An Emotion-based Agent Architecture Application with Real Robots[*]

Rui Sadio and **Gonçalo Tavares** and **Rodrigo Ventura** and **Luis Custódio**
Institute for Systems and Robotics — Instituto Superior Técnico
Lisbon, Portugal
lmmc@isr.ist.utl.pt

Abstract

In this paper, an architecture for an emotion-based agent and its application to a real robot, situated in a semi-structured environment, is presented. This architecture corresponds to an improved version of the one proposed by Maçãs *et. al.*. Both were developed based on a particular interpretation of the neuro-physiological findings of Damásio and LeDoux, namely the concepts of stimuli parallel processing, by LeDoux, and somatic marking, by Damásio. Given several applications of this architecture with virtual agents and in simulated environments, the goal here was to study and evaluate the utility and efficiency of this emotion-based agent architecture with real robots and real environments.

Introduction

The study of the importance and influence of emotion in human cognitive processes has been under active research, particularly, in the neuroscience and psychology areas, and recent results suggest that human emotions play in those processes a role far more significant than what one could anticipate. Nevertheless, besides these studies being far away from conclusive results, it is not straightforward to establish how the notion of emotion can be implemented in (virtual or real) robots and what can be gained by doing this, in terms of robot behaviours.

The work presented in this paper follows previous works presented in (Ventura & Pinto-Ferreira 1998), (Ventura & Pinto-Ferreira 1999), and (Custódio, Ventura, & Pinto-Ferreira 1999). More recently, Marcia *at. al.* have proposed an emotion-based agent architecture called DARE (Maçãs *et al.* 2001a). In (Vale & Custódio 2001) the learning and generalization functionalities of the DARE architecture were discussed. All these works are based on two main concepts: the double and parallel stimuli processing concept, proposed by LeDoux (LeDoux 1996), and the somatic marker concept, introduced by Damásio (Damásio 1994).

The sensory information acquired using the agent sensors is processed under a double perspective: a perceptual, immediate one, which allows the agent to quickly react to urgent situations, and a cognitive, elaborate one, which allows the agent to identify what is seeing given what it already knows from previous experiences. At the perceptual level, the information extracted from the stimulus is simple, basic and easily handled based on a set of built-in characteristics, which provides a fast, but rough, assessment of the stimulus (e.g., is it positive/negative, desirable/avoidable, relevant/irrelevant, urgent/not urgent). At the cognitive level, a more complex, rich, divisible, structured and hardly handled representation is used based on all information extractable from visual, audio and other sensors. This sophisticated representation should provide a comprehensive identification of the stimulus by the higher level cognitive systems.

Given these two informations, the DARE architecture incorporates a body state representation and a marking mechanism, which are used to (somatic) mark a stimulus acquired from the sensors and, afterwards, stored in memory, together with a "good-or-bad" evaluation. Therefore, an emotion-based agent is, in this work, an entity whose behaviour is guided by taking into account a rough evaluation of the stimulus goodness or badness, an stimulus identification based on past experiences, and a somatic marking mechanism which allows to recall the impact of past similar stimuli into the agent body state. This mechanism can be related with the concept of secondary emotions, as defined by Damasio (Damásio 1994), and used to anticipate action outcomes and desirability (Maçãs *et al.* 2001b).

The work presented here introduces new concepts and modules in the architecture (namely, the idea of "background feelings" — an interpretation of the homonymous concept introduced by Damásio (Damásio 1994) (Damásio 1999), a set of behaviours to complement the set of primitive actions available, the idea of motivation — a specific body state induces a particular behaviour, and a module of spatial evaluation based on sonar information. Given several applications of the DARE architecture with virtual agents and in simulated environments, the goal here was to study and evaluate the utility and efficiency of this improved emotion-based agent architecture with a real robot in a real environment.

Environment

In this implementation, the robotic agent moves in a micro-world defined as a closed area with some obstacles (boxes)

[*]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.

inside it (Figure 1). The main purpose of the agent is to keep its energy and strength internal levels above predefined (instability) thresholds. In order to do so, there are colourful signs, distributed along the micro-world, indicating either presence of food or a place to rest (blue sign). Agent decisions are taken according to its internal needs, in response to external stimuli. When energy and strength levels show values above the thresholds, the agent just rambles in the micro-world and it plays with an orange ball, in case it finds it. If one of those levels shows a value below the threshold, the agent ought to satisfy that need, by looking for food or some place to rest.

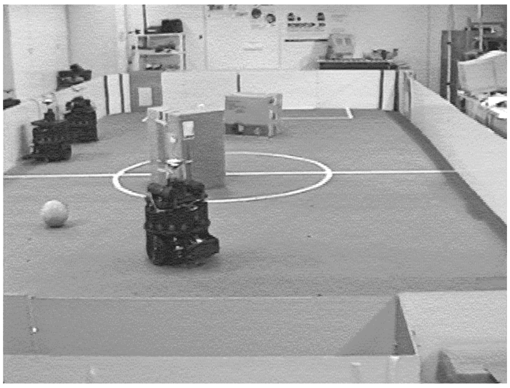

Figure 1: The world and the robot used for experimentation.

In this world, there are two kinds of food: good food (represented by an yellow and turquoise sign), which increments agent internal level of energy, and rotten food (represented by an yellow and rose sign), that does not produce any change in the agent body state. The idea is to force the agent to learn, by experience, which food sign it must choose when it has to eat.

The robot used in this application was the 'Nomad Scout II' model from Nomadic Technologies. Its most important characteristics include: a ring of 16 sonars around it, an odometric system and a CCD camera.

Agent Architecture

As mentioned before, the agent architecture described here is based on the one presented in (Maçãs *et al.* 2001a). Figure 2 presents the new proposed architecture.

This architecture incorporates the following main modules: i) Stimulus Processing, ii) Innate Part, iii) Stimulus Internal Image, iv) Memory, v) Internal Stimulus Evaluation, vi) and Behaviour Anticipation and Selection.

Stimulus Processing

The process starts with the stimulus processing stage. External stimuli are composed by three components that result from the three sensors the agent has: the visual image, provided by the CCD camera, the information given by the ring of sonars and the odometric data.

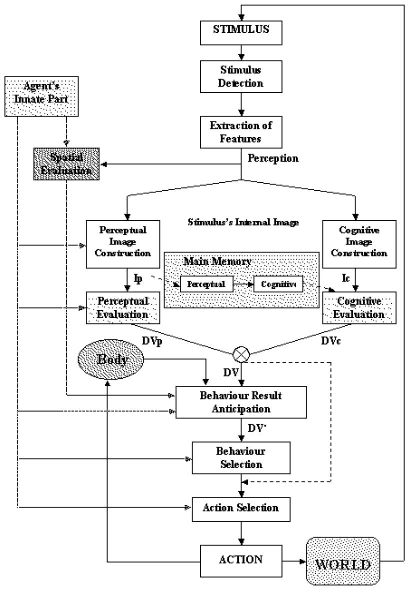

Figure 2: The emotion-based agent architecture.

Agent Innate Part

In order to bootstrap the robotic agent, some information have to be provided beforehand. This predefined (innate) information includes a set of relevant features, a set of innate meanings, one for each feature, a set of equilibrium values for the body state components, and a set of predefined behaviours.

Innate Relevant Features This set of features is used by the perceptual level of the architecture to extract basic, but relevant, information from the stimulus. In this application, a relevant feature is a color in the following set $\{yellow, blue, orange, rose, turquoise\}$ and the minimum distance to an obstacle given from the three frontal sonars.

Innate Meanings The perceptual meaning of a stimulus is established by associating an innate meaning with each one of the extracted relevant features — colors meaning and spatial information (distance) meaning. An amount of yellow and rose (or turquoise) in an image indicates the presence of food. Similarly, an amount of blue or orange indicates, respectively, the presence of a place to rest or a ball to play.

Concerning the minimum frontal distance, if it is below a predefined threshold — the contact distance — then it means the agent is in contact with an obstacle in front of it. A distance value lower than another predefined threshold — the security distance (greater than the contact distance) — means that there is an obstacle near the robot. According to the information provided by the ring of sonars, the obstacle direction may be: forward, backwards, first quadrant, second quadrant, third quadrant or fourth quadrant. If there is an obstacle in front, or in the first or second quadrant, and the amount of a certain colour, in the image, is greater than a predefined threshold, then the robot is considered to be close to an object of that colour.

Body State Innate Tendencies The agent body state has two components: energy and strength. Agent innate tendency is defined as an homeostatic vector (HV), with a similar structure. The HV vector can be seen as the optimal body state and its values remain constant all the time. When the execution is started, its body state is stable, that is to say, the HV and the body state components have equal values. As the robot moves around the world, the energy and strength of its body state decrease. If one of these components shows a value below the instability threshold then the agent is facing a basic need, hunger or tiredness, which the agent ought to satisfy as soon as possible.

By taking advantage of energy and strength variations, the concept of the background feeling, introduced by António Damásio (Damásio 1994), has been incorporated in this architecture. Damásio has stated that background feelings come from background emotions. Although these emotions are directed to the body inside, one might observe them, outside, through a behavioural point of view. In order to make use of this concept, in this particular application, four background feelings were defined: enthusiasm, apathy, vigour and fatigue. Figure 3 shows how these feelings are established based on the current levels of energy and strength.

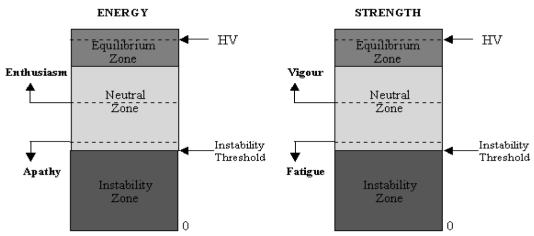

Figure 3: Background feelings definition.

The current background feeling has an influence upon the robot movement speed: the robot moves faster when it has simultaneously background feelings of enthusiasm and vigour, it moves slower when it has a background feeling of apathy or fatigue (inclusive or), and it moves with an intermediate velocity otherwise.

Another concept introduced in this architecture was the concept of motivation. A motivation is just a specific body state that leads to a particular behaviour. For this application three motivations were defined: hunger and tiredness, which lead the agent to eat and to rest, respectively, and the ”wish to play” motivation, which leads the agent to play with an orange object (e.g., a ball), as soon as it sees it. This motivation is only active when the agent is not hunger or tired, *i.e.*, the former motivations have a higher priority than the latter one.

Behaviours A behaviour is a set of elementary actions performed to achieve a certain goal. The agent may choose behaviours from the following list:

- **Approach** The agent moves forward towards a sign;
- **Execute** The agent eats or rests when it is close enough to a sign;
- **Play** The agent thrusts the ball forward;
- **Deviate** The agent moves forward turning to the left (or right);
- **Avoid** The agent stops and turns around;
- **Plan** The agent tries to reach a place (sign), where it have already been in the past;
- **Ramble** The agent just keeps moving on, randomly;

Stimulus Internal Image

After extracting features from a stimulus a perception is created and it will be used for internal image construction process. Since there are two parallel stimulus processing levels, there are, also, two distinct representations: a perceptual image (I_p) and a cognitive image (I_c). These images are defined aiming at the evaluation of the stimulus both in a perceptual and cognitive levels.

The perceptual image is defined based on the set of relevant features. Thus, the perceptual image has six components whose contents are the number of pixels from each recognized colour and the minimum frontal distance, in meters. On the other hand, the cognitive image is just the complete visual image acquired from the CCD camera.

Internal Stimulus Evaluation

After constructing stimulus internal image, a Desirability Vector (DV), associated with the perceptual and cognitive images, is determined. The DV structure consists of two boolean components: the first one referring to pain and the second one referring to pleasure, both caused by external stimuli. The pleasure component results from a logic sum of three fields in which it is subdivided: pleasure associated with the presence of food or a place to rest or a ball, whereas the pain component will be one if the robot is in contact with an obstacle

Perceptual Evaluation The perceptual evaluation makes use of perceptual innate meanings to estimate the perceptual DV (DV_p). If the perceptual image component corresponding to blue (or orange) color has a non-zero value, the DV pleasure value associated with “rest” (or the ball) is set to one. In the same way, if the perceptual image components corresponding to yellow and rose (or turquoise) have a non-zero value, the DV pleasure value associated with food is set to one. In other words, the simple detection of a relevant stimulus (ball, food, rest) is enough to associate pleasure with the stimulus. The agent will also associate pain with the stimulus whenever the minimum frontal distance has a value below the contact distance. In that case, the DV pain component is set to one. All stimuli that generate a very strong (perceptual) DV (one with the pain component equal to one) require a fast reaction through an Avoid behaviour. One of the key ideas of this architecture is that when facing an urgent situation, the perceptual level can deliver a fast decision (by choosing an adequate behaviour) without waiting for the cognitive evaluation.

Cognitive Evaluation In this evaluation, the cognitive DV (DV_c) is estimated by looking, into the memory, for similar stimuli processed in the past. If there is no need for urgent decision, the cognitive DV might override the perceptual DV.

The perceptual image of the current stimulus is used to address the cognitive memory, in the sense that similar perceptual images are more likely to be associated with similar cognitive images. This could speed up the search for similar cognitive images. A cognitive image saved on a memory frame is considered to be similar to the current cognitive image if the difference between the data for the two images (the one in memory and the current one) is within a predefined range. This process is simply called cognitive matching. In the case there is a cognitive match, the remaining frame information will be analysed. If the result corresponding to energy (or strength) is marked with success or insuccess, then the DV pleasure component associated with food (or rest) is set to one or zero, respectively. When there is no cognitive matching, the perceptual DV will remain intact.

A perceptually desirable stimulus may occult a bad experience in the past. The cognitive evaluation has to ascertain that desirability given past experiences. This allows for the agent to learn what is really desirable or not. The existence of two different kinds of food in this environment is a way to test this basic learning mechanism.

Spatial Evaluation The spatial evaluation makes use of innate meanings to produce a spatial assessment that provides the agent with a more reliable perspective about obstacles in its vicinity. The resulting spatial evaluation is a vector composed by eight boolean components indicating the presence of obstacles, in each one of the six predefined sonar-based directions, and the agent proximity to a blue or yellow object.

Memory

The agent has two different kinds of memory: a main memory, where information related to the most important stimuli processed in the past is saved, and a medium-term memory, FIFO like, where the coordinates of relevant places found when the robot is traveling around the world are stored.

Main Memory The purpose of this memory is to avoid repeating undesirable situations for the agent body state equilibrium. It is divided in two complementary memories: the cognitive and the perceptual memories. The contents of the second one is used to address the first one, as explained before.

Each cognitive memory position points to a structure called frame. Frames are introduced in cognitive memory grouped in sequences. A sequence starts when the agent chooses the Approach behaviour and it ends up when the Execute behaviour is executed. If an Execute behaviour does not follow an Approach one, the sequence is considered to be incomplete and therefore discarded from memory.

A frame contains the following information:

Cognitive data : The coordinates, in the image, of the center of mass for each recognized colour;

Energy Result : The impact of the sequence, in terms of success or insuccess, on agent energy level;

Strength Result : The impact of the sequence, in terms of success or insuccess, on agent strength level;

The perceptual memory is addressed using the information related to colours present in the stimulus. Given the number of pixels of each colour, the perceptual memory will return a list of pointers to the cognitive memory, where cognitive images, with similar amounts of each colour, are stored

Medium Term Memory This memory is used to provide information for the Plan behaviour. Each memory position indicates the coordinates (x, y, θ) of a relevant location and the type of object existing there (food or a place to rest). This information is gathered when the agent moves close to a sign that gives it pleasure, even when it does not eat or rest.

Behaviour Anticipation and Selection

As soon as the DV, the agent motivations and the spatial description are known, the process proceeds with the behaviour anticipation and the consequent behaviour selection. The agent anticipates the result of each one of the behaviours in a sequential way, by the following order: Approach, Execute, Deviate, Avoid, Play, Plan and Ramble. For each anticipation, a new DV, named DV*, is estimated. The DV* has the same structure as the DV and it represents how desirable that behaviour is according to present circumstances and previous experiences. The agent will choose the first behaviour that allows it to anticipate pleasure and no pain. If no pleasure is anticipated, then the agent will choose the Ramble behaviour. This anticipation process is performed based on a set of predefined rules, given in Figure 4.

BEHAVIOUR	**DV*[PAIN]** **⇔ Anticipates pain if...**	**DV*[PLEASURE]** **⇔ Anticipates pleasure if...**
Approach	It is close from a yellow / blue object.	It is hungry / tired and the stimulus is desirable from the food / rest point of view.
Execute	It is not close from a yellow / blue object.	It is hungry / tired and the stimulus is desirable from the food / rest point of view.
Deviate	It has detected an obstacle in front.	It has detected an obstacle in the 1° or 2° quadrants.
Avoid	It has detected an obstacle backwards and in the 3° and 4° quadrants.	It has detected an obstacle in front and in the 1° and 2° quadrants.
Play	It has detected an obstacle in front.	It wishes to play and the stimulus is desirable from the ball point of view.
Plan	It is not hungry / tired.	It is hungry / tired and there's a place with food / rest saved in memory.
Ramble	-	-

Figure 4: Behaviours *vs.* pain and pleasure evaluation.

Results

For the sake of paper length, only a pair of snapshots taken during a run is included in this paper. Each figure includes the vision image acquired from the CCD camera, and the current state information, which includes the body state values, the robot posture, the quantity of pixels for each colour and the minimum frontal distance, the DV vector, the motivation and the selected behaviour.

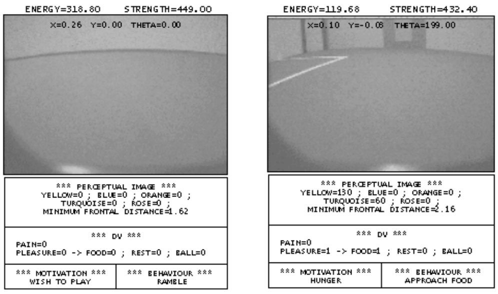

Figure 5: On the left, a snapshot at the beginning of the run, and on the right, seeing and approaching a food sign.

At the beginning of the run, the robot is positioned on a random location and its body state values are equal to the HV (equilibrium) ones (320 for energy and 450 for strength). Its motivation is "wish to play", but as it does not see the ball, it starts rambling. When the agent gets very close to an obstacle, it avoids it, and also avoids the places where it collides with something. As soon as the agent gets a glimpse of the ball (and it wishes to play), it is going to play. If the energy component of agent body state falls to a value below the instability threshold, the motivation changes to hunger and the agent will approach the first food sign it detects (a sign with yellow colour).

When the agent is close enough to the food sign, it starts eating. The way the eat (and also the rest) behaviour has been implemented was by making the robot rotate around himself during a certain period of time. While the agent is eating, the energy value increases till it reaches the HV value. After that, the agent stores the location coordinates in the medium term memory and selects a new behaviour. In what concerns rotten food, the energy value does not change, so the agent avoids that place after trying to eat for 10 seconds. If the agent is hungry and it can not see any food sign but it has, in its main memory, information about a location where it has seen a food sign, the agent elaborates a simple plan (based on odometric data) to return there. The same happens when the agent is tired.

All the way to the location target, the agent will avoid any obstacles that came along. Moreover, if by any chance the agent sees another food sign, different from the planned one, it will interrupt the plan and approach this new sign. If any of these signs turn to be rotten food and the agent have already experienced that kind of food, the agent will recall this bad experience and avoid the rotten food sign. shows the situation where the robot has already the food sign in its line of view and adopts the Approach behaviour. Furthermore, notice that the robot is seeing the ball but it ignores it, because the "hunger" motivation is stronger that the "wish to play" one.

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 is needed. Recent publications of psychology (Goleman 1996) and neuroscience research results (Damásio 1994; LeDoux 1996) suggest a relationship between emotion and rational behaviour.

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 work of Velásquez (Velásquez 1998a; 1998b), who developed a pet-robot based on Damásio's ideas, and the work of Breazeal (Breazeal 1999), who presented Kismet, a socially situated robot, based on a control architecture integrating synthetic emotions.

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 1998), 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; Aubé 1998). Some authors are now trying to formalize the notion of emotion using different frameworks, namely, the category theory (Arzi-Gonczarowski 2000), and decision theory (Gmytrasiewicz & Lisetti 2000).

Conclusions and Future Work

In what performance concerns, the results have shown that in most of the experiments, the agent was able to survive for a long period of time, by fulfilling its needs of food and rest. The survival rate depends heavily on the environment dimension and the number of signs available. The utilization of the medium term memory allows the agent to formulate plans for handling difficult situations. On the other hand, the cognitive evaluation, based on the memory of cognitive images, enables the agent to make adequate decisions when it is potentially facing a bad (previously seen) experience.

The proposed architecture allowed the implementation of an autonomous robotic agent, (i) where the goal definition results from the agent 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.

In what concerns the application of this architecture in social domains, one on-going experiment is the introduction of another robot in the environment playing the role of predator. Besides fulfilling its needs, the agent must pay attention to the predator and run away from it. To implement this, two different approaches have been tested: i) first, by considering that the agent does not know what a predator is; so the agent has to learn how to deal with it based on the consequences of being in contact with the predator, and ii) second, by assuming a new relevant colour (the black colour as it is the colour of the robots) with negative desirability; hence, the perceptual and cognitive levels will act by taking into account the urgency of the situation (estimated based on the distance to the predator).

Another work in progress is the application of this architecture in domains where social relationship is a key aspect (e.g., a market world, where there is a set of agents acting as consumers and/or vendors and a set of goods available for selling or trading). This work has two main goals, on the one hand, to study which kind of behaviours emerge from this architecture and the type and extension of innate knowledge needed in such environment. The experiments performed by Damásio, with normal people and patients with prefrontal lobes lesions, using a simple deck game (Damásio 1994), suggests that the somatic marking mechanism, implemented in this architecture, could be sufficient to handle simple trade decisions, particularly when the agent survival is involved. However, as an environment of this kind could raise much more complex social relations, which bring the need for rational decisions, for instance based on cost/benefit evaluations, the second goal is to study how a rational decision maker could be integrated within the proposed architecture. Some interesting social relationships to be studied are cooperation, competition, negotiation and, of course, communication.

References

Arzi-Gonczarowski, Z. 2000. A categorization of autonomous action tendencies: The mathematics of emotions. In *Proceedings of the EMCSR2000, Symposium 'Autonomy Control: Lessons from the Emotional'*, volume 2, 683–688.

Aubé, M. 1998. A commitment theory of emotions. In Cañamero, D., ed., *Emotional and Intelligent: the Tangled Knot of Cognition. 1998 AAAI Fall Symposium.* AAAI. 13–18.

Breazeal, C. 1999. Robot in society: Friend or appliance? In *Workshop on Emotion-based Agent Architectures — Agents99*. 18–26.

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., *Proceedings of the first International Conference on Autonomous Agents*, 148–155. ACM Press.

Custódio, L.; Ventura, R.; and Pinto-Ferreira, C. 1999. Artificial emotions and emotion-based control systems. In *Proceedings of 1999 7th IEEE International Conference Emerging Technologies and Factory Automation (ETFA'99)*, volume 2, 1415–1420. ACM Press.

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

Damásio, A. 1999. *The Feeling of What Happens:body and emotion in the making of consciousness*. Harcourt Brace.

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

Gmytrasiewicz, P. J., and Lisetti, C. L. 2000. Using decision theory to formalize emotions for multi-agent systems. In *Second ICMAS-2000 Workshop on Game Theoretic and Decision Theoretic Agents*.

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

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

Maçãs, M.; Ventura, R.; Custódio, L.; and Pinto-Ferreira, C. 2001a. DARE: An emotion-based agent architecture. In *Proceedings of the FLAIRS-2001 Conference*.

Maçãs, M.; Ventura, R.; Custódio, L.; and Pinto-Ferreira, C. 2001b. Experiments with an emotion-based agent using the DARE architecture. In *Proceedings of the Symposium on Emotion, Cognition, and Affective Computing, AISB'01 Convention*.

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

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

Sloman, A. 1998. Damasio, descartes, alarms and metamanagement. In *Proc. on Int. Conference on systems, Man, and Cybernetics (SMC98)*, 2652–7. IEEE.

Staller, A., and Petta, P. 1998. Towards a tractable appraisal-based architecture. In Caamero, 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.

Vale, P., and Custódio, L. 2001. Learning individual basic skills using an emotion-based architecture. In *Proceedings of the Symposium on Emotion, Cognition, and Affective Computing, AISB'01 Convention*.

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

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

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

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