When Robots Weep: Emotional Memories and Decision-Making

Juan D. Velásquez

MIT Artificial Intelligence Laboratory
545 Technology Square, NE43-935
Cambridge, Massachusetts 02139
jvelas@ai.mit.edu

Abstract

We describe an agent architecture that integrates emotions, drives, and behaviors, and that focuses on modeling some of the aspects of emotions as fundamental components within the process of decision-making. We show how the mechanisms of primary emotions can be used as building blocks for the acquisition of emotional memories that serve as biasing mechanisms during the process of making decisions and selecting actions. The architecture has been implemented into an object-oriented framework that has been successfully used to develop and control several synthetic agents and which is currently being used as the control system for an emotional pet robot.

Introduction

The traditional view on the nature of rationality has proposed that emotions and reason do not mix at all. For an agent to act rationally, it should not allow emotions to intrude in its reasoning processes. Research in Neuroscience, however, has provided evidence indicating quite the contrary, showing that emotions play a fundamental role in perception, learning, attention, memory, and other abilities and mechanisms we tend to associate with basic rational and intelligent behavior [Damasio 1994; LeDoux 1966; Adolphs 1966]. In particular, recent studies of patients with lesions of the prefrontal cortex suggest a critical role for emotions in decision-making [Bechara et al. 1997; Churchland 1996; Damasio 1994]. Although the studied patients can perform well on a variety of intelligence and memory tests, when faced with real-life situations they seem to be unable to make "good" decisions. Apparently, these patients lack intuition abilities, which, as many researchers think, may be based on memories of past emotions. These findings indicate that contrary to popular belief, intuition and emotions play significant roles in our abilities to make smart, rational decisions.

To this date, the field of Artificial Intelligence has largely ignored the use of emotions and intuition to guide reasoning and decision making. Several models of emotions have been proposed, but most of the work in this area has focused on specific aspects, such as recognizing emotions [Picard 1997], synthesizing emotions as the primary means to create believable synthetic agents [Bates 1994; Blumberg 1994; Elliot 1992; Maes 1995; Reilly 1996;], or synthesizing emotions and some of their influences in behavior and learning [Frijda 1996; Kitano 1995; Pfeifer 1988; Velásquez 1997].

The work described in this paper derives from, and extends previous research on computational models of emotions [Velásquez 1997]. Our main contribution is to show how drives, emotions, and behaviors can be integrated into a robust agent architecture, that uses some of the mechanisms of emotions to acquire memories from past emotional experiences that serve as biasing mechanisms while making decisions during the action-selection process.

Emotions as Biasing Mechanisms

The studies on patients with lesions in the prefrontal cortex mentioned above, motivated Damasio and colleagues to suggest that human reasoning and decision-making involves several mechanisms at different levels, extending from those that perform basic body regulation, to those that deal with more cognitive control of complex strategies. An interesting and novel component of this view is that reasoning depends also on emotions and the feelings accompanying them, which involve images that relate to the state of the body [Damasio 1994].

According to Damasio, part of this process includes the use of a covert, nonconscious biasing mechanism that directs us towards the "right" decision. This biasing step is known as the *somatic marker hypothesis*. The main idea behind this hypothesis is that decisions that are made in circumstances similar to previous experience, and whose outcome could be potentially harmful, or potentially advantageous, induce a *somatic* response used to *mark* future outcomes that are important to us, and to signal their danger or advantage. Thus, when a negative somatic marker is linked to a particular future outcome it serves as an alarm signal that tell us to avoid that particular course of action. If instead, a positive somatic marker is linked, it becomes an incentive to make that particular choice.

These ideas inspired the model described below.

The Computational Model

This section describes *Cathexis*, a computational model of emotions and action-selection inspired by work in different fields, including Neuropsychology, Artificial Intelligence, and Ethology. In particular, it has a strong influence from neuropsychological theories about the functional organization of the prefrontal lobes and their interaction with other neural systems involved in mediating emotions and sensorimotor responses that guide decision-making [Damasio 1994; Adolphs et al. 1996; Churchland 1996; Altman 1996]. Figure 1 provides a high level view of the model's architecture.

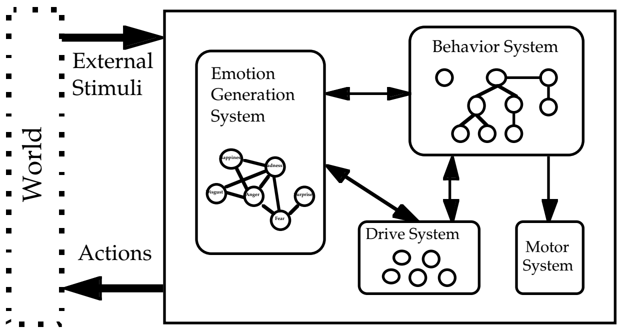

Figure 1 The Model's Architecture.

The Drive System

The Drive System consists of a set of motivational systems, or drives, representing urges that impel the agent into action. For instance, a *Hunger* drive will aid in controlling behaviors that directly affect the level of food intake by the agent.

Each Drive includes a set of Releasers which filter sensory data and identify special conditions which will either increase or decrease the value of the drive they belong to. These releasers represent control systems that maintain a controlled variable within a certain range. Drive control systems measure this variable through some of the agent's sensors and compare it to a desired value or set point. If its value does not match the set point, an error signal is produced. This error signal is fed to the appropriate drive, in which it can be combined with error signals from other relevant control systems. For instance, a *TemperatureRegulation* drive would combine the signals from two different control systems: one controlling peripheral temperature and one controlling brain temperature.

The value of each drive is determined by a linear combination of its control systems as described in Equation (1):

$$D_{it} = \sum_{k} (C_{ki} \cdot W_{ki}) \qquad (1)$$

Where D_{it} is the value of Drive i at time t; C_{ki} is the value of Control System k, and W_{ki} is the weight associated to control system k, where k ranges over the releasers for Drive i.

The Emotion Generation System

Following a neuropsychological perspective, the Emotion Generation System bears resemblance to some of the aspects in which the interactions between neural systems involving the amygdala, the hippocampus, and the prefrontal cortices have been considered to mediate emotions, such as assigning an emotional valence to different stimuli, activation of emotional behaviors, and emotional learning [Damasio 1994; LeDoux 1993; Panksepp 1995].

This system consists of a distributed network of self-interested emotional systems representing different families of related affective states, such as Fright, Fear, Terror, and Panic. Each member of an emotion family shares certain mechanisms and characteristics, including similarities in antecedent events, expression, likely behavioral response, and physiological patterns. These characteristics differ between emotion families, distinguishing one from another.

Drawing upon ideas from different theorists [Ekman 1992; Izard 1991; Johnson-Laird and Oatley 1992], we have identified and created explicit models for six different emotion families: *Anger*, *Fear*, *Distress/Sadness*, *Enjoyment/Happiness*, *Disgust*, and *Surprise*. The selection of this core set of emotion types is not arbitrary, but rather it is based upon evidence suggesting their universality, including distinctive universal facial expressions, as well as eight other properties [Ekman 1992].

Emotional Systems have a set of Releasers that constantly check for the appropriate conditions that would elicit the emotion they belong to. Influenced by Izard's multi-system for emotion activation [Izard 1993], we consider both cognitive and noncognitive releasers and divide them into four different groups:

- *Neural*: Includes the effects of neurotransmitters, brain temperature, and other neuroactive agents that can lead to emotion and which can be mediated by hormones, sleep, diet, and environmental conditions. For instance, there is a great deal of evidence that shows that decreased levels of norepinephrine and serotonin are associated with depression [Meltzer et al. 1981]. Similarly, it is clear that several chemical agents, such as carbon dioxide, yohimbine, and amphetamines produce anxiety in humans by activating the noradrenergic system [Charney and Redmon 1983].
- *Sensorimotor*: This system covers sensorimotor processes, such as facial expressions, body posture, muscle action potentials, and central efferent activity, that not only regulate ongoing emotion experiences but can also elicit emotion. Some evidence supporting this type of elicitors comes from neuropsychological studies in which experimenter-directed manipulation of facial muscles, composing a specific emotional expression, produces the subjective feeling corresponding to that emotion, as well as emotion-specific patterns of autonomic nervous system (ANS) activity [Ekman, Levenson, and Friesen 1993].
- *Motivational*: This system includes all motivations that lead to emotion. In this model, motivations include drives (e.g. *Thirst* and *Hunger*), emotions (e.g. *Anger*, and *Happiness*), and pain regulation. Some examples of elicitors in this system include the innate response to foul odors or tastes producing disgust, as measured in neuropsycho-

logical studies by [Fox and Davidson 1986], pain or aversive stimulation causing anger, and emotions like sadness eliciting others such as anger.

- *Cognitive*: This system includes all type of cognitions that activate emotion, such as appraisal of events, comparisons, attributions, beliefs and desires, memory, and so on. In previous work, these elicitors were based on a cognitive appraisal theory (See [Velásquez 1997] for details). In an effort to design a more plausible model, we have revised the Emotion Generation System so that it does not include any pre-wired cognitive elicitors, but rather allows for them to be learned through emotional experiences, as the agent interacts with its environment.

Besides its releasers, each Emotional System includes two different thresholds. The first one, α, is used to determine when an emotion episode occurs. That is, once its intensity goes above this threshold, the Emotional System releases its output signal to other Emotional Systems and to the Behavior System which in turn selects and controls an appropriate behavior according to the agent's motivational state. The second threshold, ω, specifies the level of saturation for that emotion. This is consistent with real life emotional systems in which levels of arousal will not exceed certain limits. In addition to these parameters, each Emotional System has a function, $\Psi()$, which controls the temporal decay of its intensity. These kinds of mechanisms contribute to the nonlinear behavior exhibited by the model. Figure 2 illustrates these ideas.

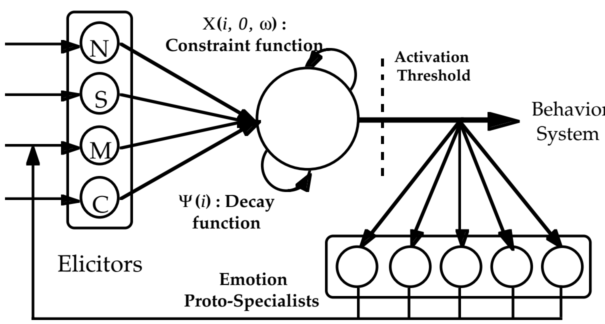

Figure 2 Emotional Systems.

The intensity of an Emotional System depends on all the factors that contribute to it, including its previous level of arousal, the contributions of each of its elicitors, and the interaction (inhibitory and excitatory) with other Emotional Systems. This is summarized in Equation (2):

$$I_{et} = \chi\left(\Psi(I_{et-1}) + \sum_{k} L_{ke} + \sum_{l} (G_{le} \cdot I_{lt}) - \sum_{m} (H_{me} \cdot I_{mt})\right) \quad (2)$$

Where I_{et} is the value of the intensity for Emotion e at time t; I_{et-1} is its value at the previous time step; $\Psi()$ is the function that represents how Emotion e decays; L_{ke} is the value of Emotion Elicitor k, where k ranges over the Emotion Elicitors for Emotion e; G_{le} is the Excitatory Gain that Emotion l applies to Emotion e; I_{lt} is the intensity of Emotion l at time t; H_{me} is the Inhibitory Gain that Emotion m applies to Emotion e; I_{mt} is the intensity of Emotion m at time t; and $X()$ is the function that constrains the intensity of Emotion e between 0 and its saturation value.

This model of Emotional Systems allows for the distinction between different affective phenomena. For instance, primary emotions are modeled as the activation of one particular Emotional System such as *Sadness* or *Disgust*. Emotion blends, such as *Jealousy*, emerge as the co-activation of two or more of these Emotional Systems. Similarly, and following a psychobiological perspective [Panksepp 1995], moods are differentiated from emotions in terms of levels of arousal. While emotions consist of high arousal of specific Emotional Systems, moods may be explained as low tonic levels of arousal within the same systems (i.e. levels below the α threshold). This representation is consistent with the enormous subtleties of human moods and feelings, as well as with the common observation that moods seem to lower the threshold for arousing certain emotions. This occurs because emotional systems that are aroused, as it happens in the representation of moods, are already providing some potential for the activation of an emotion. Finally, it is consistent with the observation that the duration of moods appears to be longer than that of the emotions, since at low levels of arousal, the intensity of the Emotional Systems will decay more slowly.

Finally, temperaments are modeled through the different values that parameters (e.g., thresholds, gains, and decay rates) within each Emotional System can have. Thus, for instance, if we want to model a depressed agent, we might lower the activation threshold and decay rate for the *Sadness* emotion as well as lowering the inhibitory gain between *Happiness* and *Sadness*. The result is a flexible, distributed model that can synthesize a variety of affective phenomena simultaneously.

The Behavior System

Following Damasio's view, reasoning and decision-making define a domain of cognition in which an agent must choose how to respond to a situation. Concepts like *decision*, *reason*, *action-selection*, and *rationality* are fundamentally linked to behavior. When an agent faces a situation, a decision of what to do next must be made. This choice is responsibility of the Behavior System. The action-selection process is mediated by the reasoning engaged. If the outcome of the reasoning and the selected behavior are adaptive (oversimplified in this model as selecting a behavior that avoids negative outcomes), the choice is considered rational.

The Behavior System is a distributed network of self-interested behaviors, such as "*approach human*", "*play*", "*request attention*", and "*avoid obstacle*".

Like Drives and Emotional Systems, Behaviors have Releasers that obtain and filter sensory data in order to identify special conditions which will either increase or decrease the value of the Behavior. Releasers might represent objects and conditions such as “battery recharger is present”, and motivational states such as “battery level is low” and “distress is high”, which would most likely increase the value of a “*recharge batteries*” behavior.

Behaviors may mutually inhibit or excite each other. For instance, “*wag the tail*” might inhibit “*running*” and vice-versa. Whereas behaviors such as “*play with human*” might excite lower-level ones like “*find human*”.

In earlier work, the Behavior System followed a winner-take-all strategy in which only one behavior could be active at a time. This made it impossible for non-conflicting Behaviors, such as “*walk*” and “*cry*” to execute at the same time. Given the parallelism of the model, we revised the Behavior System so that active, non-conflicting Behaviors can issue motor commands simultaneously. The value for each Behavior is computed as described in Equation (3):

$$B_{jt} = \sum_{n}(R_{nj} \cdot W_{nj}) + \sum_{l}(G_{lj} \cdot B_{lt}) - \sum_{m}(H_{mj} \cdot B_{mt}) \qquad (3)$$

Where B_{jt} is the value of Behavior j at time t; R_{nj} is the value of releaser n and W_{nj} is the weight for releaser n, where n ranges over the releasers for Behavior j; G_{lj} is the Excitatory Gain that Behavior l applies to Behavior j, and B_{lt} is the intensity of Behavior l at time t, where l ranges over the set of behaviors that excite Behavior j; H_{mj} is the Inhibitory Gain that Behavior m applies to Behavior j, and B_{mt} is the intensity of Behavior m at time t, where m ranges over the set of behaviors that inhibit Behavior j.

Integrating Emotional Memories

The mechanisms described above for primary emotions do not describe the whole range of emotions we experience. In fact, most of our emotional experiences can be considered secondary which occur after we begin experiencing feelings and start making orderly associations between objects and situations, and primary emotions. Thus, for instance, whereas a loud noise might activate an innate fear response (primary emotion), thinking about not making a paper deadline might activate a learned one (secondary emotion). This latter kind of emotion requires more complex processing, including in most cases, the retrieval of emotional memories of similar previous experiences. Following the ideas behind both Damasio’s somatic-marker hypothesis [Damasio 1994], and LeDoux’s work on fast (low-road) and slow (high-road) pathways for emotion activation [LeDoux 1993], we have extended our model to consider not only pre-wired, stimuli-driven emotions, but also, more cognitive, memory-based, learned ones.

These secondary emotions have been modeled with an associative network comparable to Minsky’s K-lines [Minsky 1986], in which primary emotions are connected to the specific stimuli (e.g., executed behaviors, objects, or agents) that have elicited them during the agent’s interaction with the world. The connections between primary emotions and different stimuli specify the amount (averaged throughout different occurrences) of emotional energy (intensity of the Emotional System) applied to each stimulus when encountered.

For instance, supposing that an agent is about to engage in *Feeding* and the only available food is a bad-tasting soup, it is likely that once the agent eats, the *Disgust* Emotional System will become active at some particular intensity because it has a pre-wired elicitor for foul odors and tastes. Once this happens, an association is made between the primary emotion (*Disgust*) and the stimulus that provoked it (soup), and an emotional memory is created.

The emotional memory by itself is not very useful. Instead, the model described above has been extended so that it uses this information as part of the behavior selection process. Thus, the next time an agent encounters a *marked* stimulus, such as the soup in our example, the memory represented in the associative network will be relived, reproducing the emotional state previously experienced, and influencing the selection of actions to follow.

Thus, although the agent had no pre-wired aversion for soup, its previous negative experience has created a learned one for it. Furthermore, even if the agent is very hungry, it is likely that the *Feeding* behavior will not become active if soup is the only food present. Hence, the purpose of emotional memories is twofold. First, they allow for the learning of secondary emotions as generalizations of primary ones. And second, they serve as markers or biasing mechanisms that influence what decisions are made and how the agent behaves.

Implementation and Results

The model described in the previous sections has been implemented in its totality as part of an object-oriented framework for building autonomous agents. We have used this framework to develop and control various synthetic agents, including *Simón the Toddler* (See [Velásquez 1997] for a description), and *Virtual Yuppy*, a simulated emotional pet robot, shown in Figure 3.

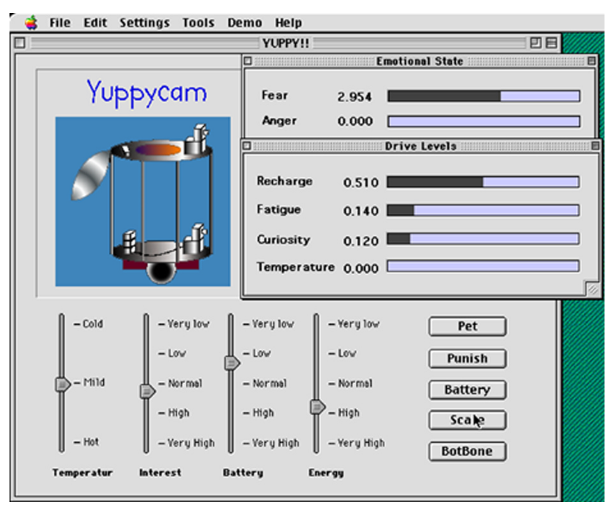

Figure 3 Virtual Yuppy

As part of our ongoing research, the same framework has been used to control Yuppy, the actual physical robot shown in Figure 4.

The implementation details and results described in this section correspond in most part to the simulated robot, but are also applicable to current observations with the physical robot.

Figure 4 Yuppy, an Emotional Pet Robot

Virtual Yuppy has different sensors, including one for very simple synthetic vision, as well as simulated tactile sensors to model painful and pleasurable stimuli.

Its Drive System is composed of four different drives: *RechargingRegulation, TemperatureRegulation, Fatigue,* and *Curiosity*, each of which controls internal variables representing the agent's *battery, temperature, energy,* and *interest* levels, respectively.

Its Emotion Generation System includes emotional systems with innate releasers for the set of basic emotion families described before. These releasers have been grouped under the Neural, Sensorimotor, and Motivational elicitor categories, modeling the primary emotions for the agent. For instance, via motivational elicitors, unsatisfied drives produce *Distress* and *Anger*, whereas satiation generates *Happiness*. Similarly, "synthetic bones" and "petting" elicit *Happiness, Fear* includes an innate releaser for darkness, and pain produces *Distress* and *Anger.*

The robot's Behavior System is composed of a distributed network of approximately nineteen different self-interested behaviors, directed in most part towards satisfying its needs and interacting with humans. Examples of such behaviors include *"search for bone"*, *"approach bone"*, *"recharge battery"*, *"wander"*, *"startle"*, *"avoid obstacle"*, *"approach human"*, *and "express emotion"*.

The user interacts with the robot in two different ways: First, by controlling its affective style, which is done by tweaking the different parameters described above (e.g., thresholds, gains, or inhibitory and excitatory connections) for each Emotional System. And second, by providing stimuli for the agents, whether in the form of internal stimuli, such as modifying the level of synthetic neurotransmitters and internal variables, or external stimuli, such as hitting the robot, showing the bone, and so on.

Using the model described before, Virtual Yuppy produces emotional behaviors under different circumstances. For instance, when its *Curiosity* drive is high, Virtual Yuppy wanders around, looking for the synthetic bone which some humans carry. When it encounters one, its level of *Happiness* increases and specific behaviors, such as *"wag the tail"* and *"approach the bone"* become active. On the other hand, as time passes by without finding any bone, its Distress level rises and sad behaviors, such as "droop the tail", get executed. Similarly, while wandering around, it may encounter dark places which will elicit fearful responses in which it backs up and changes direction.

Besides regulating action-selection and generating emotional behaviors through primary emotions, Virtual Yuppy learns secondary emotions which are stored as new or modified cognitive elicitors based on the associative network model described before. For instance, after locating a bone, the robot may approach it, thus approaching the human who is carrying it. Depending on these interactions (e.g., humans pet or hit the robot), Virtual Yuppy will create positive or negative emotional memories with respect to humans, and future selection of behaviors such as approaching or avoiding them will be influenced.

Related Work

Given the space limitations, a comprehensive review of related work is not possible, and only the most relevant work is discussed in this section. For an overview of various models the reader is referred to [Picard 1997; Pfeifer 1988; Hudlicka and Fellows 1996].

Most of the recent work on this area has focused on modeling emotions for entertainment purposes. Some excellent work in the area of believable agents includes Reilly's *Em* architecture [Reilly 1996] and Elliot's *Affective Reasoner* [Elliot 1992]. Both differ from our work in several important ways. Their approach is mostly concerned with contribution of appraisal to emotion, hence they emphasize on cognitively generated emotions. In contrast, fast, primary emotions, emergent emotions, and other affective phenomena are not explicitly modeled or are otherwise oversimplified. Also, they do not model interactions with other processes, including regulatory mechanisms (i.e. drives), decision-making, and emotional learning.

It should be noted, however, that some of these differences may be due, in part, to the specific purpose for which each model was designed. Their work is mostly aimed at designing tools and models to create believable agents. Therefore, rule-based approaches that emphasize more on cognitive generation of emotion and less on designing a plausible model, may be the most appropriate ones.

Our work also relates to models of action-selection. A number of researchers have proposed successful models of action-selection for agents with multiple goals that operate in unpredictable environments [Brooks 1986; Maes 1990, Blumberg 1994; Tyrell 1993]. However, and in contrast to the work presented here, most of these models do not consider emotions as an integral part of the action-selection process, and when they do, they do not include explicit emotion models, emotional states, or moods, but rather simplified internal variables that represent emotions as well as other motivations such as hunger or thirst.

Conclusions

We have presented a flexible agent architecture that integrates drives, emotions, and behaviors and that focuses on emotions as the main motivational system that influences how behaviors are selected and controlled. We have showed how the mechanisms of primary emotions included in the proposed model, and which have been inspired by work in Neuropsychology and other fields, can be used to create emotional memories, or secondary emotions that act as biasing mechanisms during the process of making decisions and selecting actions.

Acknowledgments

The author would like to thank Professor Rod Brooks for his support of this work. Also, thanks to the entire Yuppy team for their help. In particular, thanks to Charles Kemp for his suggestion on simultaneous activation of non-conflicting behaviors. This research is part of the Yuppy project which is funded by Yamaha.

References

Adolphs, R et al. 1996. Neuropsychological Approaches to Reasoning and Decision-Making. In: Damasio, A., et al. Eds. *Neurobiology of Decision-Making*. Berlin: Springer-Verlag.

Altman, J. 1996. Epilogue: Models of Decision-Making. In: Damasio, A., et al. Eds. *Neurobiology of Decision-Making*. Berlin: Springer-Verlag.

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

Bechara, A., et al. 1997. Deciding Advantageously Before Knowing the Advantageous Strategy. In: *Science*, 275, 1293-1295.

Blumberg, B. 1994. Action-Selection in Hamsterdam: Lessons from Ethology. In *Proceedings of SAB94*, 108-117. Brighton, England: MIT Press.

Brooks, R. 1986. A Robust Layered Control System For A Mobile Robot. In: *Robotics and Automation*, RA-2(1).

Charney, D.S., and Redmond, DE. 1983. Neurobiological Mechanisms in Human Anxiety. In: *neuropharmacology*, 22, 1531-1536.

Churchland, P.S. 1996. Feeling Reasons. In: Damasio, A., et al. Eds. *Neurobiology of Decision-Making*. Berlin: Springer-Verlag.

Damasio, A. 1994. *Descartes' Error: Emotion, Reason, and the Human Brain*. New York: Gosset/Putnam.

Ekman, P. 1992. An Argument for Basic Emotions. In: Stein, N. L., and Oatley, K. eds. *Basic Emotions*, 169-200. Hove, UK: Lawrence Erlbaum.

Ekman, P., Levenson, R., and Friesen, W. 1983. Autonomic Nervous System Activity Distinguishes Among Emotions. In: *Science*, 221, 1208-1210.

Elliot, C.D. 1992. The Affective Reasoner: A Process Model of Emotions in a Multi-Agent System. Ph.D. Thesis, Institute for the Learning Sciences, Northwestern University.

Fox, N., and Davidson, R. 1986. Taste-Elicited Changes in facial Signs of Emotion and the Asymmetry of Brain Electrical Activity in Human Newborns. In: *Neuropsychologia*, 24, 417-422.

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

Hudlicka, E., and Fellows, J. 1996. Review of Computational Models of Emotions. Tech. Report 9612, Psychometrix.

Izard, C. 1991. *The Psychology of Emotions*. New York: Plenum Press.

Izard, C. E. 1993. Four Systems for Emotion Activation: Cognitive and Noncognitive Processes. *Psychological Review* 100(1):68-90.

Johnson-Laird, P. and Oatley, K. 1992. Basic Emotions, Rationality, and Folk Theory. In: Stein, N. L., and Oatley, K. eds. *Basic Emotions*, 169-200. Lawrence Erlbaum.

Kitano, H. 1995. A Model for Hormonal Modulation of Learning. In: *Proceedings of IJCAI-95*. Montreal.

LeDoux, J. 1993. Emotional Memory Systems in the brain. *Behavioral and Brain Research*, 58.

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

Maes P. 1995. Artificial Life meets Entertainment: Lifelike Autonomous Agents. *Communications of the ACM. Special Issue on Novel Applications of AI*.

Maes P. 1990. Situated Agents Can Have Goals. *Robotics and Autonomous Systems*, 6(1&2).

Meltzer, H.Y. et al. 1981. Serotonin Uptake in Blood Platelets of Psychiatric Patients. In: *Archives of General Psychiatry*, 38:1322-1326.

Minsky. M. 1986. *The Society of Mind*. New York: Simon & Schuster.

Panksepp, J. 1995. The Emotional Brain and Biological Psychiatry. *Advances in Biological Psychiatry*, 1, 263-286.

Pfeifer, R. 1988. Artificial Intelligence Models of Emotion. In: Hamilton, V., Bower, G. H., and Frijda, N. eds. *Cognitive Perspectives on Emotion and Motivation*, 287-320. Netherlands: Kluwer.

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

Reilly S. 1996. Believable Social and Emotional Agents. Technical Report, CMU-CS-96-138, School of Computer Science, Carnegie Mellon University.

Tyrell, T. 1993. The Use of Hierarchies for Action Selection. In: *Proceedings of SAB92*.

Velásquez, J. 1997. Modeling Emotions and Other Motivations in Synthetic Agents. In: *Proceedings of AAAI-97*.