An Emotional-Evolutionary Technique for Higher Level Goal Definition in a Multi-Purpose Artificial Creature

Patrícia de Toro[1], Ricardo Gudwin[1], and Mauro Miskulin[2]

[1] DCA/FEEC/UNICAMP
P.O. Box 6101 Campinas,Brazil, 13083-970
{patbgi,gudwin}@dca.fee.unicamp.br

[2] DSIF/FEEC/UNICAMP
P.O. Box 6101, Campinas-SP,Brazil, 13083-970
mauromiskulin@terra.com.br

Abstract. Artificial Creatures are embodied autonomous agents living in a virtual world, like e.g. in a computer game or in ethological simulation studies. Usually, these artificial creatures may have general high-level goals, like survival, killing opponents, feeding, mating, discoverying the world, etc. These general high-level goals must be turned into low-level goals in time and space in order for the agent to generate its final behavior within environment. This paper presents an emotional-evolutionary technique which encodes each high-level goal into a particular emotion, and using a blend of such emotions as a fitness function in a genetic algorithm, evolves particular goals in time and space in order to drive the creature's behavior.

1 Introduction

The research field of autonomous intelligent agents comprehends a powerful abstraction for many kinds of practical applications, from mobile robotics to computer games. One of the classical problems in this context is the problem of autonomous navigation in complex environments. This problem appears e.g. in the case of mobile robotics, where a mobile robot needs to decide a trajectory, since an initial point up to a target, without colliding with obstacles (possibly minimizing the distance covered and/or the time of travelling). Another example of the same problem would be the development of intelligent opponents in computer games, where an intelligent control system must decide the actions to be provided by an agent in order to foster a good entertainment to the system user, simulating with realism the behavior of a human opponent.

Several alternative frameworks are possible in order to deal with this problem, depending on the aspects we want to emphasize. One possible form to systemize the problem is to consider it in the sense of an artificial creature which exists in a certain environment, moving itself in this environment and acting on it [1].

Classical solutions for this problem are well known in the field of artificial intelligence [14] usually involving state machines, search algorithms and, eventually, logic programming. More advanced algorithms may use neural networks, fuzzy logic or evolutionary computation.

More recently, the concept of emotion, as brought from cognitive psychology and philosophy, was suggested in the literature, as an alternative way of dealing with this problem [2, 3, 13, 12, 19, 5, 6, 20, 15, 16, 4, 10].

There is no consensus, though, on what exactly are emotions. Different approaches have different views for what it is and how to model them. For example, Ortony [11] views emotions as "valenced reactions to events, agents, or objects, with their particular nature being determined by the way in which the eliciting situation is construed". Sloman [17, 18] understand emotions as internal "alarms" which give a momentary emphasis to certain groups of signals. Damasio [7, 8] make a distinction between "emotions", which affect the body and "feelings", which are a cognitive introspection of an emotion. Other authors may have further different views for what emotions are. For some of them, emotions work like "amplifiers" for motivations. For others they are homeostatic processes related to physiological variables [5]. Some authors, instead of a single concept of emotion, develop a complex "emotional system", where many different concepts like "motivations", "drives", "impulses", "affections", "needs" and other terms are used to represent different aspects of this emotional system.

In artificial creatures [3], emotions are usually employed in order to drive behavior, being used as a criteria for action-selection mechanisms. A problem, though remains to be solved. In a multi-purpose artificial creature (e.g. a character in a computer game), there are many different higher-level goals which need to be satisfied, like survival, killing opponents, feeding, mating, discoverying the world, etc. The creature needs to decide, at each point in time and space, what to do next, and so decide the next lower-level goals, like where to move to, actions to be performed, etc.

In this article, mixing ideas from emotion-based control systems and evolutionary computation, we present a technique for higher level goal definition in the case of a multi-purpose artificial creature.

The architecture in which this work is based on, was suggested originally in [9], which defines an autonomous creature equiped with a sensorial and motor apparatus capable to navigate through an environment full of objects with different characteristics. Objects do have "colors", such that each "color" is associated to: a degree of "hardness" (varying from 0 to 1), which works like a friction coefficient in order to slow down the creature movement (or totally block it) over places where there are objects of this color; a "taste", which can be bad or good (varying from -1 to 1); and an ability to drain or supply "energy" (which also may vary from -1 to 1). The creature owns an internal battery which is rechargeable. Touching an object which supplies or drains energy would change the charge of this battery. The creature can navigate in this environment, and many goals may be prescribed. An elementary goal will be to navigate from an initial point up to a target point, avoiding collision with some undesirable objects at the envi-

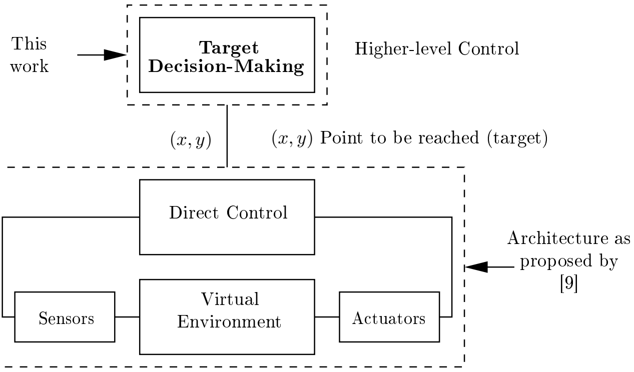

Fig. 1. Artificial Creature Architecture

ronment. There may be, however, some cases in which this collision is desirable (as in the case of objects that work as energy suppliers). The general goal would be to generate the agent behavior according to the different purposes that may be attributed to it.

Two levels of control have been considered: a lower level of direct control and a hierarchically superior level, of target decision-making (see Figure 1). At the level of direct control, the input is a target point to be reached by the creature. The creature must move from its current position up to this target point, without collisions with obstacles. At the level of target decision-making, the system must define the target points to where the creature must move itself, considering a set of higher level general purposes for the creature's existence. Together, the levels of direct control and target decision-making generate a complex behavior for the creature. The level of direct control was extensively addressed in [9], by means of a reactive/deliberative strategy which is responsible for the generation and execution of plans for the navigation to the target point. In this work, we extend the work of [9] by developing the superior level of target-decision-making.

The level of direct control only deals with the low-level goal of safe moving the creature along a set of obstacles. The level of target decision-making, allows for the consideration of a higher levels set of general purposes for the existence of the creature. Examples of higher level purposes in our case are: to explore the environment, taking care of the creature's energy balance and to learn an optimal policy regarding the distance to maintain to obstacles in order to have a safe navigation. To generate a behavior which considers together all those purposes, we made use of three emotional metaphors: curiosity, hunger and fear. Curiosity will be responsible for the creature's exploratory behavior. Hunger will be responsible for the maintenance of the creature's energy level. And fear must fine-tune the collision avoidance behavior.

2 Experiments and procedures

A sketch of the creature's sensory-motor structure can be seen at Figure 2. The creature does not have an "a priori" map of the environment. Based on the sen-

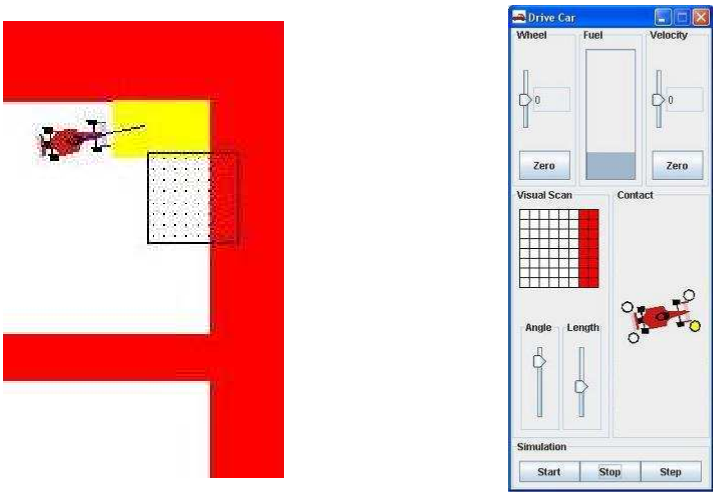

Fig. 2. Sensory-Motor Structure of the Creature

sorial information, the creature builds an incremental map of the environment, which is used to generate movement plans. A more detailed account of the creature structure and dynamic model used in the simulator and the internals of the low level direct control can be found in [9].

In this work, extending the work in [9], we deal with the higher level target decision-making sub-system, which defines at each moment, the next point to where the creature should move. According to [1], an artificial creature may have four types of standard behavior: (i) appetitive: behavior is directed towards an attractive object or situation; (ii) aversive: behavior is directed away from negative situations; (iii) exploratory: behavior is directed toward stimuli that are novel in the environment and, finally, (iv) neutral behaviors relating to objects that are neither appetitive nor aversive. It is interesting to note that our choice of "hunger", "fear" and "curiosity" corresponds to the choice of an appetitive, an aversive and an exploratory behavior. In this experiment, we do not contemplate the "neutral" behavior.

The level of direct control (as in [9]), already incorporates the emotion "fear" during the generation of a plan. This plan only considers the objects it already knows to exist at the environment. However, in [9] this fear was definided by heuristic means, encoded by a given (hard-coded) utility function. Here, we developed a more flexible mechanism, where this fear can increase or diminish, as such as the creature enters in contact with objects that represent obstacles. Or, in other words, the creature learns to have "fear" while interacting (colliding) to unpleasant objects. Moreover, this fear is also used at the determination of the creature's next target. The determination of the next target is not a simple task. Depending on the situation, different behaviors can be desirable. In the case which the creature's energy level is high, we would expect the creature to develop an exploratory behavior, going to parts of the environment which are not already known. However, in the case which the energy level is low, we would expect the creature to search within its map for a known power source and go to there. After being re-energized, the creature may return to the exploratory behavior. To implement this hybrid behavior, we developed two distinct strategies, one which is hard-coded, explicitly creating some rules of priorities among the emotions, and another one that makes an automatic blend of all emotions and derives a final decision, based on a genetic algorithm.

2.1 Description of the Hard-coded Algorithm

In our first experiment, emotions are considered in a hard-coded way, giving explicity priority on some emotions before others.

In this algorithm, if the energy level is higher than a certain level, a random point is chosen to be the next target. If the curiosity function associated to this point is below a certain level, than another random point is generated, until a qualified point is reached. The algorithm uses then the direct control defined in [9] to develop a plan which the creature puts into execution until the target is met. This step implements only implicitly the emotions of fear and curiosity. Fear is just a metaphor for the utility function which is used to derive the plan by direct control. Curiosity is also just a metaphor for the utility function which computes the degree of "awareness" of a given point. This is the standard behavior for the creature.

However, in the case which the energy level is below a certain limit, a new behavior starts to run. The creature stops seeking its standard target, saving it to resume in the future, and start looking for the nearest point (in its map) where it knows there is an energy supplier located. It uses than the direct control to build a plan to reach this point. After getting in contact with the energy supplier, the creature waits until its energy level increases up to the maximum, and after that, it restores the original target back, with a new plan given by the direct control.

It may be observed that the hard-coded algorithm imposes explicitly a fixed way of interaction among the emotions, leading to a fixed set of behaviors, according to the following rules:

- Fear (to collide with obstacles): the algorithm normally search to prevent collisions. If a collision happens, the agent increases its “caution factor”;
- Curiosity (search for unknown points at the environment): Curiosity measures whether a given point is already known by the agent;
- Hunger (to keep energy balance): Hunger leads to an exception behavior, where curiosity is disconsidered in order to generate the next target.

The hard-coded algorithm has many limitations. Perhaps, the most obvious one is that the emotional metaphor, although useful in order to inspire the algorithm development, does not allow for a generalized way to deal with different situations. The interaction among emotions need to be treated case by case. When it is desirable to include new emotions within the system, these new emotions shall require an extended specification on how it will interact with all the other already implemented emotions. The treatment of conflicting emotions

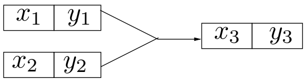

$$\begin{cases} x_3 = \alpha x_1 + (1-\alpha)\, x_2 \\ y_3 \ = \alpha y_1 + (1-\alpha)\, y_2 \end{cases}$$

Population Size = 1000 individuals
Crossover Prob.= 0,5
Mutation Prob.= 0.1

Fig. 3. Genetic Algorithm Parameters and Operators

may require the generation of exceptional behavior, which may require a special treatment nearly case by case. To solve these restrictions, we tried to develop an algorithm which could be generalized in an uniform way, not requiring special behaviors, and allowing for the addition of new emotions without too much effort. The result is the following algorithm.

2.2 Description of the Genetic Algorithm

In order to solve the aforementioned problems, a genetic algorithm (Figure 4) was conceived, where a population of possible target points shall evolve, using as a fitness function a combination of all the utility functions, one for each designed emotion. This combination results in a function of "desirability", that will mix fear, hunger and curiosity.

Initially this population of "potential targets" is generated randomly. At each evolutive step, these "potential targets" start to evolve, becoming confined to some regions of space where the "desirability" is higher. For the genetic algorithm, each target is encoded in the form of a chromosome containing a position (x,y) in the two-dimensional space. The crossover between two chromosomes is made by means of a linear combination between the two points represented by each chromosome, resulting in a point that is located somewhere in the straight line which bounds the two points (see Figure 3). A mutation operator translates a given point somewhere into its near neighborhood. All the possible goals are evaluated regarding its "desirability", and a process of elitist selection only keeps in the population the potential targets with greater "desirability".

Figures 5, 6 and 7 present the details for an example of simulation using the genetic algorithm. Figure 5 shows the environment where the creature develops its navigation. Figures 6 and 7 represent the internal map constructed by the creature, in different instants of the simulation. Figure 6 shows a situation in the

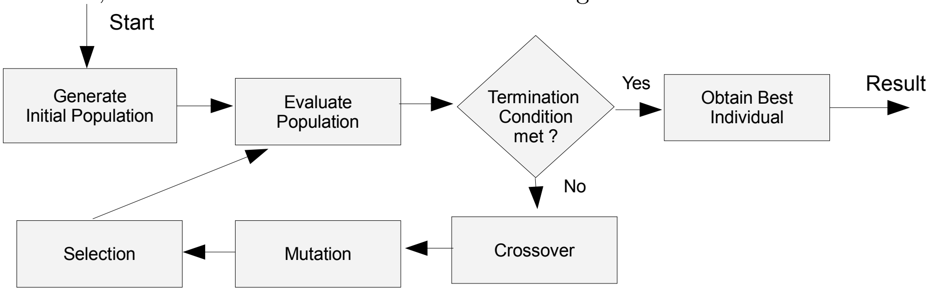

Fig. 4. The Genetic Algorithm

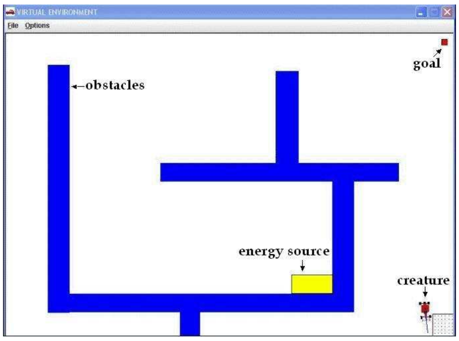

Fig. 5. A Plain View of the Environment

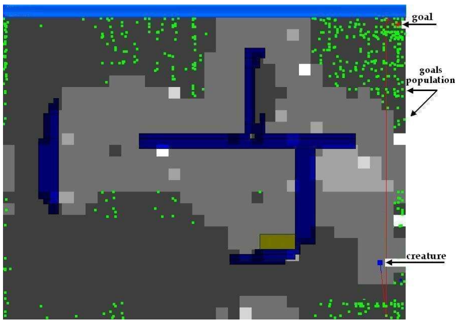

Fig. 6. Representation of the Emotional Factor *Curiosity* (different levels of gray). Green spots are the possible targets

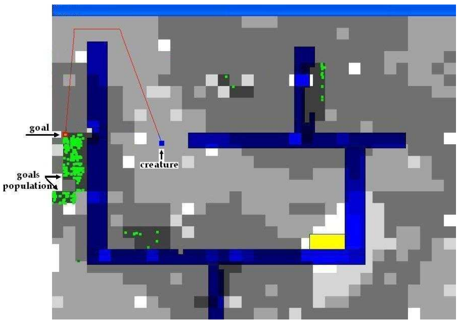

Fig. 7. Optimized Situation: creature with a higher knowledge of the environment

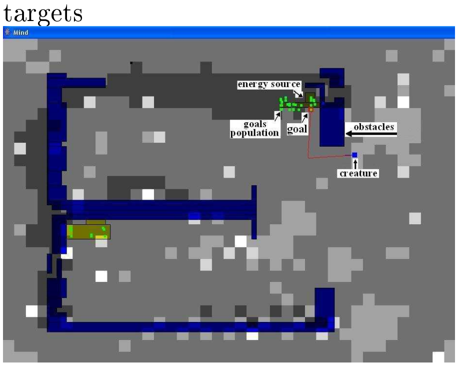

Fig. 8. Situation where the energy level is low: Targets are concentrated around energy suppliers

beginning of the simulation, where the dark areas represent unknown parts of the environment, implying in great curiosity and so great desirability for possible targets. Please pay attention to the distribution of light-green points inside these areas. They represent the population of possible targets being evolved by the genetic algorithm. Figure 7 shows a further time in simulation, where the creature already had the chance to know other areas. Observe that in this case, the number of dark areas is lower, and the possible targets are concentrated in a smaller range of locations.

Figure 8 shows another simulation (in a different environment), in a situation while the creatures's energy resources are too low. Notice that, in this case, the locations where there are objects with energy supply do attract a bigger desirability. Pay attention to the concentration of possible targets under these objects.

3 Simulation Results

In order to evaluate the simulation results, we generated some measurings, which we present in the following. Figures 9 and 10 show the variability of the creature's

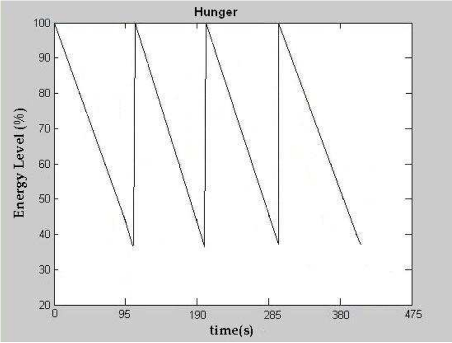

Fig. 9. Variation of Creature's Energy Level over time: Hard-coded Algorithm

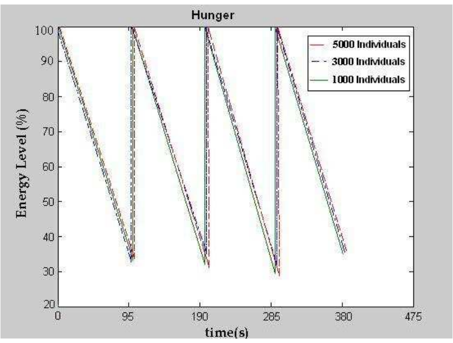

Fig. 10. Variation of Creature's Energy Level over time: Genetic Algorithm

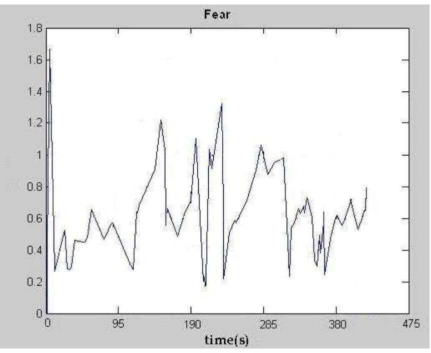

Fig. 11. Variation of Creature's Fear over time: Hard-coded Algorithm

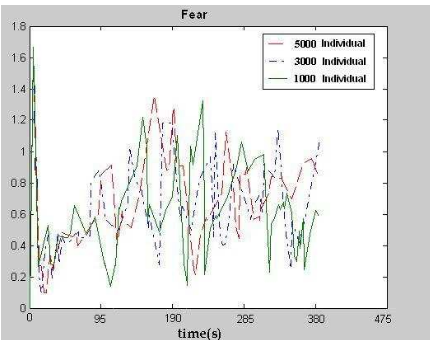

Fig. 12. Variation of Creature's Fear over time: Genetic Algorithm

energy level over time. Figure 9 shows the results for the hard-coded algorithm and Figures 10 shows the result for the genetic algorithms with populations of 1000, 3000 and 5000 individuals. Observe the behavior in both cases. The creature starts will full power and have the energy decreasing until it reaches the level of around 30 or 40 % when it then get refueled. The behavior is more or less the same both for the hard-coded algorithm and the genetic algorithm.

Figures 11 and 12 show the variability of the creature's fear over time. Fear is computed as a function which calculates the distance to the closest undesirable object at each time. We could not get any kind of correlation between these graphics, showing that completely different solutions where performed by each case of the algorithms.

Figures 13 and 14 show the variability of the unknown parts of the world over time. Figure 13 shows the case for the hard-coded algorithm and Figure 14 shows the case for the genetic algorithms, with populations of possible targets of 1000, 3000 and 5000 individuals.

We can see that in both cases, the creature increases its knowledge of the environment up to 100% after some time. In this case, we detected a small difference in efficiency among the many algorithms we simulated. The hard-

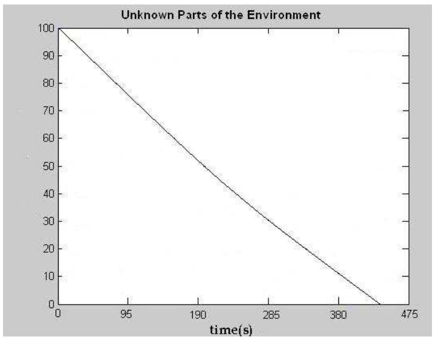

Fig. 13. Variation of Unknown Parts of the Environment over time: Hard-coded Algorithm

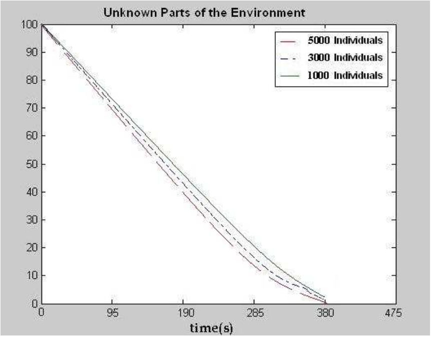

Fig. 14. Variation of Unknown Parts of the Environment over time: Genetic Algorithm

coded algorithms is less efficient, gaining total knowledge of the environment only around 430s. The genetic algorithm with population size of 5000 individuals get the same 100% with around 380s. The genetic algorithm with population sizes of 3000 and 1000 individuals also performed better than the hard-coded algorithm

As a general evaluation, we may conclude that both the hard-coded and the genetic algorithm do have a similar behavior, achieving their main generic goals of exploring the environment, maintaining their energy level and avoiding collisions with obstacles, with a slight difference to the genetic algorithm. But, on the other side, the genetic algorithm has a great advantage in terms of scalability. For the hard-coded algorithm, for each new emotion that we would like to include, all the algorithm needs to be redesigned. In the case of the genetic algorithm, only a small change in the desirability function is required, if we intend to include a newer emotion. As a conclusion, we may reach that the genetic algorithm is more advantageous than the hard-coded one.

4 Conclusions

In this paper, we presented an emotional-evolutionary technique for higher level goal definition in a multi-purpose artificial creature. This technique encodes each higher level goal into a separate emotion, and then a genetic algorithm is used to evolve a population of candidate lower-level goals, in order to drive the final creature behavior. This strategy, up to a certain point, would be comparable to how lower-level human goals are chosen amongst a population of candidate ideas. Each idea is evaluated by means of the emotional value it is able to arouse, and the most prominent idea, in emotional terms, is chosen to be performed. There is a random component, brought by genetic algorithm, and also some kind of intelligent behavior, as the emotions are higher-level goals encoded in a normalized form. We understand that the current experiments only suggest this insight, being still not definitive. But the present study shows the viability of such technique, and motivates us to continue this line of investigation.

5 Acknowledgments

Patrícia de Toro is thankful to the CAPES for the financial support in the form of fellowship. Ricardo Gudwin would like to thank the Brazilian National Research Council (CNPq) for grant #304649/2004-0

References

1. C. Balkenius. *Natural Intelligence in Artificial Creatures.* Lund Univ. Cognitive Studies 37, 1995.
2. J. Bates. The role of emotion in believable agents. *Communications of the ACM*, 37(7):122–125, 1994.
3. B. M. Blumberg. *Old Tricks, New Dogs: Ethology and Interactive Creatures.* PhD thesis, MIT Media Lab, Cambridge, MA, 1996.
4. D. Budakova and L. Dakovski. *Computer Model of Emotional Agents*, volume 4133. Springer Berlin / Heidelberg, 2006.
5. L. Cañamero. A hormonal model of emotions for behavior control. In *Proc. of the 4th ECAL (ECAL'97)*, 1997.
6. L. Cañamero. Issues in the design of emotional agents. In A. Press, editor, *In: Proc. of Emotional and Intelligent: The Tangled Knot of Cognition. AAAI Fall Symposium*, pages 49–54, Menlo Park, CA, 1998.
7. A. R. Damásio. *O Erro de Descartes: Emoção, Razão e Cérebro Humano.* Europa-América, 1994.
8. A. R. Damásio. *The Feeling of What Happens: Body and Emotion in the Making of Consciousness.* New York: Harcourt, 1999.
9. R. R. Gudwin. *Contribuições ao Estudo Matemático de Sistemas Inteligentes.* PhD thesis, DCA-FEEC-UNICAMP, 1996.
10. J. J. C. Meyer. Reasoning about emotional agents. *Int. J. Intell. Syst*, 21(6):601–619, 2006.
11. A. Ortony, G. Clore, and A. Collins. *The Cognitive Structure of Emotions.* Cambridge Univ. Press, 1998.
12. R. W. Picard. *Affective Computing.* MIT Press, Cambridge, 1997.
13. W. S. N. Reilly. *Believable Social and Emotional Agents.* PhD thesis, Carnegie Mellon University, Pittsburgh, PA, 1996.
14. S. J. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach.* Prentice Hall, Englewood Cliffs, NJ, 1995.
15. L. M. Sarmento. An emotion-based agent architecture. Master's thesis, FC University of Porto, 2004.
16. C. Septseault and A. Nédélec. A model of an embodied emotional agent. In *IVA*, page 498, 2005.
17. A. Sloman. Damasio, descartes, alarms, and meta-management. In *In Proc. of the IEEE International Conf. on Systems, Man, and Cybernetics*, pages 2652–2657, 1998.
18. A. Sloman. Beyond shallow models of emotion. In *Cognitive Processing*, volume 2, pages 177–198, 2001.
19. J. Velásquez. A computational framework for emotion-based control. In *In Proc. of the Grounding Emotions in Adaptive Systems - workshop SAB'98*, Zurich, Switzerland, 1998.
20. R. Ventura. Emotion-based agents. Master's thesis, Inst.Sup.Tec.-Lisboa,Portugal, 2000.