BUILDING MAPS FROM INCOMPLETE ENVIRONMENT INFORMATION: A COGNITIVE APPROACH BASED ON THE GENERATION OF EXPECTATIONS

Luís Macedo[1,2]
Amílcar Cardoso[2]

[1]*Departamento de Eng. Informática e de Sistemas – Instituto Superior de Engenharia de Coimbra*
Quinta da Nora, Coimbra, Portugal
lmacedo@isec.pt

[2]*Centro de Informática e Sistemas da Universidade de Coimbra, Departamento de Eng. Informática da Universidade de Coimbra,*
Pinhal de Marrocos, Coimbra, Portugal
{macedo,amilcar}@dei.uc.pt

Abstract: Unfortunately, the real world is not crystal clear to agents. This makes mapping a hard and complex process. This paper describes a cognitive approach for mapping that relies heavily on the generation of assumptions/expectations for the missing observational information. In addition, we present the architecture for an explorer-agent that uses this approach to build maps and whose behaviour is guided by the emotions, drives and other motivations that it may "feel". We describe an experiment conducted in simulated environments in order to evaluate our approach for mapping.

Keywords: Mapping, Exploration, Architectures for Intelligent Autonomous Agents, Multiagent Systems, Cognition.

1. INTRODUCTION

Exploration involves acquiring spatial models of the physical environment, which is a process that is called mapping. Although several techniques have been successfully applied for mapping (see (Thrun, 2002) for a survey), there is still much to be done specially to build maps for dynamic, three-dimensional environments. Unfortunately, the real world is not crystal clear to agents. Agents almost never have access to the whole environment, mainly because of the incompleteness and incorrectness of their perceptual and understanding components. Actually, it is too much work to obtain all the information from a complex and dynamic world, and it is quite likely that the accessible information suffers distortions. Nevertheless, since the success of agents depends heavily on the completeness of the information of the state of the world, they have to pursue alternatives to construct good models of the world even (and specially) when this is uncertain. According to psychologists, cognitive scientists and ethologists (Kline, 1999), humans and, in general, animals attempt to overcome this limitation through the generation of *assumptions* or *expectations*[1] to fill in gaps in the present observational information. When the missing information becomes known to the agent, it might happen an inconsistence or conflict between it and the assumptions or expectations that the agent has. This inconsistence gives rise to the process of updating beliefs called belief revision. Gärdenfors (1994) defended that expectations are defeasible beliefs that are necessary to everyday reasoning. With respect to their cognitive origins, Gärdenfors argued that they are much like summaries of previous experiences. Thus, he defended that they are the result of inductive reasoning.

Psychological and neuroscience research over the past decades suggests that emotions play a critical role in decision-making, action and performance, by

[1] Although some authors use the terms assumption and expectation as synonyms, there are authors that make a distinction between them defending that an expectation has to do with future world states while assumptions are related with the current world state.

influencing a variety of cognitive processes (e.g., attention, perception, planning, etc.). Actually, on the one hand, recent research in neuroscience (Damásio, 1994)) supports the importance of emotions on reasoning and decision-making. On the other hand, there are a few theories in psychology relating motivations (including drives and emotions) to action (Izard, 1991). For instance, in the specific case of emotions, as outlined by (Reisenzein, 1996), within the context of the belief-desire theories of action (the dominant class of theories in today's motivation psychology) there have been proposals such as that emotions are action goals, that emotions are or include action tendencies, that emotions are or include goal-desires, and that emotions are mental states that generate goal-desires.

A series of experiments (e.g.: (Berlyne, 1950)) have shown that in the absence of (or despite) known drives, humans tend to explore and investigate their environment as well as seek stimulation. Curiosity is the psychological construct that has been closely related with this kind of behaviour. Sharing similar ideas with Berlyne, Shand (1914) defined curiosity as a primary emotion consisting of a simple impulse to know, controlling and sustaining the attention, and evoking the bodily movements that allow one to acquire information about an object. These approaches are closely related to the emotion concept of interest-excitement proposed by the differential emotions theory to account for exploration, adventure, problem solving, creativity and the acquisition of skills and competencies in the absence of known drives (Izard, 1991).

Moreover, as argued by Berlyne, in addition to novelty, other variables such as change, surprisingness, complexity, uncertainty, incongruity and conflict also determine this kind of behaviour related to exploration and investigation activities.

In this paper we describe an approach to mapping that is based on the generation of expectations or assumptions to fill in gaps in perceptual information. In our approach, mapping is the result of an emotional-based exploration of the environment carried out by a cognitive agent.

The next section presents the architecture of the agent. Subsequently, a simulated environment is described. Then, we present an experiment carried out to evaluate the role of expectations and assumptions in mapping. Finally, we make a discussion, present conclusions and the future work.

2. AGENT'S ARCHITECTURE

The architecture that we adopted for an agent (Figure 1) is based on the belief, desire, and intention (BDI) approach (Rao & Georgeff, 1995). Besides, the agent is of motivational kind (Bates, 1994), exhibiting a module of emotions, drives and other motivations. These play a central role in reasoning and decision-making since they may be thought as action goals (Reisenzein, 1996). As in many other agents' architectures, the architecture adopted in our work includes the following modules: memory; motivations (emotions, drives and other motivations); intentions/goals and desires; and, deliberative reasoning/decision-making.

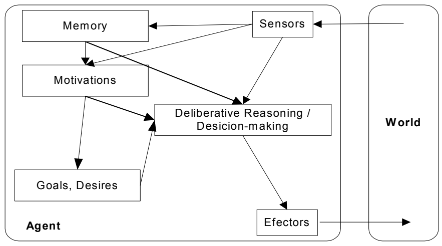

Fig. 1. Agent's architecture .

The perceptual information of the agent may be of three kinds: location (inferred from a sonar or from an optic sensor), distance (provided by an infrared sensor or sonar) and visual description of the entities (provided by an optic sensor) that surrounds it. This information is provided to the modules of memory, motivations, goals and desires, and reasoning/decision-making so that: it can be stored, it can elicit motivations, it is used to generate new goals, and it can be taken into account in the process of selecting an action for execution.

2.1 Memory

The memory of an agent stores information about the world. This information includes the configuration of the surrounding world such as the position of the entities (objects and other animated agents) that inhabit it, the description of these entities themselves, descriptions of the sequences of actions (plans) executed by those entities and resulting from their interaction, and, in generally, beliefs about the world. This information is stored in several memory components. Thus, there is a metric (grid-based) map (Thrun, 2002) to spatially model the surrounding physical environment of the agent. Descriptions of entities (physical structure and function) and plans are stored both in the episodic memory and in the semantic memory (Aitkenhead & Slack, 1987). We will now describe in more detail each one of these distinct components.

Memory for Entities. The descriptions of the entities perceived from the environment are stored in the component of memory called *episodic memory of entities*. Thus, this is a set of descriptions of entities (see Figure 2). Each element of this set (i.e., each description of an entity) is of the form *<ID,PS,F>*, where *ID* is a number that uniquely identifies the entity in the environment, PS is the physical structure of the entity, and *F* is the function of the entity. As we said above, the sensors may provide incomplete information about an entity (for instance, only part of the physical structure may be seen or the function of the entity may be undetermined). In this case the missing information is filled in by making use of the Bayes's rule (Shafer & Pearl, 1990), i.e., the missing information is estimated taking into account the available information and descriptions of other entities previously perceived and already stored in the *episodic memory of entities*. This means some of the descriptions of entities stored in memory are uncertain or not completely known (e.g.: element 4 of Figure 2).

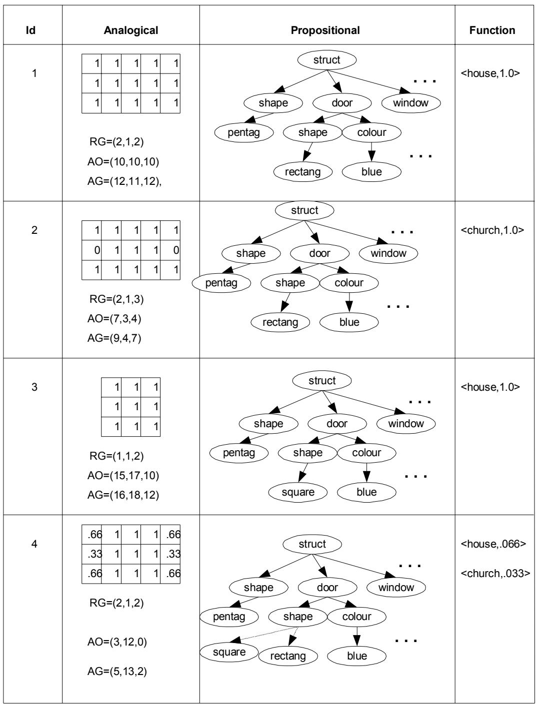

Fig. 2. Example of the episodic memory of entities.

The physical structure of an entity may be described analogically or propositionally (Aitkenhead & Slack, 1987; Eysenck & Keane, 1991). The analogical representation reflects directly the real physical structure while the propositional representation is a higher level description (using propositions) of that real structure (see Figure 2 for an illustration).

The analogical description of the physical structure of an entity is a tuple <*M,RG,AG,AO*>, where: *M* is the physical structure itself of the entity, which is represented in a three-dimensional matrix - the entity referential (a submatrix of the three-dimensional matrix of the environment) -, such that each cell is set to a value that expresses the probability of being occupied by the entity; *RG* represents the coordinates of the centre-of-mass of the entity in the three-dimensional entity referential; *AG* represents the coordinates of the centre-of-mass of the entity in the three-dimensional environment referential; and, *AO* represents the coordinates of the origin of the entity referential (origin of the three-dimensional matrix of the entity) in the environment referential.

The propositional description of the physical structure of an entity relies on the representation through semantic features or attributes much like in semantic networks or schemas (Aitkenhead & Slack, 1987). Entities are described by a set of attribute-value pairs that can be graph-based represented.

The function is simply a description of the role or category of the entity in the environment. For instance, a house, a car, a tree, etc. Like the description of the physical structure, this may be probabilistic because of the incompleteness of perception. This means, this is a set $F = \{<function_i, prob_i>: i=1,2, \ldots, n$, where n is the number of possible functions and P("function" = $function_i) = prob_i\}$.

Concrete entities (i.e., entities represented in the episodic memory) with similar features may be generalized or abstracted into a single one, an abstract entity, which is stored in the *semantic memory for entities*. Figure 3 presents a semantic memory obtained from the episodic memory of entities shown in Figure 2.

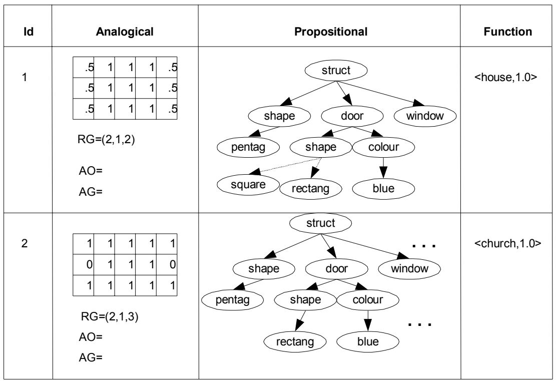

Fig. 3. Example of the semantic memory of entities.

Memory for Plans. Like entities, we may distinguish two main kinds of plans: concrete plans, i.e., cases of plans, and abstract plans (e.g.: (Bergmann & Wilke, 1996)). Concrete plans and abstract plans are interrelated since concrete plans are instances of abstract plans and these are built from concrete plans.

We represent plans as a hierarchy of tasks (a variant of HTNs (e.g., (Erol, Hendler, & Nau, 1994)) (see Figure 4). Formally, a plan is a tuple $AP = <T, L>$, where T is the set of tasks and L is the set of links. This structure has the form of a planning tree (Lotem & Nau, 2000), i.e., it is a kind of AND/OR tree that expresses all the possible ways to decompose an initial task network. Like in regular HTNs, this hierarchical structure of a plan comprises primitive tasks or actions (non-decomposable tasks) and non-primitive tasks (decomposable or compound tasks). Primitive tasks correspond to the leaves of the tree and are directly executed by the agent, while compound tasks denote desired changes that involve several subtasks to accomplish it. For instance, the leaf node *PTRANS* of Figure 4 is a primitive task, while *visitEntity* is a compound task. A task t is both conditional and probabilistic (e.g.:(Blythe, 1999; Macedo & Cardoso, 2004a, 2004b)).

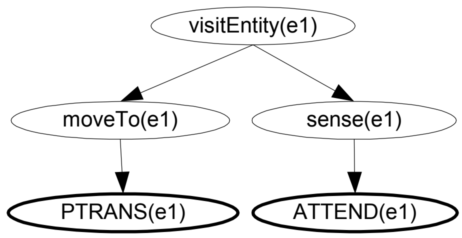

Fig. 4. Example of plan. Primitive tasks are represented by thick ellipses while non-primitive tasks are represented by thin ellipses.

The Metric Map. In our approach, a (grid-based) metric map of the world is a three-dimensional grid in which a cell contains the information of the set of entities that may alternatively occupy the cell and the probability of this occupancy. Thus, each cell <*x*,*y*> of the metric map of an agent *i* is set to a set of pairs $\phi^i_{x,y} = \{< p^i_1, E^i_1 >, \quad < p^i_2, E^i_2 >, \quad ..., \quad < p^i_{n_i}, E^i_{n_i} >, < p^i_{n_{i+1}}, 0 >\}$, where E^i_j is the identifier of the j^{th} entity that may occupy the cell <*x*,*y*> of the metric map of agent *i* with probability $p^i_j \in [0,1]$, and such that $\sum_{j=1}^{n_{i+1}} p^i_j = 1$. Note that the pair $< p^i_{n_{i+1}}, 0 >$ is included in order to express the probability of the cell being empty. Cells that are completely unknown, i.e., for which there are not yet no assumptions/expectations about their occupancy, are set with an empty set of pairs $\phi^i_{x,y} = \{\}$. Note also that each entity may occupy more than a single cell, i.e., there might be several adjacent cells with the same E^i_j. Figure 5 presents an example of a two-dimensional view of a metric map.

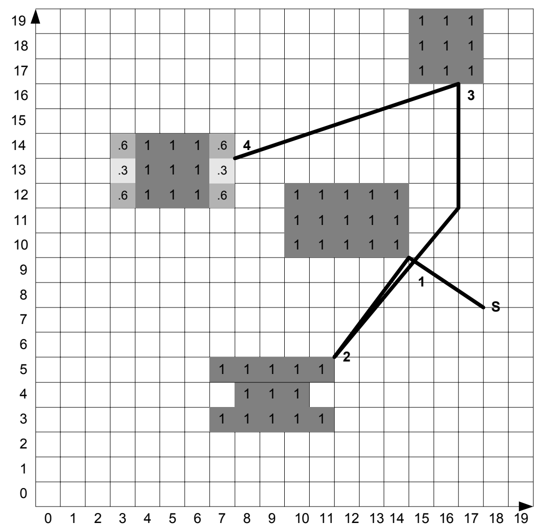

Fig. 5. An example of a metric map. Although metric maps are of three-dimensional kind, for the sake of simplicity, it is represented here only in two dimensions. For the same reason the identifier of the entities are not represented. The path followed by the agent to explore this environment (comprising buildings) is also depicted.

2.2 Emotions, Drives and Other Motivations

This module receives information from the current state of the environment and outputs the intensities of emotions such as surprise, sadness, happiness, anger, etc. Intensity of drives such as curiosity or hunger are also computed (see (Macedo & Cardoso, 2001, 2004a, 2004b) for more details about these computations). These feelings are of primary relevance to influence the behaviour of an agent.

2.3 Goals/Intentions and Desires

Desires are states of the environment the agent would like to happen, i.e., they correspond to those states of the environment the agent prefers. This preference is implicitly represented in a mathematical function that evaluates states of the environment in terms of the positive and negative feelings they elicit in the agent. This function obeys to the Maximum Expected Utility (MEU) principle (Russel & Norvig, 1995). The agent prefers always those states that make it feel more positive feelings (more positive emotions and the satisfaction of drives). Goals or intentions may be understood as something that an agent wants or has to do. These might be automatically generated by the agent or given by other agents.

2.4 Deliberative Reasoning and Decision-Making

The reasoning and decision-making module receives information from the internal/external world and outputs an action that has been selected for execution.

Roughly speaking, the agent starts by computing the current world state. This is performed taking into account the information provided by the sensors (which may be incomplete) and generating expectations or assumptions for the missing information. Then, new intentions/goals are generated and their Expected Utility (EU) computed. According to this EU, the set of goals of the agent are ranked, and the first one, i.e., the MEU goal is taken and a HTN plan is generated for it in case there is not yet one plan. The plan is then executed, primitive task by primitive task. Every time a primitive task is executed, this reasoning process is repeated. Sometimes, the ranking of the goal tasks is changed because there was an execution failure of a plan or because the EU of the goals changed.

We will now describe in more detail the step related with the generation of assumptions/expectations. The generation of plans is performed much like in HTN approaches (see (Erol et al., 1994; Macedo & Cardoso, 2004a)]). For the step related with the generation and ranking of agent's goals see (Macedo & Cardoso, 2004b).

Generating Assumptions/Expectations. As we said before, it is very difficult for an agent to get all the information about the surrounding environment. One reason is that the perceptual information is incomplete. However, taking as evidence the available information it is possible to generate expectations/assumptions for the missing information using the Bayes' rule (Shafer & Pearl, 1990):

$$P(H_i \mid E_1, E_2, \ldots, E_m) = \\ = \frac{P(E_1 \mid H_i) \times P(E_2 \mid H_i) \times \ldots \times P(E_m \mid H_i) \times P(H_i)}{\sum_{l=1}^{n} P(E_1 \mid H_l) \times P(E_2 \mid H_l) \times \ldots \times P(E_m \mid H_l) \times P(H_l)} \qquad (1)$$

where E_1, E_2, …, E_m are pieces of evidence, i.e., the available information, and H_i, i=1,2,…,n, are mutually exclusive and collectively exhaustive hypotheses for a specific piece of the missing information. The set of H_i's is the exhaustive set of instances assigned to that specific part of the missing information in past cases of entities. Each conditional probability $P(E|H)$ is given by the number of times E and H appeared together in the entities stored in memory divided by the number of times H appeared in those entities. In our work the evidence is the propositional description of the physical structure of the entities: the shape of an entity (rectangular, squared, etc.), the shape of their constituent parts (in case there are any), colour, etc. The hypotheses could be not only for parts of the descriptions of the physical structure but also for the function or category of the entity. In this case, the result is a probability distribution for the function of the entity (e.g., *P(Function=house)*=0.666; *P(Function=church)*=0.333). Based on this distribution, the analogical description of the entity may be now estimated taking into account the analogical descriptions of the abstract entities with those functions. This means that we are considering the reference class as comprising the entities with the same function. Notice that this resulting analogical description is probabilistic. Thus, considering the semantic memory presented in Figure 3 and the probability distribution for the function of an entity [*P(Function=house)*=.66, *P(Function=church)*=.33], the resulting analogical description is similar to that of entity 4 depicted in Figure 2. This is computed as follows. For all function X: (i) take the analogical description of each possible entity with function *X* and multiply the occupancy value of each cell by *P(Function=X)*; (ii) superimpose the analogical descriptions obtained in the previous step summing the occupancy values of the superimposed cells.

3. THE VIRTUAL ENVIRONMENT

In order to test the features of the agent's architecture presented above, we have developed a multi-agent environment in which, in addition to inanimate agents (objects such as buildings), there are explorer-agents whose goal is to explore the environment by analyzing, studying and evaluating it. In this simulated environment and in comparison to the real world a few simplifications were made such as the following: a parameter was defined for the visual range of the agents, i.e., objects out of that range are not visible by agents; for the sake of simplicity, the optic perception is confined to the shape of the visible part of the structure; the function of an entity is not accessible or can not be inferred from visual information unless the agent is at the same place of the entity; when an entity is perceived, its propositional description is provided directly by the virtual environment to the viewer agent (note that the agent's architecture does not include a module to transform the visual information into propositional information).

4. EXPERIMENT

We have performed an experiment to evaluate the role of expectations in mapping. Therefore, we ran the agent in 10 different environments (Figure 6). The agent had a time limit to explore the environments so that it can't explore exhaustively the whole environment. In all the environments, the agent started at location *S* and stopped after exploring entities 1, 2 and 3. Notice that these entities are maintained in all the environments. The rest of the entities (unvisited) change from environment to environment so that they correspond to environments with different degree of complexity. Environment 1 is the less complex since these unvisited entities contains only one entity (entity 4) that is different from those visited ones. Environment 2 had more diversity than environment 1 since it contains 2 entities that differ from entities 1, 2 or 3. Environment 10 is the most complex because it contains 10 different unvisited entities. Therefore, after exploring entity 3, the memory of the agent was the one depicted in Figure 2 plus probabilistic cases for the rest of the unvisited entities (just like element 4 in Figure 2). We compared the maps of the 10 different environments with the 10 maps built by the agent after exploring those 10 environments. Notice that these maps were built by computing assumptions/expectations for the unvisited objects. The difference or inconsistency between two maps was measured summing the difference between the occupancy values of any two correspondent cells (cells with similar coordinates) of the two maps. Figure 7 presents the results of this comparison. The agent built maps that are on average 70.62 (standard deviation = 17.92) different from the real maps. Best results were obtained with environment 1 and worse results in environment 10. The inconsistency between real and built maps monotonically increases with the increasing of the diversity of the entities in the environment.

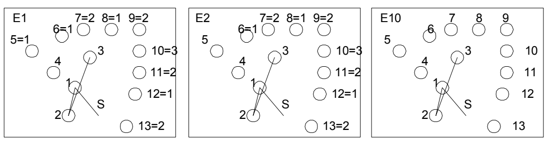

Fig. 6. Three of the 10 environments used in the experiment.

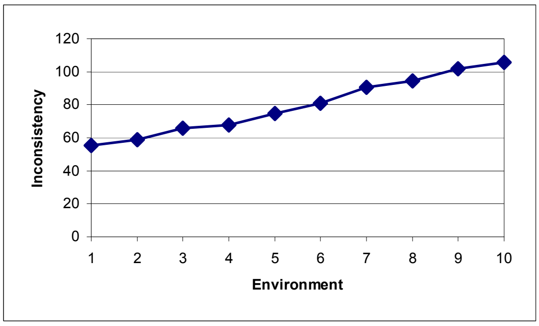

Fig. 7. Inconsistency between real and built maps after running the agent in 10 different environments.

5. DISCUSSION, CONCLUSIONS AND FUTURE WORK

The main advantage of the map learning process described in this paper is that it requires less time and less energy than that of involving a complete exploration of the environment. Actually, the agent does not have to explore all the regions of the environment, such as the invisible side of the entities, since it is able to predict that inaccessible information. The disadvantage of this approach is that the learned maps may be more inconsistent than those learned from an exhaustive exploration of the environment. However, this inconsistency is almost insignificant since, for instance, in the three-dimensional environment of the experiment with 18000 cells the average of inconsistency was 70.62 (0.4% of the environment). Moreover, we think these results depend heavily on the richness of the memory of the agent. More tests are required to prove this. We intend to test the agent's architecture presented in this paper in a robot. However, this requires the inclusion of a module to convert visual information into propositional information.

AKNOWLEDGEMENTS

The PhD of Luís Macedo is financially supported by PRODEP III.

REFERENCES

Aitkenhead, A., & Slack, J. (1987). *Issues in cognitive modelling*. London: Lawrence Erlbaum Associates.

Bates, J. (1994). The Role of Emotion in Believable Agents. *Communications of the ACM, 37*(7), 122-125.

Bergmann, R., & Wilke, W. (1996). On the Role of Abstraction in Case-Based Reasoning. In I. Smith & B. Faltings (Eds.), *Advances in Case-Based Reasoning - Proceedings of the Third European Workshop on Case-Based Reasoning* (Vol. 1168, pp. 28-43). Berlin: Springer Verlag.

Berlyne, D. (1950). Novelty and curiosity as determinants of exploratory behavior. *British Journal of Psychology, 41*, 68-80.

Blythe, J. (1999). Decision-Theoretic Planning. *AI Magazine, Summer 1999*.

Damásio, A. (1994). *Descartes'error, Emotion Reason and the Human Brain*. New York: Grosset/Putnam Books.

Erol, K., Hendler, J., & Nau, D. (1994). UMCP: A sound and complete procedure for hierarchical task-network planning, *Proceedings of the International Conference on AI Planning Systems* (pp. 249-254).

Eysenck, M., & Keane, M. (1991). *Cognitive psychology*. London: Lawrence Erlbaum Associates.

Gärdenfors, P. (1994). The role of expectations in reasoning. In M. Masuch & L. Polos (Eds.), *Knowledge Representation and Reasoning Under Uncertainty* (pp. 1-16). Berlin: Springer-Verlag.

Izard, C. (1991). *The Psychology of Emotions*. NY: Plenum Press.

Kline, C. (1999). *Observation-based Expectation Generation and Response for Behavior-based Artificial Creatures.* Unpublished MSc Thesis, Massachusetts Institute of Technology, Cambridge, USA.

Lotem, A., & Nau, D. (2000). New advances in GraphHTN: Identifying independent subproblems in large HTN domains, *dings of the International Conference on AI Planning Systems* (pp. 206-215).

Macedo, L., & Cardoso, A. (2001). SC-EUNE - Surprise/Curiosity-based Exploration of Uncertain and Unknown Environments., *Proceedings of the AISB'01 Symposium on Emotion, Cognition and Affective Computing* (pp. 73-81). York, UK: University of York.

Macedo, L., & Cardoso, A. (2004a). Emotional-based Planning. In C. Johnson (Ed.), *Proceedings of the AISB'04 Symposium on Emotion, Cognition, and Affective Computing*. Leeds: SSAISB.

Macedo, L., & Cardoso, A. (2004b). A Motivation-based Approach for Autonomous Generation and Ranking of Goals in Artificial Agents. In D. Kudenko (Ed.), *Proceedings of the AISB'04 Fourth Symposium on Adaptive Agents and Multi-Agent Systems*. Leeds: SSAISB.

Rao, A., & Georgeff, M. (1995). BDI Agents: From Theory to Practice, *Proceedings of the First International Conference on Multiagent Systems*.

Reisenzein, R. (1996). Emotional Action Generation. In W. Battmann & S. Dutke (Eds.), *Processes of the molar regulation of behavior*. Lengerich: Pabst Science Publishers.

Russel, S., & Norvig, P. (1995). *Artificial Intelligence - A Modern Approach*. Englewood Cliffs, NJ: Prentice Hall.

Shafer, G., & Pearl, J. (Eds.). (1990). *Readings in Uncertain Reasoning*. Palo Alto, CA: Morgan Kaufmann.

Shand, A. (1914). *The foundations of character*. London: Macmillan.

Thrun, S. (2002). Robotic mapping: A survey. In G. Lakemeyer & B. Nebel (Eds.), *Exploring Artificial Intelligence in the New Millenium*. San Mateo, CA: Morgan Kaufmann.