Experiments with an Emotion-based Agent using the DARE Architecture

Márcia Maçãs, Rodrigo Ventura, Luis Custódio and Carlos Pinto-Ferreira
Instituto Superior Técnico, Portugal
Institute of Systems and Robotics
{marcia,yoda,lmmc,cpf}@isr.ist.utl.pt

Abstract

The work presented in this paper is based on neurophysiology findings suggesting that efficient decision-making depends heavily on the emotions underlying mechanism. In particular, in this paper is described an agent architecture inspired on António Damásio's research which proposes that alternative courses of action in a decision-making problem are emotionally (somatic) marked as good or bad. These emotional marks not only guide the decision process, but also prunes the options leaving only the positive ones to be considered for further scrutiny. Some preliminary results of the research on an emotion-based agent architecture supported on the Damásio's somatic marker hypothesis are presented. In order to evaluate it, two implementations were developed and are described here. The first experiment was carried out in a simulated maze, requiring flexible behaviour and autonomy. The second one, although simulated, intended to test the architecture into an environment closer to the robotic one, being a preliminary implementation of the architecture on agents running in the RoboCup soccer simulator. The corresponding results of these experiments are discussed.

1 Introduction

Adequate decision making under difficult circumstances (in unpredictable, dynamic and aggressive environments) raises some interesting problems from the point of view of the implementation of artificial agents. At the first sight, in such situations, the agent should (ideally) be capable of performing deductive reasoning rooted on well established premises, reaching conclusions using a sound mechanism of inference, and acting accordingly. However, in situations demanding urgent action such an inference mechanism would deliver "right answers at the wrong moment". The research of António Damásio (Damasio, 1994) has shown that even in simple decision making processes, the mechanism of emotions is vital for reaching adequate results.

In particular, António Damásio's research proposes that alternative courses of action in a decision-making problem are emotionally (somatic) marked as good or bad. These emotional marks not only guide the decision process, but also prunes the options leaving only the positive ones to be considered for further scrutiny. In this paper an agent architecture (DARE[1]) inspired on this research is described. The basic idea underneath this architecture is the hypothesis that stimuli processing is made *simultaneosly* under two different perspectives: a *cognitive*, which aims at finding out *what the stimulus is* (by a mechanism of pattern matching), and another one, *perceptual*, intending to determine *what the agent should do* (by extracting relevant features of the incoming stimulus). As this latter process is much more rapid (in terms of computation) than the former, the agent can react even before having a complete cognitive assessment of the whole situation. Following the suggestions of Damasio, a somatic marker mechanism should associate the results of both processing sub-systems in order to increase the efficiency of the recognition process in similar future situations. On the other hand, the ability of anticipating the results of actions is also a key issue as the agent should "imagine" the foreseeable results of an action (in terms of a somatic mark) in order to make adequate decisions.

To illustrate the advantages of this approach in decision making, an implemented system shows a simple agent in a maze performing behaviours that result in survival while seeking the maze exit. Another example of an experiment made with the architecture is also described, requiring a more complex behaviour in a more dynamic simulated environment, the RoboCup soccer simulator (Kitano *et al.*, 1997). The goal of this last implementation was not to test possible multi-agent or social properties of the architecture, but rather to test how the architecture functions at an individual level in a more complex world and how it can be implemented so that produces different behaviours in different agents.

[1]DARE stands for (inverse order) "**E**motion-based **R**obotic **A**gent **D**evelopment". This work has been supported by the project PRAXIS/C/EEI/12184/1998 entitled "DARE: Study of Methodologies for the Development of Emotion-based Robotic Agents" funded by the FCT - Portuguese Foundation for Science and Technology under the PRAXIS program.

2 Emotion-based agent architecture

The proposed architecture for an emotion-based agent includes three levels: stimulus processing and representation, stimulus evaluation, and action selection and execution. Figure 1 represents the architecture with the main relationships among blocks represented by solid arrows (dashed arrows symbolize information accessing operations).

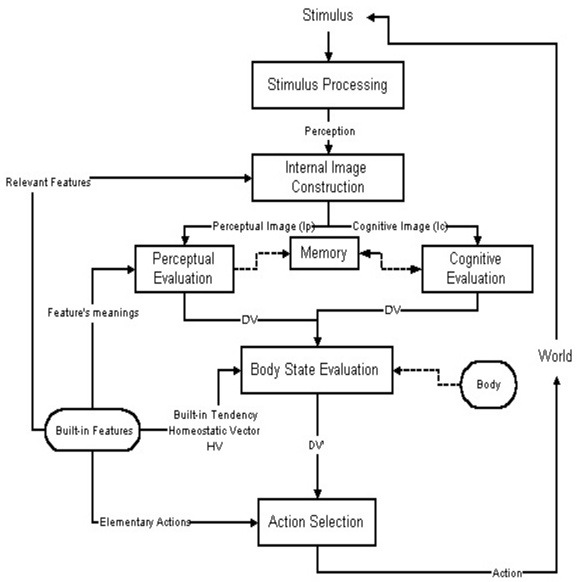

Figure 1: The proposed architecture.

The environment provides stimuli to the agent and as a consequence of the stimulus processing, the agent decides which action should be executed. During this stimulus-processing-action iterative process, decisions depend not only on the current incoming stimulus and internal state of the agent (body state) but also on the results got from previous decisions.

After the reception of a stimulus, a suitable internal representation is created and the stimulus is simultaneously analysed by two different processors: a *perceptual* and a *cognitive*. The former one generates a perceptual image (I_p) which is a vector containing the amount values of the relevant features extracted from the stimulus. For instance, for a prey the relevant features of the predator image might be the colour, speed, sound intensity, and smell, characteristics that are particular to the corresponding predator class. The definition of what are relevant features and their corresponding values of desirability is assumed to be built-in in the agent. This perceptual, feature-based image, as it is composed of basic and easily extracted features, allows the agent to efficiently and immediately respond to urgent situations. The cognitive processor generates a cognitive image (I_c) which is a more complex representation of the stimulus. This processing aims at performing a pattern matching of the incoming stimulus w.r.t. cognitive images already stored in memory. As this processor performs a more heavy computation processing, the cognitive image is not suitable for urgent decision-making.

With the two images extracted from the stimulus, the process proceeds through a parallel evaluation of both images (I_p and I_c). The evaluation of the perceptual image consists of assessing each relevant feature included in the I_p. From this evaluation results what is called the perceptual Desirability Vector (DV). This vector is computed in order to establish a first and basic assessment of the overall stimulus desirability. In the perceptual evaluation, the DV is the result of a mapping between the built-in desirability of each feature and the amount of the feature found in the stimulus. The cognitive evaluation differs from the perceptual in the sense that it uses past experience, stored in memory. The basic idea is to retrieve from memory a DV already associated with cognitive images similar to the present stimulus. Since a cognitive image is a stimulus representation including all extractable features of it, two stimuli can be compared using an adequate pattern matching method.

When the evaluation of the I_p does not reveal urgency the cognitive evaluation is performed. It consists of using the I_p as a memory index to search for past obtained cognitive images similar to the present I_c. It is here hypothesized that it is likely to have the current I_c similar to others with the same dominant features (I_p). If the agent has been already exposed to a similar stimulus in the past, then it will recall its associated DV, being this the result of the cognitive evaluation. This means that the agent associates with the current stimulus the same desirability that is associated with the stimulus in memory (which can be different from the perceptual DV). If the agent has never been exposed to a similar stimulus, no similar I_c will be found in memory, and therefore no DV will be retrieved. In this case, the DV coming from the perceptual evaluation is the one to be used for the rest of the processing.

Let I_p^c be the current stimulus perceptual image, I_c^c the current cognitive image, DV_p^c the current DV result of the perceptual evaluation, and I_p^m, I_c^m, and DV^m those stored in memory, finally DV^c is the result of both evaluations of the stimulus. $\varepsilon_{urgency}$ is a vector of thresholds that defines what values are considered urgent in each component of the DV, and $\varepsilon_{similar}$ is a threshold that defines the similarity between cognitive images.

$$DV^c = \begin{cases} DV_p^c & \text{If } \exists i : DV_{i_p}^c > \varepsilon_{i_{urgency}} \text{ or} \\ & \nexists I_c^m \in MEMO : |\Delta(I_c^c, I_c^m)| < \varepsilon_{similar} \\ & (MEMO = \phi \text{ or failed search}), \\ DV^m & \text{otherwise.} \end{cases}$$

In this architecture is introduced the notion of body which corresponds to an internal state of the agent, *i.e.*, the agent's body is modelled by a set of pre-defined variables and the body state consists in their values at a particular moment. The internal state may change due to the agent's actions or by direct influence of the environment. The innate tendency establishes the set of body states considered ideal for the agent, through the definition of the equilibrium values of the body variables — the Homeostatic Vector (HV). In other words, this comprises a representation of the agent's needs. The body state evaluation consists of an anticipation of the effects of alternative courses of action: "will this action help to re-balance a particular unbalanced body variable, or will it get even more unbalanced?". This action effect anticipation may induce a change on the current stimulus DV, reflecting the desirability of the anticipated effects according to the agents's needs.

Let $Action = \{a_1, a_2, ..., a_k\}$, be the built-in set of elementary actions that the agent can execute. Consider BS^c the current body state, BS'_a is the anticipated body state for action $a \in Action$, by retrieving from memory the body state change (ΔBS^m_a) caused by the execution of a with similar stimulus, $BS'_a = BS^c + \Delta BS^m_a$. When no past experience exists the anticipation is based on the relevant features of the stimulus (I_p) which give a rough estimate (DV_p) of the possible change in the body.

Given BS'_a, let p be the number of components of the body state, DV and HV, $\beta = \beta_1, \beta_2, ..., \beta_p$ the possible increments to each component of the DV, and $\alpha = \alpha_1, \alpha_2, ..., \alpha_p$ be the possible decrements to each component of the DV (corresponding to the ones in the body state). So,

$$(DV^c_i)'_a = \begin{cases} DV^c_i + \beta_i & \text{If } \Delta(BS'_{i_a}, HV_i) < \Delta(BS_i, HV_i), \\ DV^c_i - \alpha_i & \text{If } \Delta(BS'_{i_a}, HV_i) > \Delta(BS_i, HV_i), \\ DV^c_i & \text{If } \Delta(BS'_{i_a}, HV_i) \simeq \Delta(BS_i, HV_i). \end{cases}$$

Where $(DV^c_i)'_a$ is the ith component of the DV resulting from the anticipation of a, DV^c_i is the ith component of the current DV, $BS'i_a$ is the ith component of the anticipated body state for action a, and HV_i is the ith component of the homeostatic vector.

This anticipation results in changes in the DV so that:

- if the anticipated action leads to a more balanced body state the DV associated with the stimulus and action is increased ($DV' > DV$),
- if it leads to a worse unbalance than the current one, then the DV decreases ($DV' < DV$),
- finally, if no relevant changes happen to the body state, the DV' is equal to the DV.

As the agent's decisions depend on a particular body state — the one existing when the agent is deciding, it will not respond always in the same manner to a similar stimulus. On the other hand, the existence of a body representation forces the agent to behave with pro-activeness — because its internal state drives its actions — and autonomy — because it does not rely on an external entity to satisfying its needs.

After finishing the evaluation process, the agent will select an adequate action to be executed. The action with the best anticipated contribution for the agent's overall body welfare will be selected as the one to be executed next.

Let C be a set containing indexes of components of the body state, i.e. $C = \wp(\{1, 2, ..., p\})$, which will represent those that are unbalanced. At a given time it contains the indexes i that satisfy $|\Delta(BS_i, HV_i)| > \varepsilon_{need_i}$, where ε_{need_i} is a threshold defining the unbalance in the ith body component .

Let $W = [w_1, w_2, ..., w_p]$ be the weights given to each DV component (related to the body component) at a given time, so that $w_i > w_j$, $\forall i \in C \wedge \forall j \ni C$ and $\sum_k w_k = 1$.

Let

$$DV_a = \sum_i^p (DV^c_i)'_a \cdot w_i,$$

where DV_a is the final DV associated with action a, so

$$A = \arg\max_a DV_a, \ \forall a \in Action,$$

where A is the action to execute.

After the selected action being executed, the changes in the environment will generate a new stimulus to be processed.

3 The Maze implementation

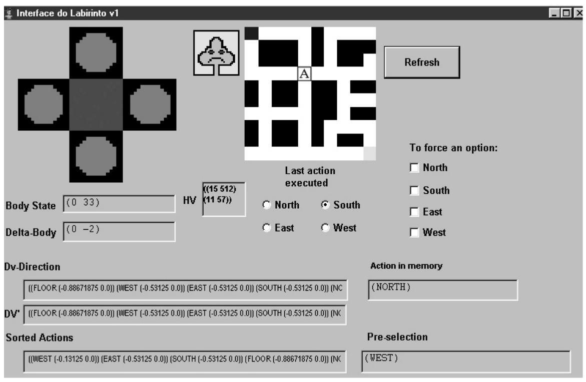

Figure 2: The interface of the maze implementation.

In order to evaluate the proposed architecture, an agent in a maze trying to survive and find the exit was implemented (Maçãs *et al.*, 1999). The maze is represented by a grid (Figure 2), where the agent proceeds by discrete steps and, when it reaches a new position, it perceives new images, representing walls, ways to other positions, or food (examples are shown in figure 3). The agent can perform the elementary actions of moving towards North, South, East or West, and eating.

3.1 Implementation

The agent's body internal state consists of two variables, light and energy, whose ideal values are defined by the homeostatic vector (HV). The difference between a body state and the HV implicitly defines the agent's needs. The energy need arises because every action executed by the agent leads to a reduction of energy in its body state. To compensate this loss of energy, the agent must seek for food and feed itself. On the other hand, inside the maze there is light only near the exit and even there it does not completely satisfy the agent needs. The need of light will make it find the exit, since only outside of the maze there is enough light to satisfy the agent.

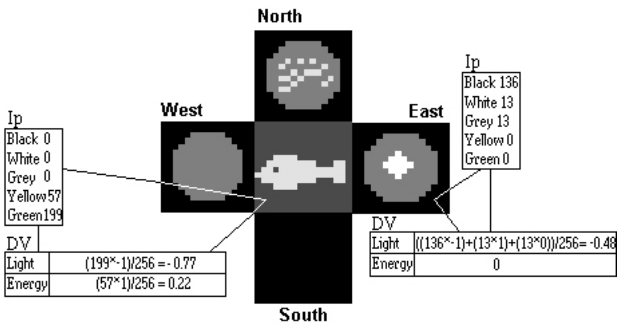

Figure 3: The visual stimuli in a particular position and some of its associated I_p and DV resulting from perceptual evaluation. The South stimulus shows a wall (all black pixels). The West stimulus signals a way (circle of grey pixels with black background). The North stimulus represents food smell (yellow pixels in the way). The East stimulus represents some light (white pixels in the way) and its I_p and the DV are shown . Finally, the floor stimulus (the middle image) represents food (yellow pixels with green background) and its I_p and DV are shown.

The stimuli perceived by the agent in the maze are bitmaps (figure 3). In this implemented system the relevant features to be extracted in the perceptual processing are colours of pixels in the bitmap. The agent gives (built-in) relevance to some of the colours, while others are not considered relevant. The perceptual image (I_p) is a vector containing the amount of each relevant colour present in the percept. In the present case, the cognitive image (I_c) of a stimulus is the bitmap itself.

Each relevant feature represented in the I_p leads to a built-in assessment. For example, yellow and white colours, representing respectively food and light, are considered to be good, black colour means darkness so is considered very negative. The desirability vector (DV) has two components, one measuring the desirability of the stimuli towards the need of light and the other to the need of energy.

The cognitive evaluation was implemented as a search for cognitive images in memory that are similar to the I_c of the current stimulus. The basic memory structure is a frame containing the I_p, I_c and DV resulting from previous processed stimulus. The current I_p is compared to those already in memory. When a match is found, *i.e.*, an I_p with the same number of pixels of each relevant colour, the I_c associated with it is compared with the current one. Two cognitive images are similar if the difference between the two bitmaps is under a pre-defined threshold. When they are found to be similar, the DV of the I_c in memory will be the DV of the present stimulus.

This process establishes a basic learning mechanism, since the DV stored in memory may be different from the perceptual one. In fact when the execution of a chosen action results in a very different (worst or better) body state than the anticipated one, the DV in memory is modified accordingly. For instance, the cognitive evaluation allows the agent to make the discrimination between good food and rotten food (figure 4). Both images have positive perceptual evaluations (yellow colour), but they have different shapes. After eating the rotten food and experiencing the decrease of energy in the body state, the DV associated with the cognitive image is re-marked in memory. Therefore, in the future, if the agent finds rotten food, the negative DV will be retrieved and, despite the perceptual evaluation, the agent will avoid to eat the rotten food.

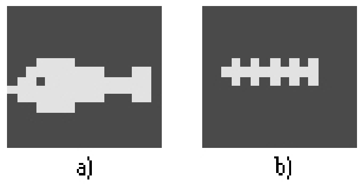

Figure 4: The visual stimuli of good food (a) and rotten food (b). Both consist in yellow pixels with green background resulting in a good perceptual evaluation. After experiencing the bad effect of rotten food (b), its shape is recognized and results in a bad cognitive evaluation.

The implementation of the body state evaluation is basically made by comparing the current body values of light and energy to the ideal body state (HV). The anticipation of the effects of possible actions is made in order to find out if an action will approach the body state to the HV or if it will make it more unbalanced. In the process of anticipation, given the current stimulus and the set of possible actions, for each pair stimulus/action a new DV (DV') is determined by increasing or decreasing the DV according to the body state anticipation. The anticipation is made either by recalling the change in the body state in previous similar experiences or by perceptually anticipating the possible action effect. In the former case, the memory is used to retrieve the change in the body state occurred when the current anticipated action was chosen with a similar stimulus (similar I_c). This change is applied to the current body state and the result is evaluated. On the other hand, in the latter case, if no match is found in memory then the perceptual image is used, changing the body state according to its contents and the associated DV. Since there is a relationship between the features of the I_p (colours), the components of the DV and the components of the body state (light and energy), this rather rough anticipation can sometimes be useful and extremely fast.

When current stimuli, corresponding to the current position of the agent in the maze, are evaluated the decision is made by choosing the action with maximum DV. Since each DV has two components (light and energy), the selection of the best action is made by giving different weights to the DV components. The weights are determined by the current body state: the more unbalanced component corresponds to the most weighted component of the DV. This way the selection of the best action is guided by the current agent needs. When the comparison results in several equally good actions the choice is random among them. After the decision making, the selected action is executed, and the effect is new stimuli coming from the environment.

In this implemented system it is possible for the agent to make and execute plans. When a need arises, the agent searches in memory for stimuli that lead to its satisfaction. If a sequence of actions found in memory ends in a (relevant) stimulus with high DV in the component corresponding to the current need, a meta-sequence containing that sequence and those that followed it (in temporal order) is used to form a plan. The execution of the plan results on the return of the agent to the desired location. The plan is simplified in order to avoid the lost of energy, specially in what refers to redundant sub-sequences (dead-ends) that might compose the plan. These are eliminated from the plan, since they are not important to achieve the goal. Moreover, the plan is executed only if there is energy to do so. The execution of a plan is made by slightly increasing the DVs of the planned direction/action. Nevertheless, the agent is continuously evaluating the stimulus along with the execution of the plan. Therefore, if the environment changes drastically, due to this continuous evaluation, the agent can either find some direction with a better DV than the planned one, or it can react accordingly when the stimulus in the planned direction has changed, being no longer desirable. This kind of planning can be seen as a layer above the architecture, which does not rely only in the image evaluation but also on the higher level memory contents and organization.

3.2 Results

Several tests were made with the implemented system, revealing an acceptable overall behaviour. Figure 5 shows parts of the path followed in one of these tests.

The agent starts at the most top left location on the maze, with a balanced level of energy and a very unbalanced level of light. The need of light arises and to satisfy it the agent must seek the exit. Initially, the memory is empty, so only the perceptual processing is used to make

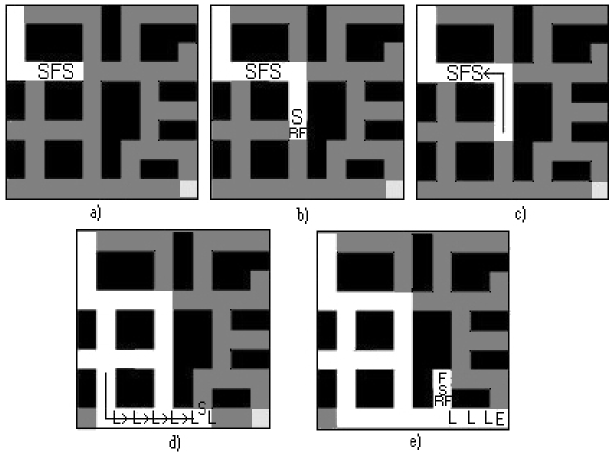

Figure 5: An Example of the agent behaviour inside the maze. The letters represent the following: “S” – food smell; “F” – food in the floor; “RF” – rotten food in the floor; “L” – light; “E” – exit of the maze. The white places are those where the agent had already been and arrows represent the direction of the path.

decisions. The stimuli representing walls are perceptually very negative (black pixels) so the directions where they come from are immediately eliminated from decision, without cognitive evaluation.

Along the way (figure $5a$)), the agent continues to avoid walls and the memory starts to be filled up with the acquired stimulus information, allowing cognitive evaluation. In figure $5a$), the agent finds smell and food for the first time but passes through them without eating due to a level of energy still acceptable. Continuing to move in the maze, in figure $5b$), the agent finds another place with smell (equal to the first one) and this time (after the execution of several actions) the energy level has been reduced under the threshold that defines the “hungry” body state. Following the smell, the agent finds rotten food, but since there is no similar stimulus in memory, there is also no cognitive evaluation. Therefore, as the perceptual evaluation results positive because of the yellow pixels, the agent eats the rotten food. The result of eating rotten food is a decrease of energy instead of an increase when the food is good. This unexpected change in the body state results in the change of the DV associated with the image of rotten food in memory to a negative DV.

Afterwards the agent is even more hungry, and since it has found food earlier in other place, it uses the memory to make a plan to return to it. The plan execution (figure $5c$)) results in finding good food and the agent chooses to eat it due to its positive DV. The body energy increases and the next need to satisfy is light again. The agent wanders in the maze searching for light and reaches the location in figure $5d$). In this part of the maze there is some light because it is near the exit, but not enough to satisfy the body need. The first choice is to follow the light but in the meanwhile the agent enters in the “hungry” state again. The agent can not execute any plan because those found in memory are too long and it has no energy to execute them. So it goes on until it finds a cross-road with smell and light stimulus in different directions. The agent chooses to move towards the smell direction[2], and finds rotten food for the second time (figure 5e)). The cognitive evaluation of the rotten food stimulus results on a negative DV retrieved from memory. Therefore it does not eat the rotten food and follows another smell towards North, leading to find good food and eat it.

After the need of energy being satisfied only the light need remains. Since light was found in the past, a plan is made to return to it. After the plan execution, the agent continues to follow light until it reaches the exit, and the simulation is over.

In a nut shell, the implemented agent is able to find the maze exit (therefore fulfill its need for light), to search for food when needed, eat when its level of energy is below a threshold, to recall positions in the maze where the agent have seen food and establish a plan to go back to those positions, and to identify and learn to avoid rotten food. The obtained behaviours emphasize the importance of the anticipation of action effects in the process of decision making.

4 The RoboCup soccer implementation

In the RoboCup (Kitano *et al.*, 1997) soccer implementation, the architecture was used in order to implement agents that can play soccer at an individual level, instead of exploring the social or multi-agent aspects of this environment. Each agent has the same underlying architecture, but the built-in components are different from each other, determining the differences between the roles that players assume in a team.

Therefore, the perceptual mapping between the relevant features of stimuli and their assessment differ according to the role of the player. This means that each agent reacts in different ways to stimuli, since the perceptual DV's will result in different preferred actions. In this implementation, the stimuli considered were the visual ones (received from the server) and consist of a set of spacial features (relative distance, direction, and so on) of the objects (ball, players, lines, flags, goals) seen by the agent in the field. The perceptual representation is composed by the built-in relevant features, which are the distance and direction of the seen objects, and its evaluation maps the values of those features into built-in desirability values. For instance, a goal-keeper will associate a small distance to the ball with a very negative DV, since it means the ball is close to its own goal, and a forward will associate the ball in the opposite goal with a very positive DV.

[2]When the two needs are unsatisfied, the energy has higher priority, in the form of a higher weight to the energy DV component.

The decision-making process relies on the anticipation of possible actions (kick, dash, turn, and so on) and parameters associated with them (direction and/or power), selecting the ones that result in more desirable situations, *i.e.* in anticipated images with the best DV.

As an example, consider the forward near the opposite goal, when it anticipates the several directions to kick, the goal direction has the maximum DV, since the ball in the opposite goal is built-in associated with a high DV. Due to the complexity and dynamic properties of the environment this anticipation had to be very fast and simple but not always accurate. This preliminary implementation considered mainly the perceptual level of the architecture using the elementary actions provided by the server. The cognitive processing relies on the representation of the seen objects in absolute coordinates, instead of relative ones, and the matching process uses all the objects in the I_c of the current stimulus and searches in memory, indexed by the I_p, for similar situations.

The results obtained in this simple implementation showed that the perceptual processing is useful for obtaining behaviours with low complexity in an efficient way. On the other hand, the degree of change in the environment made the cognitive processing, as proposed in the architecture, difficult to apply since the situations almost never matched due to the low abstraction level. New matching processes and more abstract cognitive representations are being studied in order to implement the cognitive level in this kind of environments. Also the planning and multi-agent concepts will be introduced in this implementation in order to improve the results, either by joining other theories with the emotion-based architecture or by expanding it so that it would include those concepts.

5 Related work

The discussion concerning the relevance of emotions for artificial intelligence is not new. In fact, AI researchers as Aaron Sloman (Sloman & Croucher, 1981) and Marvin Minsky (Minsky, 1988) have pointed out that a deeper study of the possible contribution of emotion to intelligence was needed. Recent publications of psychology (Goleman, 1996) and neuroscience research results (Damasio, 1994; LeDoux, 1996) suggest a relationship between emotion and rational behaviour, which has motivated an AI research increase in this area. The introduction of emotions as an attempt to improve intelligent systems has been made through different ways. Some researchers use emotions (or its underlying mechanisms) as a part of architectures with the ultimate goal of developing autonomous agents that can cope with complex dynamic environments. In this set is included the Velásquez work (Velásquez, 1998a; 1998b), who developed a pet-robot based on Damasio's ideas. Also based on Damásio's theory, an architecture for emotion-based agents has been proposed by Ventura, Pinto-Ferreira and Custódio (Ventura & Pinto-Ferreira, 1998a; 1998b; Ventura, Custódio, & Pinto-Ferreira, 1998a; 1998b; Ventura & Pinto-Ferreira, 1999). Another architecture (Tabasco) was proposed by Staller and Petta (Staller & Petta, 1998), which is based on psychological theories of emotions.

Other researchers focused their work on the adaptation aspects of emotions, using it in reinforcement learning (Gadanho & Hallam, 1998). There are researchers who defend that emotion is a side effect of an intelligent system (Sloman & Croucher, 1981), others defend the opposite, *i.e.*, emotion is the basis of emergent intelligent behaviour (Cañamero, 1997). The social role of emotion has also been explored by several researchers using it to improve societies of intelligent agents (Cañamero & de Velde, 1999; Staller & Petta, 1998). Another path of research is the one that explores the capabilities of emotion as a mean of communication in human-machine interaction (Picard, 1997), resulting in the utilization of emotions on the implementation of believable agents to interactive fiction and virtual reality (the EM architecture in the OZ project (Reilly & Bates, 1992)), synthetic characters (Moffat, 1997; Martinho & Paiva, 1999) and pedagogical agents. This research path has been based mainly on the OCC theory of emotions (Ortony, Clore, & Collins, 1988) and (or) on Frijda's work (Frijda, 1986) in this area, being both result from psychological research on emotions.

6 Conclusions and Future Work

The proposed architecture allowed the implementation of autonomous agents, (i) where the goal definition results from the agent's behaviour and needs, *i.e.*, it is not imposed or pre-defined; (ii) where the agent is capable of quickly reacting to environment changes due to the perceptual level processing; (iii) where the agent reveals adaptation capabilities due to the cognitive level processing; and finally (iv) where the agent is capable of anticipating the outcomes of its actions, allowing a more informed process of decision making.

The results of the first implementation showed the basic learning capability introduced by changing the DVs in memory, its utilization by the cognitive evaluation for object identification purposes and also the action planning for need satisfaction using DV information. The obtained behaviours also emphasize the importance of the anticipation of action effects in the process of decision making.

This architecture is now being tested in more complex and dynamic environments, namely using real robots interacting with semi-structured environments. Moreover, the maze world is being extended in order to accommodate different needs, objects and environment reactions. In a more distant horizon, the study of the social role of emotions in a multi-agent system (the RoboCup environment), the study of non-verbal reasoning mechanisms (e.g., pictorial reasoning) and its relation with the emotion processes, will be addressed.

References

Cañamero, D., and de Velde, W. V. 1999. Socially emotional: using emotions to ground social interaction. In Dautenhahn, K., ed., *Human Cognition and Social Agent Technology*. John Benjamins Publishing Company.

Cañamero, D. 1997. Modeling motivations and emotions as a basis for intelligent behaviour. In Johnson, W. L., ed., *Proc. of the First Int. Conf. on Autonomous Agents*, 148–155. ACM Press.

Damasio, A. 1994. *Descartes' Error: Emotion, Reason and the Human Brain*. Picador.

Frijda, N. 1986. *The Emotions*. Cambridge University Press.

Gadanho, S., and Hallam, J. 1998. Emotion-triggered learning for autonomous robots. In Cañamero, D., ed., *Emotional and Intelligent: the Tangled Knot of Cognition. 1998 AAAI Fall Symposium Technical Report FS-98-03*. AAAI. 31–36.

Goleman, D. 1996. *Emotional Intelligence*. Bloomsbury.

Kitano, H.; Kuniyoshi, Y.; Noda, I.; Asada, M.; Matsubara, H.; and Osawa, E. 1997. Robocup: A challenge problem for AI. *AI Magazine* 18(1):73–85.

LeDoux, J. 1996. *The Emotional Brain*. Simon and Schuster.

Martinho, C., and Paiva, A. 1999. Pathematic agents. In *Proc. of Autonomous Agents'99*. ACM Press.

Maçãs, M.; Couto, P.; Ventura, R.; Custódio, L.; and Pinto-Ferreira, C. 1999. Sociedade de agentes e inteligência sensória-motriz. Technical Report 704-99, Systems and Robotics Institute.

Minsky, M. 1988. *The society of mind*. Touchstone.

Moffat, R. 1997. Personality parameters and programs. In *Creating personality for synthetic actors*. Springer.

Ortony, A.; Clore, G. L.; and Collins, A. 1988. *The Cognitive Struture of Emotions*. Cambridge University Press.

Picard, R. 1997. *Affective Computing*. MIT Press.

Reilly, W. S., and Bates, J. 1992. Building emotional agents. Technical Report CMU-CS-92-143, Carnegie Mellon University, School of Computer Science.

Sloman, A., and Croucher, M. 1981. Why robots will have emotions. In *Proc. 7th Int. Joint Conference on AI*, 197–202.

Staller, A., and Petta, P. 1998. Towards a tractable appraisal-based architecture. In Cañamero, D.; Numaoka, C.; and Petta, P., eds., *Workshop: Grounding Emotions in Adaptive Systems Int. Conf. of Simulation of Adaptive Behaviour, from Animals to Animats, SAB'98*. 56–61.

Velásquez, J. 1998a. Modeling emotion-based decision-making. In Cañamero, D., ed., *Emotional and Intelligent: the Tangled Knot of Cognition. 1998 AAAI Fall Symposium Technical Report FS-98-03*. AAAI. 164–169.

Velásquez, J. 1998b. When robots wheep: Emotional memories and decision-making. In *Proc. of AAAI-98*, 70–75. AAAI.

Ventura, R., and Pinto-Ferreira, C. 1998a. Emotion-based agents. In *Proc. of AAAI-98*, 1204. AAAI.

Ventura, R., and Pinto-Ferreira, C. 1998b. Meaning engines - revisiting chinese room. In Cañamero, D.; Numaoka, C.; and Petta, P., eds., *Workshop: Grounding Emotions in Adaptive Systems Int. Conf. of Simulation of Adaptive Behaviour, from Animals to Animats, SAB'98*. 69–70.

Ventura, R., and Pinto-Ferreira, C. 1999. Emotion based agents: Three approaches to implementation. In Velásquez, J., ed., *Workshop on Emotion-based Agents Architectures, EBAA'99*. 121–129.

Ventura, R.; Custódio, L.; and Pinto-Ferreira, C. 1998a. Artificial emotions - goodbye mr. spock! In *Proc. of the 2nd Int. Conf. on Cognitive Science*, 938–941.

Ventura, R.; Custódio, L.; and Pinto-Ferreira, C. 1998b. Emotions - the missing link? In Cañamero, D., ed., *Emotional and Intelligent: the Tangled Knot of Cognition. 1998 AAAI Fall Symposium Technical Report FS-98-03*. AAAI. 170–175.