Integrated Attention for Cognitive Robotics

Raúl Arrabales, Agapito Ledezma, and Araceli Sanchis

Departamento de Informática, Universidad Carlos III de Madrid
Avda. Universidad, 30. 28911 Leganés. Spain.
Email: {rarrabal; ledezma; masm}@inf.uc3m.es

Abstract: An autonomous robot is usually expected to perform its assigned task in complex and unstructured environments, where the high amount of sensory information acquired by sensors cannot be entirely processed in real-time. Like in living organisms, an attention mechanism is required in order to select most relevant stimuli from the environment. The outcome of this selection process comprises all the information that will be available to the control system of the robot. Therefore, an optimal operation of the attention mechanism is necessary for robot mission success. In this work, we present a novel cognitive architecture which permits the implementation and integration of efficient attention mechanisms for autonomous robots. Additionally, the application of this approach to the domain of unknown environment exploration is discussed and results are analyzed. Specifically, the influence of attention in the generation of efficient robot behavior is demonstrated.
Keywords: attention, cognitive modeling, intelligent agents.

1. Introduction

Attention is a cognitive or mechanical process developed in living beings as a mean to improve critical stimuli response time. Attention mechanisms have evolved in nature as an adaptation to complex and highly dynamic real world environments. The same sort of mechanism is also a clear need in the design of artificial autonomous agents. The aim of artificial attention mechanisms is twofold: on one hand, sensors have to be directed to the most interesting sources of information; on the other hand, the available sensory information has to be filtered in order to prevent unnecessary information processing. As a result of the application of the attention mechanism, the behavior of the robot is optimized. From the point of view of time, the robot is able to react faster to mission-related stimuli because processing resources are not wasted with non mission-related stimuli, which are ignored. From the point of view of space, the robot will only head towards areas of interest for its assigned task.

A significant number of works exist on attention, typically in the field of artificial vision [7], but also applied to autonomous mobile robots [5]. These works are usually inspired in neurobiological mechanisms like foveation and eye saccades [9]. However, in the present work we have considered some selected cognitive theories of consciousness as the main source of inspiration for the design of an artificial attention system for mobile robots. Although the proposed approach is currently being tested in the domain of unknown environment exploration, our aim is to design a multimodal attention mechanism flexible enough to be applied to different problem domains.

The aforementioned benefits can only be effectively demonstrated if the attention mechanism is properly integrated into the control architecture of the robot. In this work, we present a cognitive architecture for the control of autonomous robots which is inspired in psychological models of consciousness, and therefore exhibit a natural integration of cognitive processes like attention.

In section 2 we briefly describe the computational model of consciousness that we have implemented and how the artificial attention mechanism integrates within it. Then, section 3, covers the software architecture and design parameters in detail. In section 4, the experimental design and obtained results are discussed. Finally, we conclude in section 5 with an analysis of artificial attention benefits under the light of preliminary results and potential applications.

2. Computational Models of Consciousness and Attention

From the broad range of most established scientific theories of consciousness [3], we have focused on those which can effectively be translated into computational models. Furthermore, we have adopted a pragmatic machine consciousness approach, identifying the functional aspects of consciousness and using main psychological models as the inspiration for a practical implementation. As in Minsky's Society of Mind [10], a key feature common to most relevant theories of consciousness is the consideration of a large set of specialized processors or agents running concurrently in our brains. These small processors have evolved to perform highly specific tasks. Even though these processors are completely mindless, intelligence and consciousness emerge from their interrelations. A model for the organization and interaction of these processors called *Reasoning Consciousness Model* has been defined, which establishes a framework for the cooperation and competition of *specialized processors* [1]. The main sources of inspiration for such a framework are the Multiple Draft Theory (MDT) [6] and the Global Workspace Theory (GWT) [4]. From a purely functional perspective, both theories argue that consciousness emerges from dynamic coalitions of processors. Although these theories do not cover the actual implementation of the required underlying mechanisms, they provide metaphors that characterize the way processors collaborate and compete.

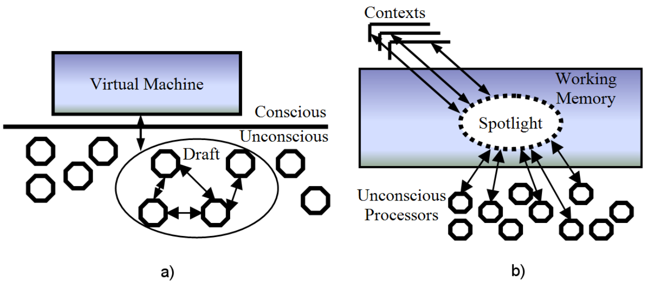

Fig. 1. Multiple Draft Theory and Global Workspace Theory

A deep comparison between MDT and GWT is out of the scope of the present paper. However, as they are the main inspiration of the proposed model we will focus on the concrete mechanisms that are used to implement their essential functionality. According to MDT, at any given time a coalition of unconscious processors, a *draft*, wins the multiple draft competition and the associated content becomes the conscious content of the mind (see Figure 1a). In GWT, consciousness is illustrated with a theater spotlight simile that represents the focus of consciousness directed by attention. Working memory takes the role of the scene, and behind scenes, unconscious contextual systems form the events taking place in the spotlight (see Figure 1b). There is a common denominator in both theories: the set of parallel processes collaborating and competing in a common workspace in the search of a global solution. As we explain below, this kind of frameworks can be implemented following the same design principles as used in other similar AI paradigms, like blackboard systems. Hofstadter's Copycat model is an illustrative example [8]. However, our implementation differs from classical blackboard systems and other cognitive architectures like Copycat, as it is integrated with a real-time control architecture and oriented to its direct application in the field of autonomous robotics.

3. THE *CERA/CRANIUM* ARCHITECTURE

CERA (*Conscious and Emotional Reasoning Architecture*) is a three-layer software architecture that allows the integration of cognitive components for autonomous agent control [2] (see Figure 2). CERA Physical Layer implements agent specific sensorimotor low-level control. CERA Instantiation Layer encloses the mission specific representation and processing. Finally, the CERA Core Layer implements a computational model of consciousness which includes several modules that represent different cognitive processes. CERA has been designed to work as a flexible test-bed for testing different cognitive processes and their interactions. Therefore, each module in the CERA Core Layer can be independently activated or deactivated. The discussion and results presented in this work are mostly related with one of these processes: the attention mechanism.

The computational model of consciousness implemented in CERA determines how the cognitive processes in the Core Layer interact with each other, and more significantly here, how these modules influence the way the lower layers work. Basically, the functional model of consciousness that we have implemented allows the attention module to send real-time

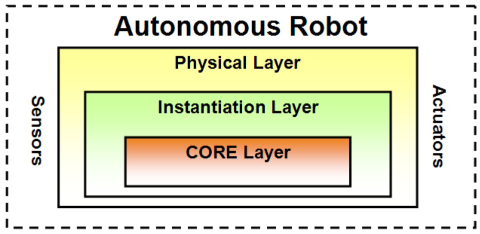

Fig. 2. CERA Architecture Layers

commands from CERA Core Layer to lower layers, thus modifying the processes of sensory information representation (perception) and next action selection (behavior). In other words, the attention module produces a cognitive bias within the CERA architecture (see Figure 3).

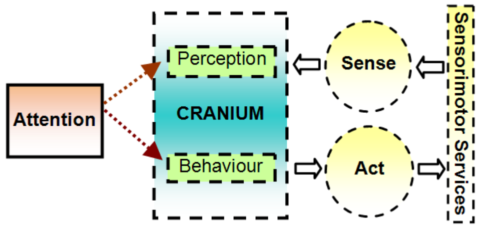

Fig. 3. Attentional Bias

CRANIUM (*Cognitive Robotics Architecture Neurologically Inspired Underlying Manager*) provides a framework in which CERA can execute thousands of asynchronous but coordinated concurrent processes [1]. It is actually inspired by the way the brain works from a systems level perspective, where specialized regions process information coming form the senses and/or other specialized regions. This scheme is analogous to the *pandemonium* architecture described by Dennett [6], where many processors (or *demons*) compete with one another for activation. The CRANIUM Workspace is a particular implementation of a pandemonium where each of these processors is designed to perform a specific function. The level of activation of processors is calculated based on a heuristic of how much they can contribute to the global solution sought in the workspace. In the current design we use two CRANIUM Workspaces: the first one is located in the CERA Physical Layer where specialized processors are fed with information coming from the sensors, the second workspace is located in the CERA Instantiation Layer and its processors are fed with more elaborated information coming from the Physical Layer (See Figure 4).

The raw data coming from the sensors is initially processed in the CERA Physical Layer by specific sensor preprocessors. These preprocessors build single representations units called *Simple Percepts*. These mono-modal percepts are composed of sensor data plus other associated physical state that characterizes the perceived stimulus. Basically, a time-stamp and relative position of the source of the percept are calculated and packaged together with the sensor data in a Simple Percept object. All simple percepts produced by sensor preprocessors are submitted to the Physical CRANIUM Workspace, where they become accessible to all the physical specialized processors. The processors located in the physical layer are able to

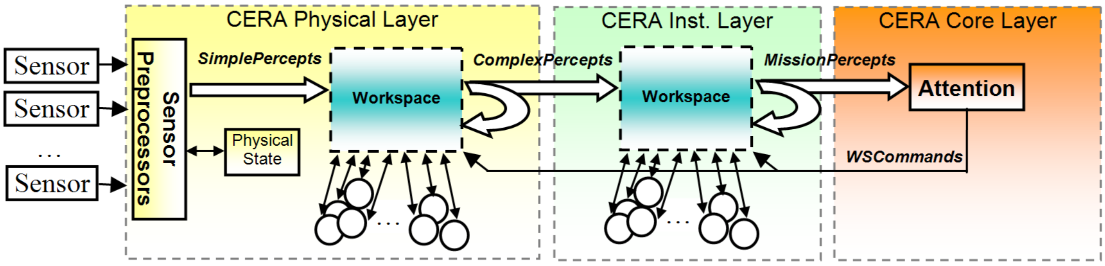

Fig. 4. CRANIUM Workspaces

combine both different modality and same modality simple percepts and build the so-called *Complex Percepts*, which are more elaborated multi-modal representations of the physical environment. The complex precepts generated in the physical workspace are both re-submitted to the same workspace (so other processors can use it for further processing) and also submitted to the CERA Instantiation Layer workspace. An analogous process takes place in this higher level CRANIUM Workspace. Complex percepts are processed by instantiation layer processors and more elaborated mission-specific representations are obtained. These new objects, called *Mission Percepts*, are both re-submitted to the same workspace for further processing and submitted to the CERA Core Layer. Although the same processing could be carried out using just one workspace, we have decided to use two separate workspaces in order to keep the independency between physical robot specific processors (Physical Layer) and mission specific processors (Instantiation Layer).

The cognitive functionality modules located in the core layer receive this flow of elaborated perceptual information; however they are not passive listeners. The way the lower levels workspaces make available the data to the specialized processors is modulated by commands sent from the CERA Core Layer. Specifically, the attention module can send context commands to the workspaces (*WSCommands*) indicating what sort of information should be under the spotlight of the working memory. Note that the CRANIUM Workspace also represents the working memory (or short-term memory) of the robot, where all the content available for reasoning is stored. The contents that the (explicit) specialized processors can receive are only those which are "illuminated" by the "spotlight" of attention. In order to preserve robot physical safety and allow physical level goals ("reflexes") to be triggered when needed, another class of specialized processor called implicit processor has been defined in the physical layer. Implicit processors receive all the simple percepts that they have subscribed for, independently of the perceptual bias induced by the attention module. This mechanism allows an implicit monitoring of robot status and promptly execution of higher priority low-level goals.

The incremental composition of perception knowledge, from raw sensory information to simple percepts, then complex percepts, and finally mission percepts, is distributed across the operation of the different specialized processors. This composition is directed by processors activation, which in turn is biased by attention. Each processor has an activation function and a set of parameters that can be used to calculate

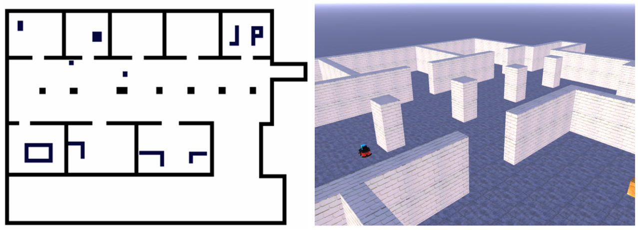

Fig. 5. Simulated indoor office-like environment

the degree of membership to a given context. Although the concept of context defined in the GWT is much broader, the contexts we have used are typically defined using criteria established by simple robot physical and mission goals. Basic criteria are restrictions on the dimensions of time and space. For instance, there could be a context for the concept of "*now*", which involves all the percepts contained in the workspaces that have been marked with a time-stamp close enough to CERA current internal time. Analogously, a context for the concept of "*far*" can be easily defined in CERA/CRANIUM having the processors to look for percepts whose relative position parameter is above a certain threshold. Additionally, processors calculate their own activation level in terms of the relevance of the percepts they are building to the current active contexts. An example of the application of this approach to the particular problem domain of navigation in unknown environments is presented in the next section.

4. APPLYING ATTENTION TO NAVIGATION

The results presented in this paper are all obtained using a (real and simulated) robotic base *Pioneer 3-DX* equipped with an onboard computer, frontal and rear bumpers and a frontal sonar ring. The eight frontal sonar transducers cover an angular range of 196 degrees and are used to measure distances to obstacles and map the surroundings of the robot.

In order to test the influence of our proposed artificial attention mechanism in the generation of optimal robot behavior we have confronted our prototype to the problem of unknown environment exploration. Given a typical indoor office environment - and its corresponding computer-simulated version (see Figure 5) - we have thoughtfully tested the performance of the robot governed by CERA/CRANIUM. In order to specificaly analyze the effects of the attention mechanism in the generated behavior, we have purposely neglected the typical localization problems by configuring a perfect odometry simulation. The following autonomous control strategies have been implemented and tested in order to analyze the influence of the attention mechanism in behavior features like path generation and mapping efficiency:

1-Context Behavior. With the purpose of having a reference of poor behavior performance, a simple and random non-attentional wander behavior has been analyzed. This non-attentional behavior simply apply some physical level goals ("*reflexes*") to keep the robot away from obstacles. Basically, no functionality from CERA Instantiation and Core layers is used, and only one possible context is considered. In the physical level workspace the specialized processors *Nearest-Obstacle-Processor* and *Possible-Impact-Processor* play a key role in the generation of this simple behavior (Figure 7 describes the typical outcome of this *1-Context behavior* in terms of explored areas and followed path).

2-Contexts Attention. This behavior implements a simple form of attention to local environment. A second context is added to be active when the robot is free of any collision risk. A physical level processor called *Open-Space-Processor* continuously calculates the angle (relative to the robot) of the direction where more open space is available. In this case, there is a basic mission goal activated in the instantiation layer that makes the robot to drive in the direction where more open space is available. Additionally, physical level goals are only triggered when the physical processors gain more activation than the instantiation processors, i.e. when the robot is too close to an obstacle, attention module switches the active context and the mission-level goal is vetoed (Figure 8 shows the outcome of this behavior).

3-Contexts Attention. In order to improve the exploration performance, a third context is defined to be active when the robot is not exploring unknown areas. A new processor called *Best-Heading-Processor* calculated the direction relative to the robot in which more unknown space can be discovered. Additionally, a mission-level goal is activated to move towards unknown areas of the map. The attention module induce a new bias (the third context) to move the robot away from already known areas (Figure 9 shows the typical outcome of this strategy).

Figure 6 shows a comparison of the performance of the three aforementioned behaviors. In average, 1-Context behavior is able to discover 0.19 square meters per second, while 2-Context and 3-Context behaviors map at a rate of 0.29 and 0.36 m^2/s respectively. The areas of the graph where there is no increment in discovered area correspond to the robot being "trapped" in a room. This is caused by the lack of a global navigation planning mechanism. However, focusing on the local navigation behavior, the application of attention contexts provides an efficient mechanism to integrate different autonomous strategies as a function of robot's current situation. The observed path followed by the robot in each of the analyzed cases demonstrates how the attention mechanism induces an adaptation by temporarily filtering the unnecessary sensory information. When a given goal cannot be achieved due to current situation, the sensory information correlated with that particular goal is ignored. For instance, the goal of heading towards unexplored areas should not be pursued when the robot is maneuvering to avoid a wall (see Figure 9).

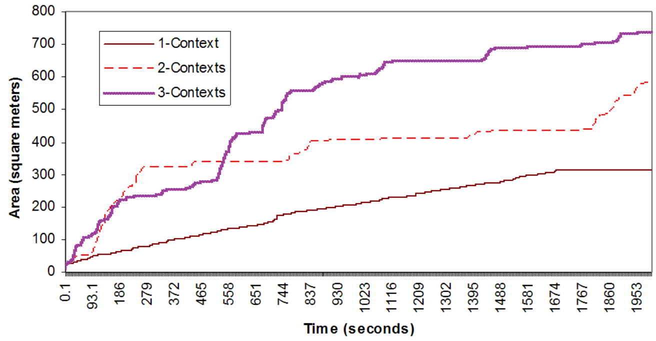

Fig. 6. Exploring Performance

5. CONCLUSIONS

A mechanism for robot attention has been developed as a highest level cognitive process of an autonomous control architecture. The cognitive bias induced by attention in the lower levels of control have been analyzed and the resulting behavior confronted to non-attentional strategies. An efficient method for integrating attention into a complex hybrid architecture has been demonstrated given the problem of unknown environment mapping. The proposed technique provides an efficient way to select areas of interest out of the available sensory space based on dynamic contexts. New problem-specific context can be easily defined within the proposed framework.

As future work, more challenging problem domains, including simultaneous localization and mapping with real odometry, are to be tested. Additionally, the interaction between attention and other higher cognitive processes is to be explored in detail as the proposed architecture is enhanced with other core capabilities like learning and planning. In fact, the addition of global path planning to our architecture could greatly enhance the performance making possible to define global navigation contexts.

References

[1] R Arrabales, A Ledezma and A Sanchis, Modeling consciousness for autonomous robot exploration. IWINAC 2007, Vol. 4527-4528 of Lecture Notes in Computer Science.

[2] R Arrabales and A Sanchis, A machine consciousness approach to autonomous mobile robotics. 5th International Cognitive Robotics Workshop. AAAI-06. 2006.

[3] A P Atkinson, M S C Thomas and A Cleeremans, Consciousness: Mapping the Theoretical Landscape, Trends in Cognitive Sciences. 2000, ch. 4.

[4] B J Baars, A Cognitive Theory of Consciousness, Cambridge University Press, New York, 1993.

[5] A Carbone, A Finzi, A Orlandini and F Pirri, Model-based control architecture for attentive robots in rescue scenarios. Auton.Robots, Vol. 24, No. 1, 2008, pp. 87–120.

[6] D C Dennett, Consciousness Explained, Little, Brown and Co, Boston, 1991.

[7] S Frintrop, Vocus: A visual attention system for object detection and goal-directed search, 2006.

[8] D Hofstadter, The Copycat Project: A model of mental fluidity and analogy-making, Fluid Concepts and Creative Analogies. Basic Books, 1995.

[9] L Itti, G Rees and J K Tsotsos, Neurobiology of Attention, Elsevier, San Diego, CA, 2005.

[10] M Minsky, Society of Mind, Simon & Schuster, New York, 1988.

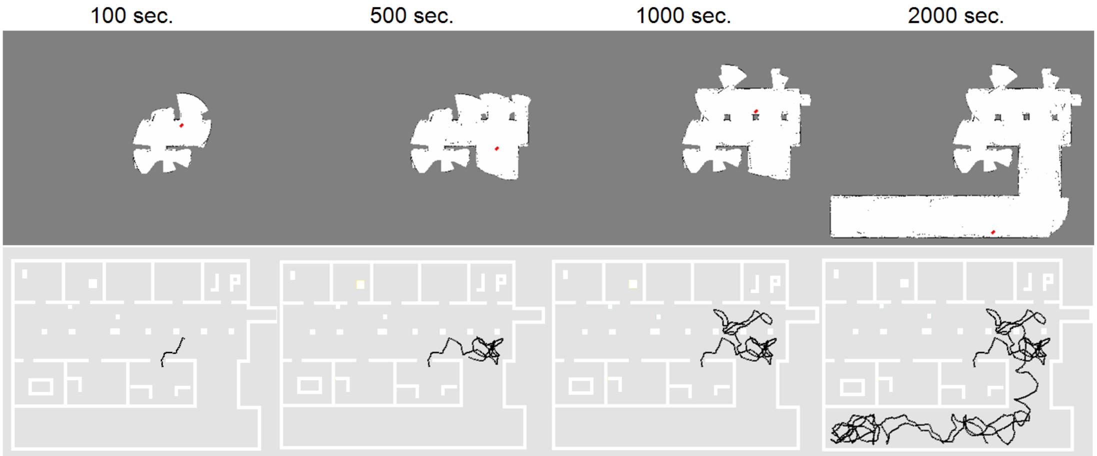

Fig. 7. 1-Context Non-Attentional Behavior

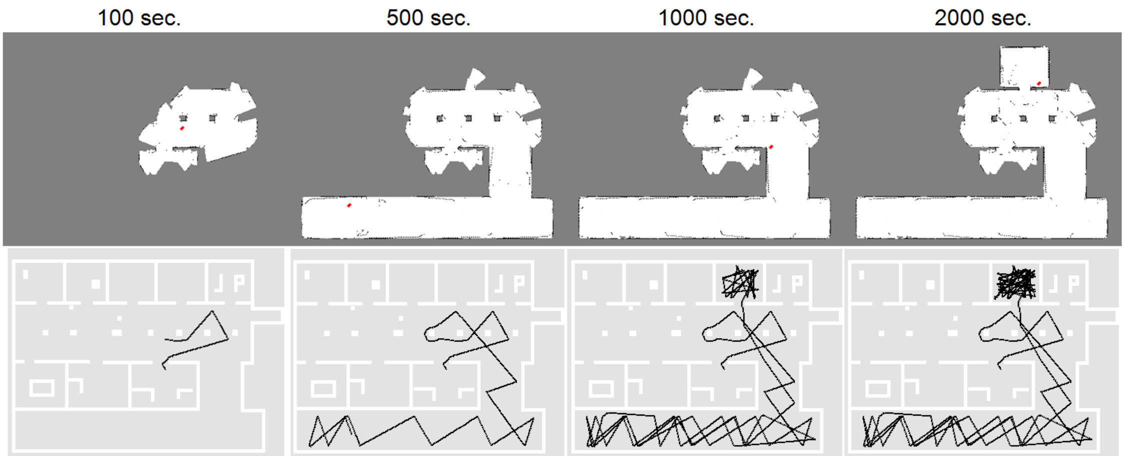

Fig. 8. 2-Contexts Attentional Behavior

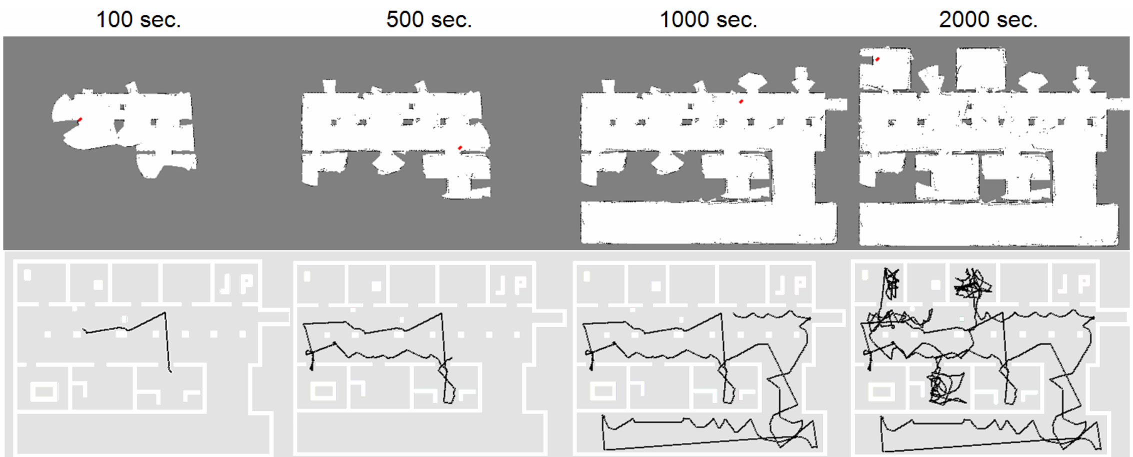

Fig. 9. 3-Contexts Attentional Behavior