognition in artificial agents is a crucial topic in the BICA literature (Chella, Cossentino, Gaglio, & Seidita, 2012; Samsonovich, 2012). This work presents a novel architecture that reconciles planning in uncertain conditions and meta-cognition expressed in terms of self-awareness and control on the coherence between the model of the world and the outer environment.

Some of the authors investigated the problems related to the interaction of an artificial agent with uncertain environment. In particular natural language interaction between humans and the agent was investigated. The problem arose when designing TutorJ (Pirrone, Cannella, & Russo, 2008) an architectural framework for building ITSs that are able to support a student in the learning process by supplying learning material customized to her cognitive needs, skills, and goals. The architecture is inspired to the Human Information Processor Model (HIPM) (Todorovski, Bridewell, Shiran, & Langley, 2005) where perceptual, cognitive and sensor-motor modules can be devised. Understanding natural language is an uncertain process, which makes not sure the meaning of the user's sentences. At first, the problem was faced up by adopting a planner agent that is able to manage uncertainty expressed through probability: the Partially Observable Markov Decision Process (POMDP) (Cannella & Pirrone, 2009). The agent's actions are a wide range of communication acts, aimed at coping with the learner's cognitive processes.

Next, the authors moved to modeling meta-cognition in such an agent. The artificial tutor mentioned above has to reflect on questions like ''how well I understood the user?'' or ''how well the user understood my sentence?'' in order to refine its next dialogue move. The agent has to access to its internal state, and to change it through proper actions. In a few words it needs meta-cognition. We used the classical definition of meta-cognition as *cognition about cognition* (Metcalfe, Shimamura, Metcalfe, & Shimamura, 1994) so the original POMDP is analyzed by the agent itself, whose internal reasoner was modeled by another POMDP that is a meta-cognitive module inserted in a two-level structure (Cannella, Pipitone, Russo, & Pirrone, 2010). The second POMDP has its own perceptions both from the external world (the dialogue flow) and from the actions issued by the cognitive POMDP. Cognitive actions are the actions in the previous version of the system. Meta-cognitive actions are communication acts aimed both to evaluate self regulated skills in the learner and to stimulate her to reflect on her meta-cognitive state in order to pursue a self-regulated learning.

In subsequent works (Cannella, Pirrone, & Chella, 2012) the authors investigated a unified management of uncertainty in Markov Decision Processes (MDPs), and presented a planner model able to manage different kinds of uncertainty together, expressed as probability, possibility, and fuzzy logics. We called this model ''uncertainty based MDP'' (u-MDP). The present work is a synthesis of the research activity described above. We present a meta-cognitive architecture for mobile robots engaged in tasks in natural environments. Such an architecture is based on a couple of u-MDPs where the cognitive MDP is devoted to deal explicitly with the external environment, while the meta-cognitive one governs the meta-cognitive abilities of the robots in terms of self-awareness and control. Choosing an implementation based on MDP instead of POMDP is not a limitation. In the outlined scenario, the autonomous robot can be supposed to work in a ''observable'' even if uncertain world. We're not dealing here with issues related to understandability of percepts. Moreover, we're currently investigating the extension of the presented architecture to the use of POMDP.

The rest of the paper is arranged as follows. Section 2 provides the reader with some theoretical background about the research on unified models of uncertainty. Section 3 presents our model for dealing with uncertainty and details the u-MDP. Section 4 presents the meta-cognitive architecture. Finally, in Section 5 conclusions are reported and future work is outlined.