Modeling Emotion-Based Decision-Making

Juan D. Velásquez

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

Abstract

This paper presents a computational approach to Emotion-Based Decision-Making that models important aspects of emotional processing and integrates these with other models of perception, motivation, behavior, and motor control. A particular emphasis is placed on using some of the mechanisms of emotions as building blocks for the acquisition of emotional memories that serve as biasing signals during the process of making decisions and selecting actions. We have successfully followed this approach to develop and control several different autonomous agents, including both synthetic agents and physical robots.

Introduction

Most theories of human reasoning and decision-making fall between two different positions. The first one argues that we make decisions in a way similar to that of solving problems in formal logic. According to this view, when faced with a problem, we form a list of all different options and their possible outcomes, and then we use logic in its best sense to perform a cost/benefit analysis that will provide us with the best possible choice. The second view considers reasoning and decision-making to be associative. That is, when confronted with a situation that requires a decision, we compare it to similar situations that have been encountered in the past, and tend to act accordingly. For a review of some of these theories see [Evans et al. 1993].

Motivated by the findings of several studies on patients with lesions in the prefrontal cortex, Damasio and colleagues have proposed a recent framework for the understanding of human reasoning and decision-making [Damasio 1994; Adolphs et al. 1996; Bechara et al. 1997]. An interesting and novel component of this view is that reasoning depends also on emotional processing and the resulting feelings, which involve images that relate to the state of the body. In contrast to traditional cognitive approaches, this proposal emphasizes that reasoning is not disembodied, but instead uses biological information to bias and steer the decision-making process toward outcomes that are advantageous, based on past experiences with similar situations [Adolphs et al. 1996].

The hypothesis behind this biasing mechanism is known as the *somatic marker hypothesis*. This hypothesis states 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 [Damasio 1994].

These findings indicate that contrary to popular belief, intuition and emotions play crucial roles in our abilities to make smart, rational decisions.

To this date, however, the field of Artificial Intelligence has generally ignored the use of emotions and intuition to guide reasoning and decision-making. Although several interesting models of emotions have been proposed, most of them are oriented towards using emotions for entertainment purposes [Bates 1994; Blumberg 1994; Elliot 1992; Maes 1995; Reilly 1996] or are focused on specific aspects, such as modeling very limited psychological problems [Abelson 1963; Colby 1975], recognizing emotions [Picard 1997], modeling physiological and hormonal influences of emotion [Cañamero 1997; Kitano 1995], or modeling influences in goals and learning [Blumberg, Todd, and Maes 1996; Frijda 1986; Pfeifer 1988; Wright 1996].

Due to space limitations, a comprehensive review of related work is not possible, and only the work most relevant to the present discussion has been mentioned. For an overview of various models, including some that were not listed here, the reader is referred to [Picard 1997; Pfeifer 1988; Hudlicka and Fellows 1996].

This paper presents an approach to the study of emotions and decision-making that has been inspired by recent findings in neuropsychology and that relies on the use of computational frameworks for what we call *Emotion-Based Control*, control of autonomous agents that relies on, and arises from, emotional processing.

A Model for Emotion-Based Decision-Making

The architecture, an extension of the Cathexis model described in previous work [Velásquez 1996; Velásquez 1997] is depicted in Figure 1. Relevant systems are represented as modules. The perceptual systems obtain information from the world and provide behavior and emotional systems with stimuli features and objects. These systems also receive error signals from drive systems. The emotional systems assess the emotional significance of stimuli and bias behavioral responses and future perception accordingly. Relevant behavior systems generate and execute appropriate motor actions.

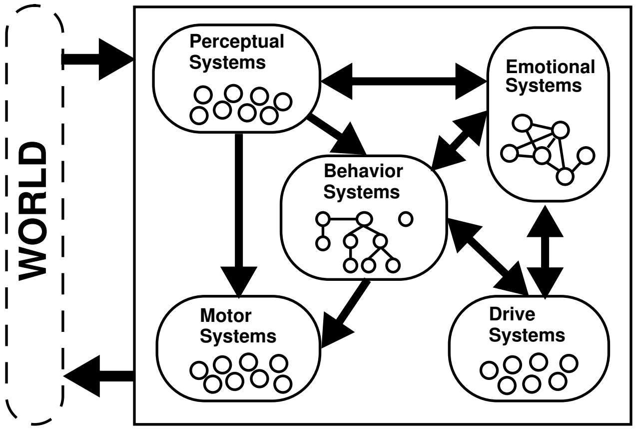

Figure 1 Schematic of the proposed framework for Emotion-Based Control.

The modular representations adopted in the figure exist solely for illustration purposes. In practice, none of these systems exist as single modules that are connected with each other through one single interface, but rather they consist of distributed networks of independent processing units which connect directly to other relevant units.

Computational Units

The basic elements of the framework correspond to nonlinear computational units that are used to represent all relevant systems (i.e., sensory systems, emotional systems, behaviors, drive systems, and motor systems).

Each of these systems can be thought of as consisting of a set of inputs, an appraisal mechanism and a set of outputs.

An important component of the appraisal mechanism consists of what we have called *releasers*, which correspond to computational units that filter sensory data and identify special conditions which will provide excitatory (positive) or inhibitory (negative) input to the system they are associated with.

Releasers can be innate and hard-wired (*natural releasers),* or they can be learned (*learned releasers*). The latter kind represent stimuli that tend to be associated with and predictive of the occurrence of natural releasers.

All releasers posses short-term memory and thus can habituate to stimuli using a model of rate-sensitive habituation similar to that proposed by [Staddon 1993].

The general form of the response or activation of each basic system is a nonlinear function of its input. Which in turn is proportional to the output of its releasers and the strengths of their connections, or weights, as described in Equation (1):

$$A_i = f\left(\sum_k (R_{ki} \cdot W_{ki})\right) \tag{1}$$

Where A_i is the activation of system i; R_{ki} is the value of releaser k, and W_{ki} is its associated weight, where k ranges over the set of releasers for system i, and f is a limiting function such as the standard ramp and logistic functions.

Drive Systems

Drives are motivational systems 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. It should be noted, however, that drive systems are differentiated from emotional systems. In fact, within the proposed framework, emotional systems constitute the main motivations for the agent and even drive systems exploit their power to bias specific behaviors. In other words, it is both the error signal produced by a *Hunger* drive, and the *Distress* caused by it, that motivates the agent to obtain food.

Drive releasers represent control systems that maintain a controlled variable within a certain range. This variable is measured through some of the agent's sensors and compared 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, in the case of a mobile robot with two different batteries powering its base and the different actuators in its body, a *RechargingRegulation* drive would combine the signals from two different control systems: one controlling the base battery level and one controlling the body battery level.

The activation of a drive system is a nonlinear combination of its control systems as described in Equation (1).

Emotional Systems

Emotional systems represent different families of related affective responses, 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 on 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*, *Joy/Happiness*, *Disgust*, and *Surprise*. The selection of this core set of emotion types is not arbitrary, but rather it is based on evidence suggesting their universality, including distinctive universal facial expressions, as well as eight other properties [Ekman 1992].

In contrast to other models proposed to date that emphasize on cognitively generated emotions, we consider both cognitive and noncognitive releasers of emotion. The latter kind tend to be more associated with "physical" aspects which may only apply metaphorically into non-embodied agents, but which are nonetheless relevant for physical agents and other synthetic agents that at least simulate physical interactions with their environments. Influenced by

[Izard 1993], we divide all releasers of emotion 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.
- *Sensorimotor*: Includes sensorimotor processes, such as facial expressions, body posture, and muscle action potentials, that not only regulate ongoing emotion experiences but can also elicit emotion.
- *Motivational*: 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 neuropsychological studies by [Fox and Davidson 1986], pain or aversive stimulation causing anger, and emotions like sadness eliciting others such as anger.
- *Cognitive*: Includes all type of cognitions that activate emotion, such as appraisal of events, comparisons, attributions, beliefs and desires, and memories.

Currently, all releasers from the first three categories are natural releasers and all cognitive ones are learned releasers. In previous work, these cognitive releasers were based on a cognitive appraisal theory (See [Velásquez 1997] for details). In the proposed framework, however, emotional systems do not include any pre-wired cognitive releasers, but rather these are learned throughout the agent's interactions with the world as it is described below.

The main reason behind this change is that although suitable to test theories of cognitive evaluation of emotional experiences, computational models of emotion appraisal are very limited for the understanding on how the brain might process emotions. This is primarily due to the fact that these models are based on post-hoc cognitive recollections of emotional experiences and they are not concerned with the neural mechanisms behind such processes.

The activation of emotional systems differs slightly from that of the general form described above. Its main differences include considering the excitatory (positive) and inhibitory (negative) input from other emotional systems, as well as considering its temporal decay. This is summarized in Equation (2):

$$A_i(t) = f\left(\psi(A_i(t-1)) + \sum_k R_{ki} \cdot W_{ki} + \sum_l \mu_{li} \cdot A_l(t)\right) \quad (2)$$

Where $A_i(t)$ is the activation of emotional system i at time t; A_{it-1} is its activation at the previous time step; $\psi()$ is the function that controls the temporal decay of the activation of emotional system i; R_{ki} is the value of Releaser k, and W_{ki} is its associated weight, where k ranges over the set of releasers for emotional system i; μ_{li} is the strength of the excitatory (positive) or inhibitory (negative) input from emotional system l, where $A_l(t)$ is its activation value at time t; and f is a limiting function such as the standard ramp and logistic (sigmoid) functions.

Fast Primary Emotions. This model of emotional systems allows for the distinction between different affective phenomena. For instance, primary emotions are modeled as the activation, via natural releasers, of one particular emotional system such as *Disgust* or *Fear.*

These primary emotions play an essential role in the preparation of appropriate emotional responses that are adaptive for the agent. Such is the case, of the Fear emotional system, which, through these fast primary emotions detects a dangerous situation and generates defensive responses that maximize the probability of surviving it.

Emotion Blends and Mixed Emotions. Although there is no consensus on this issue, most basic emotion researchers believe that there are also nonbasic emotions that are the result of blends or mixes of the more basic ones [Plutchik 1994; Izard 1991]. For instance, according to Plutchik, fear and surprise would generate alarm, whereas joy and fear would produce guilt.

The mixing of these higher order emotions is generally conceived of as a cognitive operation or as the concurrent activation of both cognitive and affective systems. Although the proposed framework does not provide explicit models for these higher order emotions, some blends might emerge as the co-activation of two or more of the basic emotional systems. In such cases, it is possible that these co-active emotional systems bias one particular system (i.e., perceptual or behavior system) or several non-conflicting ones, as it is explained below for the case of behavior systems.

Emotional Memories and Secondary Emotions. Emotional systems have the capacity of acquiring learned releasers, which, as we previously mentioned, correspond to stimuli that tend to be associated with and predictive of natural releasers. The activation of emotional systems via these learned releasers correspond to what are referred to by some researchers as *emotional memories* or *secondary emotions*[1] [LeDoux 1996; Damasio 1994]. These secondary emotions occur only after we start making orderly associations between objects and situations, and primary emotions. They require more complex processing, including in some cases, the retrieval of emotional memories of similar previous experiences.

Secondary emotions also play an adaptive role in dealing with situations that have occurred over and over throughout evolution, such as escaping danger, finding food, and mating. For instance, while primary emotions might include pre-wired mechanisms for the detection of a potential mate and the ensuing mating responses, secondary emotions provide the means to associate learned stimuli (e.g., the place where a potential mate was last seen) with these same mechanisms.

Inspired by the work of LeDoux [LeDoux 1996], these emotional memories have been modeled with an associative

---

[1] Not to be confused with the higher order emotions described in the previous section, known as emotion blends and mixed emotions.

network comparable to Minsky's K-lines [Minsky 1986], in which salient stimuli (e.g., features, and percepts representing objects and agents) are connected to primary emotions when these have become active throughout the agent's interaction with the world.

During emotional learning, connections within this network are changed according to a modified Hebbian rule that prevents saturation of the connection weight between the newly created cognitive releaser and the active emotional system. The sum of the weights of all incoming learned connections to an emotional system are kept constant, through multiplicative normalization based on the existing excitatory and inhibitory connections between emotional systems.

Figure 2 illustrates a scenario in which an emotional memory is formed when a person interacts with a robot. In this example, petting and disciplining the robot are natural releasers for the *Joy* and *Fear* emotional systems, respectively. When the *Fear* emotional system becomes active because of the punishing action, the salience of all other present stimuli is determined and the most salient one, in this case the presence of a person, is associated with the *Fear* emotional system as a new learned releaser.

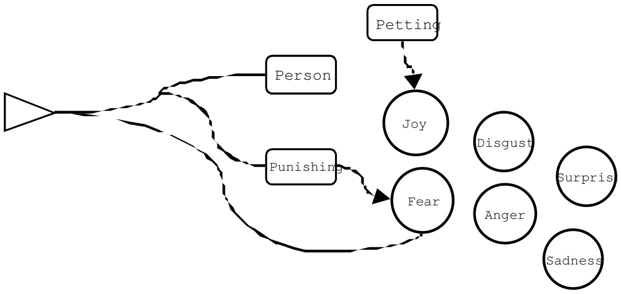

Figure 2 An Emotional Memory

These learned releasers also have the ability to act as biasing mechanisms during the action-selection process. Thus, the next time an agent encounters a *marked* stimulus, such as a person in this scenario, the memory represented in the associative network will be relived, reproducing the emotional responses that occurred previously, and influencing the selection of actions to follow.

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.

Moods. 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. 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.

Temperament. 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 "grumpy" agent, we might lower the activation threshold and decay rate for the *Anger* emotion as well as increasing those for the *Joy* emotion, and lowering the inhibitory gain between *Anger* and *Joy*.

Behavior Systems

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 [Adolphs et al. 1996]. This choice is responsibility of behavior systems, which represent interconnected, self-interested behaviors, such as *Approach-Person*, *Play*, *Request-Attention*, and *Avoid-Obstacle*.

Like drive systems and emotional systems, behavior systems also have releasers that obtain and filter sensory data in order to identify special conditions which will either increase or decrease their activity. Releasers might represent physical objects and specific conditions, such as "battery recharger is present", as well as motivational states such as "battery level is low" and "distress is high", which would most likely increase the activity of a *Recharge-Batteries* behavior.

Behavior systems may mutually inhibit or excite each other. For instance, *Wag-Tail* might inhibit *Running* and vice-versa. Whereas behaviors such as *Play-With-Person* might excite lower-level ones like *Find-Person*.

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 be executed at the same time. Given the parallelism of the model, we revised the Behavior System so that active, non-conflicting behaviors can issue motor actions simultaneously. The value for each behavior is computed as described in Equation (3):

$$B_j(t) = \sum_n (R_{nj} \cdot W_{nj}) + \sum_l (\mu_{lj} \cdot B_l(t)) \qquad (3)$$

Where $B_j(t)$ 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; μ_{lj} is the strength of the excitatory (positive) or inhibitory (negative) input from behavior system l, where $B_l(t)$ is its activation value at time t.

Implementation and Results

The framework for Emotion-Based Control described in the previous sections has been used to develop and control various synthetic agents, including *Simón the Toddler,* a synthetic character, and *Virtual Yuppy,* a simulated emotional pet robot [Velásquez 1997].

As part of our ongoing research, the same framework is currently being used to control *Yuppy,* an emotional pet robot shown in Figure 3.

The Sensory System

Yuppy has different sensors, including two color CCD cameras (currently using only one as part of its vision system), an active stereo audio system composed of two microphones mounted on Yuppy's ears, IR sensors for obstacle avoidance, an air pressure sensor used to model simple touch perception in the form of painful and pleasurable stimuli, a pyro sensor aligned with the top camera used to detect changes in temperature due to the presence of people, and a simple proprioception system.

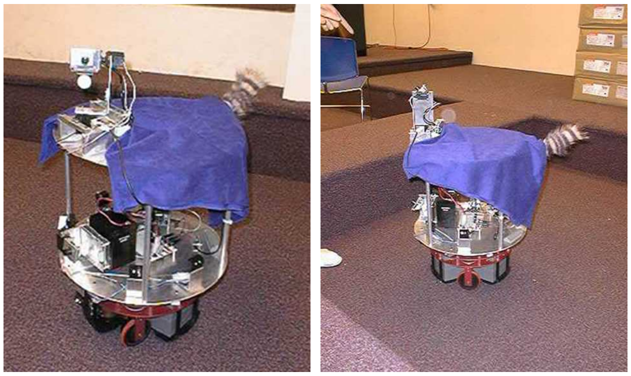

Figure 3 Yuppy, an Emotional Pet Robot

The Drive System

Yuppy's 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.

The Emotional Systems

Yuppy has emotional systems with natural releasers for the set of basic emotion families described before. Some of the most representative releasers and their associated emotional systems include:

- Interactions with drive systems: Unsatisfied drives produce *Distress* and *Anger; Distress* is also produced, to a lesser extent, when drives are oversatisfied; and Satiation of drives releases *Happiness*.
- Interactions with the environment: In general, all pink-reddish objects release *Happiness* to some extent, although pink Styrofoam "bones" elicit the most*;* Yellow objects found in the environment release *Disgust;* Darkness and blue Styrofoam horses (also known as the "Evil Pool-Ponies"), release activity in the *Fear* emotional system; and loud noises release *Surprise*.
- Interactions with People: When people interact with the robot, two of the several possible interactions include petting and disciplining the robot. These actions will generate representations of pleasure and pain, respectively. Pleasure releases *Joy*, and pain releases *Fear*, and *Anger.*

The Behavior System

The robot's Behavior System is composed of a distributed network of approximately nineteen different self-interested behaviors, directed in most part toward satisfying its needs and interacting with humans. Examples of such behaviors include *Search-For-Bone*, *Approach-Bone*, *Recharge-Battery, Wander, Startle, Avoid-Obstacle*, *Approach-Person, and Express-Emotion*.

Emotional Behaviors and Emotional Conditioning

Using the model described before, Yuppy exhibits emotional behaviors under different circumstances. For instance, when the robot's *Curiosity* drive is high, Yuppy wanders around, looking for the pink bone. When it encounters one, the activity of the *Happiness* emotional system increases and specific behaviors, such as *Wag-Tail* and *Approach-Bone* become active. On the other hand, as time passes by without finding any bone, the activity of its *Distress* emotional system rises and appropriate responses, such as *Droop-Tail*, get executed. Similarly, while wandering around, it may encounter dark places or evil pool-ponies, which will elicit fearful responses in which it backs up and changes direction.

Besides regulating action-selection and generating emotional behaviors through primary emotions, Yuppy learns secondary emotions which are stored as new or modified cognitive releasers based on the associative network model described before. In one instance, the *Fear* emotional system acquires a new releaser for loud sounds. In this classical scenario of *fear conditioning*, a natural releaser (Pain produced when a person disciplines Yuppy) generates a fearful response (*Cower* behavior). The sound stimulus by itself, however, does not produce any activation in the *Fear* emotional system, thus the *Cower* behavior does not become active either. If both stimuli are presented simultaneously, the *Fear* emotional system forms a new cognitive releaser for the sound stimulus. After only one trial, the newly formed releaser for the loud sound is capable of producing some activation of the *Fear* emotional system. After several more trials, the connection between the sound releaser and the *Fear* emotional system is strong enough to produce activation of the *Cower* behavior, and thus an emotional memory is formed.

In another scenario, after locating a pink bone (natural releaser) and approaching it, the robot interacts with the person carrying the bone. Depending on these interactions (e.g., the person pets or disciplines the robot), Yuppy will create positive or negative emotional memories with respect to people, and future selection of behaviors such as approaching or avoiding them will be influenced.

These results showed that emotional conditioning was possible under the proposed model, and that it could further be used as a biasing mechanism within the robot's action-selection process.

Conclusions

We have presented a computational approach to emotion-based decision-making that integrates models of perception, motivation, emotions, and behaviors and that focuses on emotions as the main motivational system that influences how behaviors are selected and controlled. A wide range of affective phenomena is modeled, including fast primary emotions, emotion blends, emergent emotions, mood, and temperament. Furthermore, we have showed how emotional conditioning can be used to create secondary emotions that act as biasing mechanisms during the process of making decisions and selecting actions.

Acknowledgments

Support for this research was provided in part by Yamaha Motor Corporation. Work on the physical Yuppy robot is carried out in conjunction with Professor Rodney Brooks, Charles Kemp, and Makoto Yoshida.

References

Abelson, R.P. 1963. Computer Simulation of "hot cognition". In: Tomkins, S.S., and Messick, S. eds. *Computer Simulation of Personality.* New York: Wiley.

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.

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.

Blumberg, B., Todd, P, and Maes, P. 1996. No Bad Dogs: Ethological Lessons for Learning. In: *Proceedings of SAB'96*. Cambridge: MIT Press.

Blumberg, B. 1994. Action-Selection in Hamsterdam: Lessons from Ethology. In *Proceedings of SAB'94*, 108-117. Brighton, England: MIT Press.

Cañamero D. 1997. Modeling Motivations and Emotions as a Basis for Intelligent Behavior. In: *Proceedings of AGENTS'97*. New York: ACM Press.

Colby, K.M. 1975. *Artificial Paranoia*. New York: Pergamon.Damasio, A. 1994. *Descartes' Error: Emotion, Reason, and the Human Brain*. New York: Gosset/Putnam.

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.

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

Evans, J et al. 1993. *Human Reasoning: the psychology of deduction*. Hillsdale: Lawrence Erlbaum.

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. TR-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. 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*.

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.

Plutchik, R. 1994. *The Psychology and Biology of Emotion*. New York: HarperCollins.

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

Staddon, J. 1993. On Rate-Sensitive Habituation. Adaptive Behavior, 1 (4). MIT Press.

Velásquez, J. 1996. Cathexis: A Computational model for the Generation of Emotions and Their Influence in the Behavior of Autonomous Agents. S.M. Thesis. MIT.

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

Wright, I. 1996. Reinforcement Learning and Animat Emotions. In: *Proceedings of SAB'96*. Cambridge: MIT Press.