

Emotional Valence-based Mechanisms And Agent Personality

Eugénio Oliveira¹ and Luís Sarmiento^{1,2}

¹NIADR - Faculdade de Engenharia Universidade do Porto
Rua Dr. Roberto Frias, s/n Lab. I 121, 4200-465 Porto, Portugal

²Escola das Artes – Dep. Som e Imagem - Universidade Católica Portuguesa C.R.P
Rua Diogo Botelho 1327, 4169-005 Porto, Portugal
eco@fe.up.pt, las@porto.ucp.pt

Abstract. Artificial Intelligence is once again emerging from a pragmatic cycle and entering a more ambitious and challenging stage of development. Although the study of emotion in the realm of Artificial Intelligence is not totally new (Simon, Minsky and Sloman and Croucher), much more attention has been recently devoted to this subject by several researchers (Picard, Velasquez, Wright). This renewed effort is being motivated by trends in neuroscience (Damásio, LeDoux) that are helping to clarify and to establish new connections between high level cognitive processes, such as memory and reasoning, and emotional processes. These recent studies point out the fundamental role of emotion in intelligent behavior and decision-making. This paper describes an on-going work that intends to develop a practical understanding of models backing those solutions and aims at their integration in Agent Architectures, having always in mind the enhancement of agents' deliberation capabilities in dynamic worlds.

1 Introduction

In the Artificial Intelligence field, the role of emotion in cognitive processing has been acknowledged since the late sixties by Herbert Simon [13]. Nevertheless, during the following 25 years, few were the researchers from the AI field that adventured themselves in the study of Emotion. Some notable exceptions are Marvin Minsky [7] and Aaron Sloman [14]. Recently, the work of the neuroscientist António Damásio [3] established a clear relationship between specific brain structures and emotional capabilities. Damásio's studies on his patients allowed the identification of specific brain regions (pre-frontal cortexes) that, whenever affected, would render the patient unable to respond to emotionally rich stimulus (e.g. violent or sexual content images). At the same time, these patients revealed significant difficulties in dealing with several real life situations, especially when confronted with the need to perform decisions either on a personal, social or professional level. However, in these cases, patients still keep their mathematics and speech skills intact, as well as their memory. Their performance in IQ tests remains normal and most of the times their problem is unnoticeable. Damásio's results suggest that highly cognitive tasks such as risk assessment and decision-making are somehow related to emotional processing and

that this relation is actually supported by neuronal structures. Evidence of a biological support for emotion-cognition relationship seems an extremely significant result, bringing some light over the original ideas of Simon, Minsky and Sloman.

Inspired by Damasio's work and following the work of several other researchers [11], [16], [17], [18], we started a project with the aim of endowing intelligent agents with the possibility to use an emotion-based mechanisms that strongly influence their own decision-making capabilities [15]. Furthermore, we are interested in studying how such emotion-based mechanisms can be manipulated and tuned to create different individual Agents based on the same Architecture. These Agents can be said to have distinct Personalities that could reveal themselves to be more advantageous for pursuing their own specific goals in specific environment conditions. The structure of this paper includes, besides this first introductory section, an introduction to the concept of Emotional Valence, which is at the core of our Architecture. We will then present our model of emotional mechanisms and establish its relationship with the other elements of the agent architecture. We will then try to show how emotional mechanisms, such as those we proposed, might be used to promote intelligent behavior. Finally we will present our current implementation of the architecture.

2 The Role of Emotion - Valence

Emotion is a highly complex multi-faceted phenomenon. Consequently, researchers from several fields have developed deep insights on the study of that concept. Thus, depending on the original field of the researchers the focus of the study could vary immensely. For an interesting survey about several emotion issues refer to [4]. From our perspective, as engineers and computer scientists, we are mostly interested in studying the functional aspects of emotional processes. Particularly, we aim to understand how emotional mechanisms can improve cognitive abilities, such as planning, learning and decision-making, for hardware and software Agents. We hope to develop more flexible Agent Architectures capable of dealing with highly complex, rapidly changing and uncertain environments. In a certain way, we are following the complementary direction of the work done by A. Ortony, A. Collins and G. Clore that lead to the well-known OCC model [10]. The OCC model is mainly focused on explaining "the contribution that cognition makes to emotion". The work presented in [10] discusses the cognitive processes that generate the appropriate conditions leading to given emotional states (*eliciting conditions*). We, on the other hand, seek to explore the functionality of emotional states to increase the performance of an Artificial Agent interacting with complex environments.

From a functional point of view there are several issues about emotion that we found useful to investigate and to work upon. The most fundamental functionality of emotion concerns state evaluation. In this context, Emotions can be regarded as a built-in mechanism able to provide *automatic and rapid evaluations of environment conditions together with its own internal state*. In particular, for a given Agent, with a defined set of goals and capabilities to change its environment, emotions are used to identify the *valence of the environment and of its own capabilities*. We define *valence* as a subjective measure that relates the chances of an Agent being able to fulfill its goals given a particular environment situation, its internal state and its set capabilities. Valence may be positive if environment conditions and internal state of the agent are

favorable to goal achievement, or negative otherwise. An important point to stress here about valence concept, is that agents evaluate environment not just “per se”, but according to their own current goals and motivations. Based on the emotional capability we propose several other features that we believe may be advantageous for more sophisticated Agent Architectures: (1) Valence-based Long-Term Memory, (2) Emotional Alerts, (3) Action Tendency.

In the next sections we will address all these issues in detail. We will also address the possibility of exploring variations over several Emotional parameters. In fact, despite the same internal Emotional-based Architecture, we will demonstrate that two Agents may show different behavior in the same situation, reflecting the existence of two distinct Agent Personalities and two different past histories.

3 Emotional Valence-Based Mechanisms

As mentioned in the last section, Emotions provide an automatic and quick way of evaluating the environment and the internal state of the Agent in respect to its own goals. It is important to stress that this evaluation is twofold. Firstly, it reflects the outside environment conditions by providing a valence tag to the information gathered by the perception subsystems. For example, emotional mechanisms may alert to a particular outside situation that influences critically an agent goal (an extremely negative or positive valence situation) and, therefore, requires special treatment. In this context, emotional mechanisms will try to quickly answer questions such: “How good are environment conditions to my specific goal(s)?”.

Secondly, and also related with environment evaluation, Emotions do also reflect the fitness of the Agent to cope with specific environment states [5], [8], [9]. In particular, Emotions will valence Agent’s own action set, current plans and knowledge regarding their effect on goal achievement in a given environment. Up to a certain point, this process of internal evaluation can be regarded as a basic introspective activity. Emotional mechanisms will indirectly try to answer questions such as “How fit are these plans to help me achieving my goal(s)?” or “How useful has been my knowledge in my last decisions?”.

4 Valence Functions, Accumulators and Memory Thresholds

In this section we will introduce a model of the emotional mechanisms. As we have described above, these mechanisms should receive input from internal sources, I , as well external sources, E , and produce a valence measure, V , according to what we will call *Emotional Valence Function*, EVF . Emotional Valence Functions return the valence of the situation regarding a given goal, G : $V = EVF(I,E,G)$.

An Emotional Valence Function is supposed to be a fast mechanism and therefore it should easily computed. However, it is also possible conceive some higher level EVF ’s dealing with complex inputs, such social beliefs, as long as their computation does interfere in the ability of the Agent to respond in real-time to the environment. Emotional Valence Functions can be further decomposed in a Normalized Valence

4 Eugénio Oliveira and Luís Sarmento

Function NEVF, whose values range from -1 to 1, and a Sensibility Factor S . Thus: $V = EVF_i(I,E,G) = S_i \times NEVF_i(I,E,G)$.

The Valence value returned by EVF is then used to update the agent internal state. For each EVF the agent keeps an *Emotional Accumulator* to which the Valence values are added. Emotional Accumulators exhibit a time dependent behavior. Their values decay with the passing of time, at a given *Decay Rate* (DA_i). This behavior is similar to the dynamics shown by emotions in people.

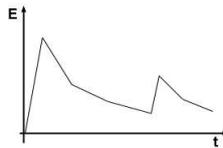


Figure 1 - Profile of an Emotional Accumulator. The rises of the curve represent (positive) updates from EVF. The value of the Accumulator decreases at a given decay each time slot.

Emotional Accumulators are fundamental elements of the internal state of the agent, and have, as shown later, direct influence on all deliberative and reactive processes. Furthermore, let's assume that the valence measure and the sources of evaluation are then associated in order to form a Valence Vector: $\langle V, I, E, G \rangle$. Valence Vectors are stored in the working memory and made available for all Agent processes for further consideration. Valence Vectors may be stored afterwards in long-term memory by a dedicated process that selects specific vectors according to their relevance. For each EVF_k , let us define MT_k as the Memory Threshold level. Then a specific Valence Vector $\langle V_j, I, E, G \rangle$ is selected to be stored in long-term memory if: $|V_j| = |EVF_k(I, E, G)| > MT_k$.

Valence Vectors that have higher valence magnitude than the corresponding MT_k , can be seen as particularly relevant and should be stored for later processing while others may be simply discarded. We will explore this issue in the following section.

In summary, for each of its goals (explicit or implicit), an Agent will have one Emotional Valence Function, the corresponding Emotional Accumulators and Memory Thresholds. Figure 2 tries to depict what we have just described.

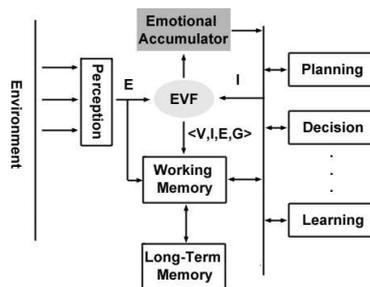


Figure 2 – Emotional Valence Function and its relationship within Agent Architecture

5 Valence-Based Long-Term Memory

By combining all the Valence Vectors that are being produced during its interaction with the environment, an Agent is able to create contextual memory maps of its past experiences. As we have seen before, Emotional Valence Functions and Memory Threshold levels allow the Agent to select which data is worth storing in long-term memory. Having in mind the purpose of Emotional Valence Function, we can say that highly valenced data is related either with good goal achievement perspectives (positive valence) or dangerous threats to specific goals (negative valence). Therefore, this selection process retains only the information that is considered particularly valuable to the goals of an Agent, while discarding less relevant, although probably much more abundant, information. Additionally, valenced long-term memory may help the search for pre-existing plans and facts. Long-term memory may be indexed by valence and then searched in an informed way. Contextually relevant information may be automatically transferred to working memory where more complex processing can be performed. For example, when facing a situation with a given calculated valence, all information coherent with that valence assessment can be decisive. Plans and facts used in situations with similar valence present a high probability of being reused or excluded according to the result (either positive or negative) they have achieved previously. Thus, the search for appropriate behaviors over a knowledge base can be pruned right from the beginning.

There is a certain similarity between the mechanism we have just described and Case Based Reasoning, although some important differences can be noted. Besides being much more simple than the overall CBR cycle [1], Valence-based Memory uses an Agent centered measure to choose which cases are to be retained: Valence. Thus, cases retained depend much more on the current performance of the Agent than on certain metrics defined during the design stage. Moreover, Valence-based Memory is not intended to store extensively all possible cases, which would be an inappropriate procedure considering the real-time demands of target environments and the memory limitations of Agents. As more recent Valence Vectors are computed, older or less significant ones may be “forgotten” so that the stored knowledge can be refreshed.

We will continue to work on this particular subject in order to develop a deeper understanding.

6 Emotion-driven agent behaviors

One important feature of emotional processes is the immediate and intuitive recognition of critical situations, which are supposed to be reflected by strong valence assessments and high Accumulator levels. These emotional evaluations may generate new motives to change current behavior or even change Agent capabilities.

Alerts: Emotional Valence mechanisms may be useful in detecting situations that may interfere (positively or not) in goal achievement, alerting the Agent’s internal processes for relevant events that may demand attention [8], [9]. When facing such events and situations, which in complex environments may not always be easily identified or expressed by beliefs, emotional alerts should drive the agent to focus on important data. This alert and focusing will motivate the agent to eventually start

classification or pattern recognition procedures and then search for appropriate actions. Emotional Accumulators, for example, may help the agent to detect situations that, although their instant valence is not particularly relevant, remain active for long periods.

Tendencies: However, more than just alerting and starting other processes, emotional mechanisms may also directly contribute to Agent response, by creating a specific internal context. As our own body and senses get prepared by the effect of fear to respond effectively (quickly or not) to a possible harmful situation (in this case, the goal is keeping physical integrity), Agents behavior at deliberative and reactive layers may also suffer similar alterations. Thus, emotion can be regarded as a mechanism capable of creating *action tendencies*. For example, emotions may be responsible for plan pre-selection, by offering deliberative layers a set of “experience tested” plans or rules. Although this first selection may eventually leave out the best choice, it does also contribute to reducing the work of deliberative layers that are then able to respond much more promptly, a condition that is usually essential in survival. In this sense, emotions contribute positively to the notion of Bounded Rational Agent [12] by allowing the Agent to behave as well as possible given its limited resources and complex environment conditions.

Moods: Emotional interference in action guidance can also be done in larger time spans. If we consider slower effect emotions that reflect themselves not in immediate actions but in new goals adoption we may be able to devise a long-term adaptation mechanism. These slower effect emotions, which remain active for longer time periods, may be seen as *moods* and their influence upon Agents is made at a higher level, namely in goal adoption. They also can be of great help in filtering current possible options selecting those which are in agreement with long term policies.

7 Testing Ideas

We are currently developing software simulations based on the platform RealTimeBattle (RTB), which is available at <http://realtimebattle.sourceforge.net>. This platform provides a simulated real-time environment where softbots fight for survival in dynamic scenarios. RTB allows the developer to program their own softbots in C/C++, as well as to create custom 2D scenarios. Simple physical properties (air resistance, friction, material hardness) are also implemented to enrich the simulation. Softbots perception is basically a set of radar events from which they can detect walls, other softbots, shots and randomly distributed energy sources and mines. Softbots can accelerate, break, rotate and shoot in a given direction. RealTimeBattle assures that Softbots have limited processor time but it demands real time response from the softbots. In this way RealTimeBattle platform seems appropriate to test some of the ideas described before. Our current Emotional Agent Architecture comprises 3 different layers. Each layer provides a set of capabilities that can be used by upper layers. Each layer has also included a set of simple Emotional Valence Functions and Accumulators are intended to reflect the success of the Agent in achieving an explicit or implicit goal.

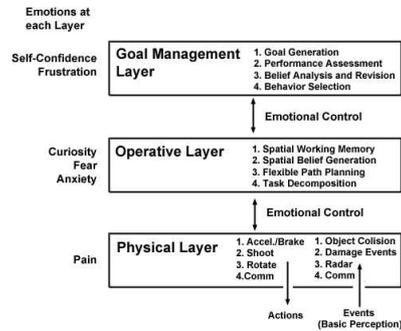


Figure 3 – Layered architecture. Emotional parameters generated at each layer are shown in the left side.

The bottom layer is the Physical Layer and is highly domain dependent. In the case of RealTimeBattle it includes the softbot sensing capabilities and all the low-level action and communication mechanisms. At this level, the robot is capable of providing simple reactive responses to the environment. For, example, the softbot can shoot a close mine, without any further consideration. In the physical layer we have included one Emotional Valence Function and the corresponding Accumulator whose objective is to measure the aggressiveness of the environment. Our purpose is mimicking the function of pain in animals. Pain is deeply and directly related with the goal of survival and physical health. Thus, damaging events, as shoot and mine collisions, are internally reflected by high values of the ‘Pain’ EVF and increases in the ‘Pain’ Accumulator. At the physical layer level these values will be reflected in some internal parameters such as the power used by the softbot when shooting or the speed of its moves. Thus, for each action A_{c_j} we have a set of preconditions P which include EVF and Accumulator (Acc) values: $P(A_{c_j}) = \{\{EVF\},\{Acc\}\}$. The action itself is also function of a given set of EVF’s and Accumulators: $A_{c_j}(\{EVF\},\{Acc\})$. At upper layers other effects are also felt, but usually in an indirect way, as we will see.

The next layer, the Operative Layer, is responsible for more complex capabilities. It receives sensor data from the Physical Layer, and analyses it in order to construct a map representation of the environment and also to track the position of other robots, mines and cookies. This layer also keeps a record of the location where it was inflicted pain. The Operative Layer also provides path-planning capabilities. The softbot is capable of planning its path to a given destination given its knowledge of the environment collected through sensing. One particular issue about this planning capability is the possibility of controlling two different parameters: the number of steps in the path and their length. This allows the softbot to choose between simple and quicker plans or elaborate, but eventually slower, plans. These parameters will be subject to the influence of another EVF/Accumulator that represents an emotion similar to Anxiety. High the values of the ‘Anxiety’ Accumulator will result in the creation of shorter plan that allow the agent to respond promptly to a given situation. The EVF of the ‘Anxiety’ accumulator uses several input parameters, in which are included the value of other Emotional Accumulators such as ‘Fear’, ‘Curiosity’ and ‘Pain’. Since ‘Fear’, ‘Curiosity’ and ‘Pain’ depend themselves on several other dynamic, time-varying, parameters (see Figure 4), it can be seen that the resulting

structure is very complex and would be difficult to implement in the form of IF-THEN rules.

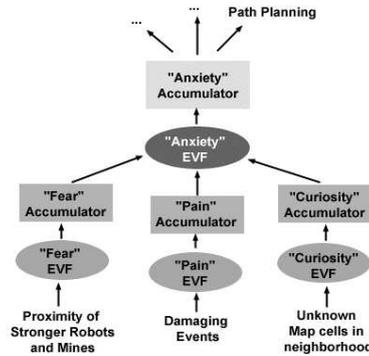


Figure 4 – Relationship between different emotional parameters from different layers.

The upper layer, called Goal Management Layer (see figure 3), is still under implementation. Its purpose is to manage all agent high-level behaviors. It is responsible for generating goals and sub-goals and tracking their execution. The goal generation process will also be dependent on the value of EVF and Accumulators. For example, the Accumulator “Curiosity” may contribute to the generation of a goal such as ‘Explore Surroundings’. This will then motivate an exploring behavior with several lower-level operative actions (look around, move to an unknown point in map). At this level we propose two different emotional dimensions related with global performance: “Self-Confidence” and “Frustration”. “Self-Confidence” should increase when the softbot is regularly achieving its goals. It will be reflected in the way softbot deals with difficult situation, such as those that are related with high levels of ‘Fear’. High levels of “Self-Confidence” will make the softbot adopt a more active behavior, such as attacking or hunting other robots. On the other hand, low levels of “Self-Confidence” accumulator will promote behaviors such as running away or hiding from enemies. Note that this behavior appears to provide a natural form of adaptation, increasing chances of survival. On the other side, “Frustration” should reflect the inadequacy of current behaviors to achieve given goals. It will indicate the softbot that a change of behavior or goal is needed. At this layer emotional mechanisms are essentially related with introspective activities.

8 Agent Personality And Evolutionary Agent Design

Within the same Emotion-based Architecture, that includes specific EVF’s, Emotional Accumulators and Memory Thresholds there are several different parameter variations that can be seen as distinct Agent Personalities. Although “Agent Personality” is certainly a difficult concept to define precisely, we may say that it is what distinguishes similar Agents (i.e. with the same Architecture) regarding their patterns of behavior. In this perspective, assuming that, for example, we change the intensity of how Emotional Accumulators interact with each other, we can expect the overall behavior of the Agent to change because of their intense relations with all the

Agent processes. In our Architecture the intensity of those interactions is ultimately controlled by the EVF's. Therefore, EVF parameters may be considered, in a rather simplified way, as part of the Personality of the Agent. An Agent Ag_i personality, indirectly governing an Agent Ag_i behavior, can then be described as the complete set of its Emotional Valence Functions and corresponding Accumulators and Memory Thresholds: $Personality_i = \{EVF_k, Ac_k, MT_k, \text{ for all } Em_k \in \{Em\} \text{ (the set of emotions) and } Ag_i \in \{Ag\} \text{ (the set of Agents)}\}$.

Let explore, for example, the Sensibility of the EVF_k . Despite the similarity of their overall internal structure, two Agents, Ag_r and Ag_s , will tend to behave differently if they have different Sensibilities factors regarding the corresponding EVF's: $Ag_r : EVF_{rk} = Srk * NEVF_k$ and $Ag_s : EVF_{rs} = Srs * NEVF_k$.

Higher sensibilities will naturally motivate the Agent to respond more quickly to a given environment stimulus. The Agent should therefore, in these cases, look more nervous and will probably change its behavior more abruptly.

We can broaden the concept of Agent Personality by also manipulating the Decay Rate of Emotional Accumulators (refer to Figure 1). Decay Rates are related with behavior stability. Slower decay rates will increase the stability of Agent's internal state, making it less dependent from environment changes. Agents with slow Decay Rates will be influenced by environment stimulus for longer periods. On the other hand faster decay rates will make the Agent surpass environment stimulus quickly.

These possibilities suggest an opportunity for tuning emotional parameters for better agent performance. Since each individual Agent has a particular set of emotional parameters, which comprise the Sensibility Factors (S) of EVF's, the Decay Rate of Emotional Accumulator (DA) and Memory Threshold Levels (MT), we may admit there exists a specific combination of these parameters that optimizes Agent performance in a given environment. This combination would be the Optimal Agent Personality. For a given Emotion-based Architecture, we shall define the Personality Set Domain (PSD) as being the set of all possible combinations of Sensibility factors, Decay rates and Memory Thresholds:

$$PSD = \{S_1\} \times \{S_2\} \dots \times \{S_n\} \times \{DA_1\} \times \{DA_2\} \dots \times \{DA_n\} \times \{MT_1\} \times \{MT_2\} \dots \times \{MT_n\}.$$

Therefore, the PSD includes every possible Agent Personality within a specific Emotion-based Architecture. Finding the Optimal Agent Personality for a specific environment can be seen as a search problem over the PSD space. From a system designer point of view this suggests an evolutionary approach to the development of Emotion-based Agents, releasing to burden of finding the best Sensibility Factors, Decay Rates or even Memory Thresholds manually. The designer would perform a search for the optimal Agent Personality by varying these parameters around some reasonable initial values over several rounds of simulations. The parameters that yield the best Agents in respect to a certain performance criteria in a specific environment would then be selected as the Optimal Agent Personality for that environment.

Note that this search process does not reduce the ability of an Agent to cope with environmental changes. It is mainly a design method to help the developer automatically tune some or the available parameters. To cope with environment changes that happen during the "lifespan" of an Agent, the proposed Architecture includes other mechanisms located at the "Goal Management Layer". Namely, both "Frustration" and "Self-Confidence" emotional mechanisms try to regulate the behavior of an Agent in order to promote adaptation (e.g. belief revision).

9 Conclusions

In this paper we have proposed an Emotional based Agent Architecture intended for Agents that operate in complex and real-time environments. Particularly we have been mainly concentrated on Emotional Valence Functions, which are mechanism that make possible for an agent to perform a fast evaluation of external and internal states regarding its chances of achieving its own goals. We have also showed how emotion-based processes could be used to direct deliberative agent processes, such as decision-making and planning. We have also discussed the possibility of exploring variations on such emotional mechanisms and its relation with the concept of Agent Personality.

References

1. A. Aamodt, E. Plaza. Case-Based Reasoning: Foundational Issues, Methodological Variations, and System Approaches. AI Communications. IOS Press, Vol. 7:1, pp. 39-59.
2. L. Custódio, R. Ventura, C. Pinto-Ferreira. Artificial Emotions and Emotion-Based Control Systems. Proc. 7th IEEE Int. Conf. Emerging Technologies and Factory Automation. 1999.
3. A. Damásio. Descartes Error - Emotion, Reason and the Human Brain. 1994.
4. P. Ekman and R. Davidson (Editors). The Nature of Emotion. – Fundamental Questions. Oxford University Press, 1994.
5. Nico Frijda. Emotions are functional, most of the time. In P. Ekman and R. Davidson (Editors). The Nature of Emotion. – Fundamental Questions. Oxford University Press, 1994.
6. J. LeDoux. The Emotional Brain: the Mysterious Underpinnings of Emotional Life. 1996.
7. M. Minsky. The Society of Mind. First Touchstone Edition. 1988.
8. D. Moffat, N. Frijda. Functional Models of Emotion. In G. Hatano, N. Okada and H. Tanabe (Ed.), 13th Toyota Conference-Affective-Minds, pgs 169-181. Elsevier, Amsterdam, 2000.
9. D. Moffat. Rationalities of Emotion. To appear.
10. Ortony, G.L. Clore, A. Collins. The Cognitive Structure of Emotions. New York: Cambridge University Press.
11. R. Picard. Affective Computing. The MIT Press, 1997
12. S. Russel, P. Norvig. Artificial Intelligence: A Modern Approach . Prentice Hall. 1995.
13. H. Simon. Motivational and emotional controls of cognition. Psychological Rev., 74. 1967.
14. A. Sloman and M. Croucher. Why Robots will have emotions. In Proc. 7th Int. Joint Conference on AI, Vancouver 1981.
15. A. Sloman. Beyond Shallow Models of Emotion. In "Cognitive Processing", Vol. I, 2001.
16. J. Velásquez. Modeling Emotion-Based decision making. In Dolores Cañanero (Editor), Emotional and Intelligent: The Tangled Knot of Cognition, pages 164-169, 1998.
17. R. Ventura. Emotion-Based Agents. Msc. Thesis. Instituto Superior Técnico, Lisboa, Portugal, 2000.
18. I. Wright. Emotional Agents. Ph.D. thesis, School of Computer Science, The University of Birmingham, 1997 (<http://www.cs.bham.ac.uk/research/cogaff/>)