BDIE: a BDI like Architecture with Emotional Capabilities

David J. Hernández and **Oscar Déniz** and **Javier Lorenzo** and **Mario Hernández**
Institute for Intelligent Systems and Numerical Applications in Engineering
Universidad de Las Palmas de Gran Canaria
Edificio Central del Parque Científico-Tecnológico
Campus Universitario de Tafira, 35017, Las Palmas, Spain
davhernandez@isusiani.ulpgc.es
{odeniz, jlorenzo, mhernandez}@dis.ulpgc.es

Abstract

In this paper, a first version of the BDIE architecture is presented. This architecture is an attempt to model affective phenomena and adds a new emotional component to a BDI agent. The new component affects to the others trying to take advantage of the use of emotions for the development of artificial agents. The architecture has been implemented in an object-oriented framework that allows to easily define new agents. As an example of an agent, initial testing has been done with a robotic head.

Introduction

In the last years emotions have been receiving increasing interest from the Artificial Intelligence research community. Several authors are turning their attention to the study of emotions mainly due to two reasons:

1. The discovered role of emotions in rational thinking (Damasio 1994): neurological experiments have demonstrated the importance of emotions in it.
2. The interest in affective computing (Picard 1997): the idea is that if computers can recognize, express and probably have emotions, the interaction with the user will be improved because the computer can analyze the affective state of the user and choose a different action course depending on it.

Today is widely accepted that emotion is important in rational human activity. As pointed out in (Alvarado, Adams, & Burbeck 2002), the question is no longer whether a designer of an architecture for artificial agents should or not incorporate emotions, but how to do so.

We are interested in the role of emotions in the field of artificial agents. First, we have worked in human-robot interaction and the introduction of emotions would improve that interaction (Déniz & others 2002) due to the robot's capability to express emotions. Second, we also have interest in the application of artificial agents to computer games, where emotional capabilities would allow a much better user experience because agents would have a more human decision making process. Finally, the use of emotions guiding the decision-making process would improve its performance as argued below.

This paper presents the BDIE architecture, an attempt to capture the effects of emotions in the human mind and reproduce them in artificial agents. The main goal is to develop and evolve an architecture that should allow the creation of artificial emotional agents and use them in several scenarios to test their suitability.

The Base: a BDI Architecture

The approach starts from the BDI architecture (Rao & Georgeff 1995; Kinny, Rao, & Georgeff 1996) and adds to it the necessary components to endow it with emotional capabilities. This decision was made based on the following criteria:

1. This architecture clearly separates what conceptually every agent must have, that is, perceptions, goals, and behaviors. This allows a clear distinction of the different data structures that in other architectures are mixed. Thus, it is possible to be focused on each component for its design and implementation.
2. The planning capabilities allow the agent to produce plans as complex as desired. The result of the decision making process can be just the next action or a complex sequence of actions to execute over time. This allows to change the planning algorithm according to the needs. Generated plans may be used to compute the effect produced by a giving action sequence over the emotional system as pointed out in (Gratch 1999).
3. The architecture is sufficiently powerful to achieve the design goals to be reached, that is, to develop an architecture that produces a more human decision making process by adding the capabilities to have, recognize and express emotions.

Figure 1 shows a basic view of the BDI architecture's components. There are two types of information components:

- Data containers which store the agent's information. Three main stores are found:
  1. Beliefs. This container stores what the agent knows about its environment (perceptions) and, possibly, about itself (proprioception).

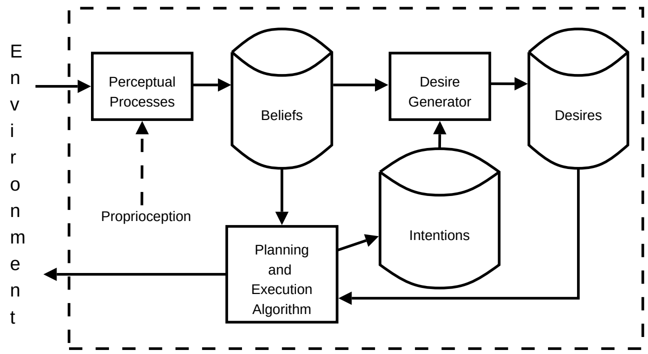

Figure 1: The BDI architecture components

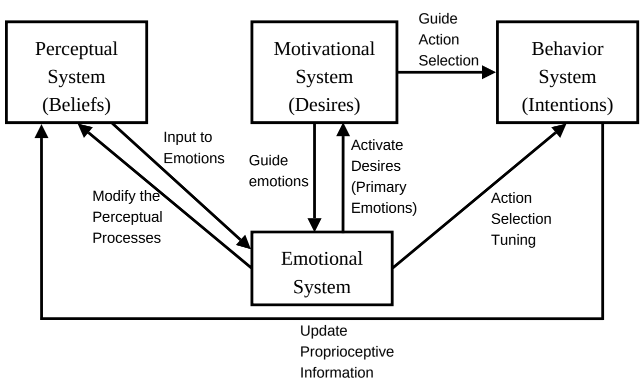

Figure 2: BDIE Functional View

2. Desires. In this container, the agent's goals are found. Goals guide the decision making process.
3. Intentions. The decision making process produces a plan that will be executed over time. This plan and its related information are stored in this container.

- Data processing algorithms:
  1. Perceptual Processes. These include the sensors that receive the information from the environment and the own agent and update the beliefs.
  2. Desire Generator. This process may create new desires given the current desires and intentions.
  3. Planning and Execution Algorithm. In this module the planning algorithm generates a plan to achieve the current goals while another algorithm executes the different components of this plan.

BDIE

Figure 2 shows a functional view of the architecture with the different modules, including the emotional one, and the relationships between them. The new module contains both data containers and processes that will be introduced below. As shown in the figure, it has dependencies with all the other components in the architecture:

- Relationships with the perceptual system:
  - Its inputs are the belief which, in turn, come from the sensors.
  - It modifies the perceptual processes, representing that the same input stimuli may be perceived differently depending on the affective state. For example, if the active emotion has an associated high arousal state, a stressful situation, the sensors may stop working.
- Relationships with the motivational system:
  - Current goals affect the current emotional state. For example, when a goal is satisfied, it would be possible to have an emotional response of happiness.
  - It activates goals when a primary emotion becomes active in order to respond in presence of a special situation. This goal will guide the decision making process to the execution of a behavior that will try to preserve the agent's welfare. Our decision to create this link came from two side. On the one hand, as argued in (Canamero 2000) emotions should be connected to intentions through goals in order to increase the architecture's modularity. On the other hand, as pointed out in (Castelfranchi 2000), emotions may be sources of goals as one of the roles in the complex relationship that exists between them. However, the fact that emotions do not directly communicate with the intentions module does not mean that fast responses, as reflexes, can not be implemented. In the BDIE architecture, whenever a primary emotion (see below) is activated, a high priority desire is unsatisfied. This produces that the planning algorithm turns all its resources to solve the current situation. Additionally, limitations in the algorithm's processing time in order to satisfy the real-time necessity of a behavior execution are planned to be introduced if needed. For example, in a behavior network (Maes 1989), that could means, that the energy is not allowed to jump from one behavior to other more than a certain amount. This produces a more reactive (vs more deliberative) response of the net.
- Relationships with the behavior system:
  - It tunes the action selection algorithm. In this case, the emotional system acts as a *metaheuristic*, a mechanism that guides the action selection algorithm through the vast space of available options to execute. Conceptually, it acts as a pruning method of the state space, removing those options that are not *emotionally good*. This allows the algorithm to concentrate on a smaller set of options, discarding the emotionally bad ones. Additionally, certain experiments have shown that the emotional state affects the capability of individuals to solve problems so that, for example, when the individual is sad it takes more time to solve a problem than when it is happy. In this line, work is being carried on a dynamic reconfigurable behavior network whose parameters, structure and update cycle will depend on the current affective state.
  - It modifies the execution of a behavior producing a more *affective output*. For example, it would be possible to speak with a different intonation or to show a

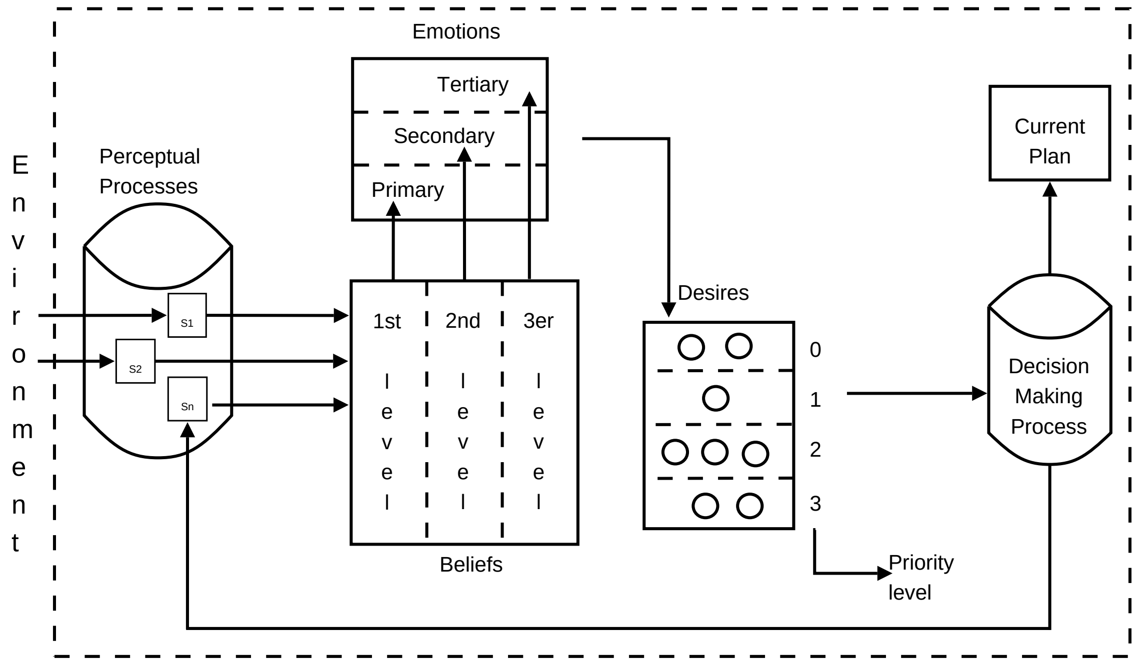

Figure 3: BDIE Structural View

facial expression.

Figure 3 shows a structural view of the architecture with all the components which will be detailed in the following sections.

Emotional system

The emotional system keeps track of the agent's affective state. In BDIE the Damasio's (Damasio 1994) and Sloman's (Sloman 2000) views have been combined. There are three types of emotions:

- *Primary emotions.* The primary emotions are related with *fundamental life tasks* so that the individual can survive. They are fast and prewired mechanisms that are activated whenever a very specific situation occurs. The primary emotions take control over the normal decision making process in order to rapidly execute a behavior that guarantees the agent survival. Examples of primary emotions are fear and surprise.
- *Secondary emotions.* This emotions are related with high level cognitive processes that interpret the perceptual information and modify the affective state according to that information. They are associated with other functions of emotions such as interaction and communication with other individuals. Typical examples of this category are happiness, sadness and anger.
- *Tertiary emotions.* Currently not implemented, they are included for completeness in the design. They are related to reflective data and may include emotions as shame or pride. They will be the next step of the BDIE architecture.

Although this classification exists in BDIE, emotions are not viewed as processes that evaluate the input stimuli. In contrast, emotions are the result of other processes which make such evaluation. These other processes perform the affective appraisal and cognitive evaluation of input stimuli (Castelfranchi 2000). Affective appraisal is a fast and often unconscious evaluation of the input stimuli, usually implemented by reactive mechanisms trained to act quickly in certain situations. In contrast, cognitive evaluation is a more elaborated evaluation of the stimuli and is usually performed by high level processes.

These concepts are mapped in BDIE introducing *evaluators*. There are first level evaluators, those which perform the affective appraisal, and second level evaluators, those that perform the cognitive appraisal. They are associated respectively with primary and secondary emotions. First level evaluators look for first level beliefs (see below) and detect special situations that fire a primary emotion. Usually, they are simple processes such as threshold exceeding tests. On the other hand, second level evaluators are more complex processes that look for several first level beliefs and combines them into a second level one. They may probably use information from other systems such as desires or long term memory.

As these processes are in charge of stimuli evaluation, emotions are just the output of them. This conceptually fits in both, the psychological and the software engineering fields (allowing a clearer design and implementation).

The Emotional Space Classically, there have been two representations for emotions in artificial agents:

- A set of discrete emotions where each one is represented separately. This representation comes mainly from the work of Paul Ekman (Ekman 1999) where there is a basic set of six emotions. This set is the base of a six dimensional space where all the other emotions can be placed as a combination of the basic ones. Examples of this representation can be found in (Velásquez 1997; 1998b; 1998a).

- A continuous two dimensional space with *valence* and *arousal* as the axes (Breazeal 2002). In this case, the space has only two dimensions and each affective phenomenon has a pair of valence and arousal associated values that allows to move the current emotional state, a point in the space, within it. The valence represents how favorable or unfavorable is the stimulus to the agent while the arousal means how arousing is it. Sometimes, a third axis associated with the stance is added to this scheme, representing how approachable is the stimulus.

Actually, the two representations are not exclusive, that is, they can be combined. (Breazeal 2002) is an example of this combined representation. In it, each emotion has an associated area of the space. If the current point, representing the current emotional state, is inside one of these associated areas, it is said that the emotion is the current one. We also have decided to use this type of representation because it allows us to easily obtain emotional information from the perceptual information. As each input stimulus has an associated pair of valence and arousal values, it can serve as an input to the emotional system that moves the point in the space representing the current emotional state according to the values of new stimuli. Primary emotions, due to its necessity to act quickly, provide an absolute position in the space so that the affective state changes immediately. Secondary emotions provide movements relative to the current affective state. Thus, it changes gradually from one emotion to another according to the input stimuli.

In BDIE, the emotional space is divided in *emotional sectors*. Each emotion delimits a sector in the space. Figure 4 shows an already divided emotional space. Each sector is described by:

- A *descriptive point*. A point in the space whose module indicates the emotion's maximum intensity. The points are used to compute the angles that conform each sector. The maximum intensity is the sector outer arc's radius.
- A *minimum intensity* that indicates when the emotion is considered to be inactive. This value is the sector inner arc's radius.
- A *decay factor*. It is applied to the current point when no input stimulus has activated an emotion. In that case, the absence of activating stimuli produces that the current emotional point tends to go to the origin. The decay factor is an amount by which the current point's components are divided.

Additionally, each primary emotion has an extra configuration parameter that does not participate in the emotional space. This parameter is the name of the desire dissatisfied when the emotion is active.

When the system is loaded, the emotional space is computed given all the emotion's descriptions. The neutral emotion is considered active when the current emotional point is in the space around the origin, where no emotion is defined. This current emotional point is located anywhere in the space and is moved according to the input stimuli. As said, it may suffer absolute or relative movements depending on the output of the affective processes.

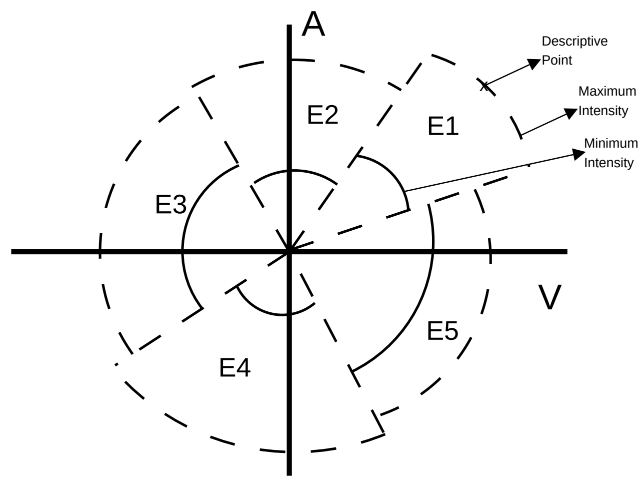

Figure 4: The emotional space

The emotion's parameters allow the agent designer to define personalities. For example, indicating a lower decay rate means that an emotion that is active, will remain active longer than another emotion with a greater decay factor. Over the time, that means that is likely the agent will experiment during more time the emotion with lower decay factor. This is a way to model a *mood*. Also minimum and maximum intensities determine if the agent is more or less sensible to certain stimuli, producing a agent which usually experiment certain emotions.

Perceptual processes

Perceptual processes include all the processes that accept input from sensors and generates first level beliefs. That includes any kind of perceptual processing, which makes them compatible with any kind of attention filter or habituation processes. The designer is able to implement different sensors and add to them the functionality he/she wants. BDIE supports threaded sensors that run in their own threads which are useful in environments as robotics. As indicated below, the retrieval of sensors' current values is the first step in the agent's update cycle.

Beliefs

As said before, beliefs represent the agent's vision of its environment and itself. They are a set of elements like perceptual information arrived from the perceptual processes. In BDIE there are three conceptual levels of beliefs each one with a higher semantic level:

- *First level* beliefs represent basic perceptual information, basic stimuli come from perceptual processes. They have not gone through any post-processing algorithm that would try to infer more information about the input. Each belief will contain the sensor information and, after passing through the emotional module, the valence and arousal evaluation. As an example, an *image* and an *audio signal* coming respectively from a camera and a microphone could be first level beliefs. The image belief could have a high arousal value if an object is detected close to the agent. Similarly, the audio signal could have a high arousal value if the volume is too loud.

- *Second level* beliefs represent more elaborated beliefs. They are the output of cognitive processes that combine several first order beliefs into one or more second level ones. Following with the previous example, an *alert* second level belief could be generated from the two first level belief *image* and an *audio signal*. The *alert* belief would obtain a resultant valence and arousal from a combination of the source's two values.
- *Third level* beliefs will be the future next step of our architecture and will be related to *tertiary emotions* (Sloman 2000). They will work in a similar fashion to second level beliefs. There will be cognitive processes that will combine different second level beliefs into one or more third level beliefs. As tertiary emotions, they are not currently implemented.

Desires

In BDIE, desires comprise goals and homeostatic variables. Goals are desires based in a predicate. Whenever a goal's predicate returns true, the systems considers that the goal is satisfied. The inclusion of homeostatic variables, as another element that is able to guide the decision making algorithm, is because they are very useful in all kind of agents. Thus, they can be found in several architectures to model concepts like thirst or hunger. Homeostatic variables act like any other goal: when an homeostatic variable is in its normal range, it is considered as satisfied so it does not try to guide the decision making process. Whenever it goes over or under the normal range it will be considered as unsatisfied and will try that the decision making algorithm executes a behavior that satisfy it.

In several cognitive architectures, desires have priority values that classify them and guide the decision making algorithm to execute behaviors that will satisfy higher priority desires. In BDIE this concept is supported but raised to a higher level. Desires, that is, goals and homeostatic variables, have an importance value but, besides, they belong to a given priority level. As shown in Figure 3, several priority levels may exist in a given agent. The decision making algorithm will try to satisfy the higher importance unsatisfied desire in the higher priority level. That means that if desire A has a higher importance than desire B but B is in a level with higher priority than A, then the decision making algorithm will try to satisfy B first.

Whenever a desire of a given level is unsatisfied, all the goals from the levels with lower priority are *deactivated*, even not allowing them to guide the decision making algorithm. This scheme allows the decision making process to be focused on the desires that really matter in the current situation. For example, in a dangerous situation the high level desires should be left away just to invest all the resources in solving such situation. That is what happens, for example, if a primary emotion related to a survival task is activated.

Desire's priority levels are dynamic, meaning that the desire can move from one level to another. This allows to cover cases where the urgency of a given desire, that is, the urgency with which the system has to try the satisfaction of it, changes. Thus, the agent is always focused on the most urgent desires, that is, the more important ones in the current situation. If a desire is important but not urgent it will be in a lower priority level and will receive its opportunity when the urgent ones are satisfied.

Intentions and the Behavior System

The intentions module contains the planning algorithm and the current plan. As said before, the BDI architecture nicely separates the different modules of which is composed. Thus, the planning algorithm is completely interchangeable. In our experiments, several algorithms such as a rule system or a very simple reactive algorithm, have been used. The connection of emotions with intentions through the goals allows to keep this division. In the future, a the dynamic reconfigurable behavior network commented before will be introduced.

BDIE's Update Cycle

Finally, the BDIE architecture's update cycle is presented. The following steps are done in each cycle:

1. Sensors are updated getting from them the current input stimuli.
2. Input stimuli, in the form of beliefs, are introduced in the first level of the belief container.
3. First level evaluators are invoked to check for special situations that may fire primary emotions.
4. If a primary emotion is activated, the emotional space is updated, its corresponding desire is unsatisfied and the execution control passes to the the planning algorithm.
5. If no primary emotion is activated then, the second level evaluators are invoked. They compute new second level beliefs given the first level ones.
6. The emotional space is updated and the current emotion computed. In case it is a primary emotion, the point four is executed.
7. The control is transfered to the planning algorithm.
8. The next behavior is executed.

Experiments

The architecture has been implemented using the Java programming language. It is an object-oriented framework that due to the facilities given by the language is heavily based on dynamic load and unload of modules and components. For example, it is possible to add a new module to the architecture without having to recompile it. The designer of an

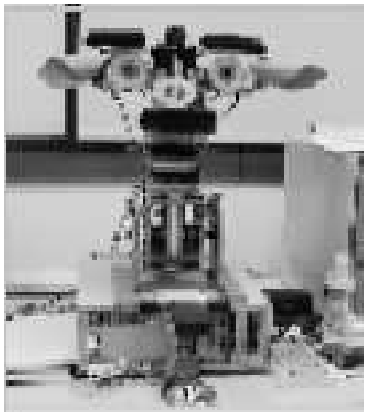

Figure 5: The Casimiro robot head

agent creates the classes extending others classes or implementing interfaces from the architecture. At load time, all the modules will be loaded and configured.

To test the basic components that the architecture already have, we implemented a simple agent in a robotic head developed in our group. The tests have allowed us to identify the design flaws and strengths. From them, changes have been proposed and will be introduced in next versions of BDIE.

Casimiro (Déniz & others 2002) is a humanoid robotic head with vision and audio sensor capabilities and several mechanical degrees of freedom. For the experiments we use one of the cameras as a sensor and the head's motors as effectors. Figure 5 is a picture of the robotic head.

The goal of the experiment was to produce facial expressions for certain colors and the activation of primary emotions given the image's luminance value. For this task, the robot's software provided us with four measures of the current image: the luminance level and the three basic color components red, green and blue.

In the emotional system there is one evaluator in the first and second levels, that is, there is one process in charge of the affective appraisal and another one in charge of the cognitive evaluation. There are five emotions: happiness, sadness, surprise, fear and anger. Surprise and fear are first level emotions while the others are second level ones. Figure 6 shows the emotions distribution in the emotional space.

The desires system counts only with one desire, representing the survival goal. It is in the highest priority level and is always satisfied unless a primary emotion is activated. In that case, the desire is unsatisfied. The example tries to focus on the emotional system and homeostatic variables have the same influence than goals on it, so they are not included.

Finally, the intentions system is based on a very simple algorithm that shows a facial expression given the current emotion.

Each update cycle, the four beliefs, each one with one of the sensed values, are created. In the next step the first level evaluator examine the belief corresponding to the luminance and determines if a first level emotion should be activated.

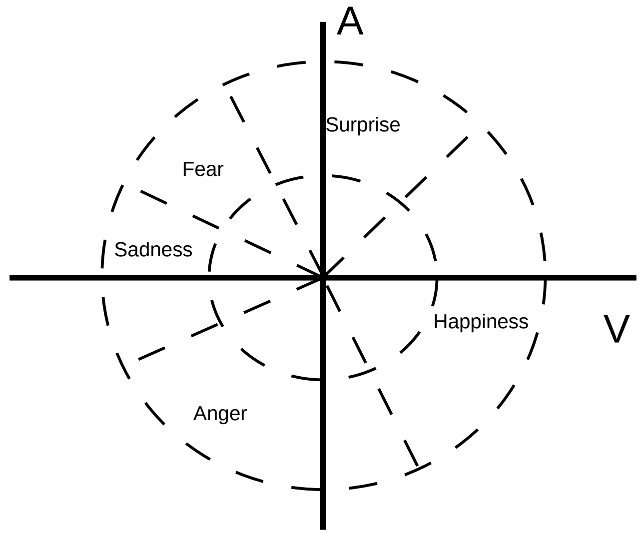

Figure 6: An emotional space with five emotions

If that is the case, the second level evaluator is not invoked and the corresponding first level emotion's desire is unsatisfied. As an example of the effect of emotions in sensors, if a first level emotion is activated the sensors stop working until the emotion becomes inactive. Then, the planning algorithm computes the action it should execute to satisfy the desire, but in this case it does nothing. When the selected behavior must be executed, the facial expression is updated given the current emotion.

If no first level emotion is activated, then the second level evaluator is invoked. It uses the three color beliefs to compute a position in the [r,g,b] space. Previously, a table with points in this space for the colors that will be shown to the robot has been created (see Table 1). Around each of these key points, a cube is defined. If the current point is inside any of these cubes, then the robot thinks that the corresponding color is being shown. Additionally, each color has an associated valence and arousal. The corresponding emotional change, output of the current three considered input stimuli, will be these arousal and valence multiplied by a factor that decreases with the distance to the key point. Thus if the current color matches the key point, the valence and arousal will be the one associated with the color. On the other hand, if the color is near a cube's border, the arousal and valence will be close to zero. Finally, a new second level belief is created with this computed arousal and valence.

After the evaluation process, the current emotional point is updated given the evaluation of the belief recently created. As a result, the point navigates through the emotional space. Note that a first level emotion may also be activated.

Finally, the planning algorithm is invoked and the execution of the current algorithm requested. As said before, in these experiments this sets the current facial expression according to the emotional state.

Table 1: The color space

	Red	Green	Blue	Valence	Arousal
Red	124	27	28	-10	+10
Green	5	105	85	0	0
Blue	11	80	125	+5	-5
Yellow	135	140	0	+5	+5
Orange	165	65	20	+7	+7
Black	7	5	5	-10	-10

Conclusions and Future Work

The initial revision of the BDIE architecture has been presented. This work has allowed us to research in the field of emotions and to produce a first model of an architecture that will evolve as the research advances. The architecture is illustrated with a simple example which allowed us to catch several design flaws and strengths and, to observe the properties of the selected valence-arousal space. Thus, several improvements, design changes and alternative representations have emerged. We will put effort in the analysis and implementation of those proposals. The experiments have proved that the architecture is valid in an application with a real robotic head but should be improved if more complex behavior is desired.

The main strength of the architecture is the fact that the emotional state affects to all the other components in the architecture. This reflects the fact that in minds, the emotions are deeply related to other systems. The intention is not to create perfect agents but to create *believable* ones and this may imply to take wrong decisions.

The chosen emotional representation is not enough for complex emotional phenomena. For example, visceral factors (Loewenstein 1996) could somehow be modeled with the current version of the architecture but there should be a mechanism that would allow an easy implementation of this phenomena as they are so common. Such a simple representation was implemented to start from the ground and earn some experience. In the future, a more advanced one will substitute the current one.

As said before, the architecture is based in a dynamic modules and components loading system. This will allow us to easily add new components to the architecture. Some of these planned modules are:

- A long term memory module. Currently, the only memory the agent has are the beliefs and they are short term. It is clear that to produce complex emotional reactions there is a need for long term memory to remember past actions and events.
- The action selection algorithm, as pointed out, is one of our main interests. As noted, we are currently developing a reconfigurable behavior network whose parameters, structure and runtime algorithm will be dynamically established based on the current emotional state.
- Desires are planned to support hierarchical trees. The designer will then be allowed to define complex desires that depend on another ones.
- An attention filter is necessary to filter out not related sensory information. This module would be affected by the emotional state. Experiments have shown that in high arousal situations, attention tend to focus in few elements than in normal situations. Thus, the input stimuli varies.
- A habituation module is also necessary. The agent should not always react the same way to the same stimuli. If some stimuli are repeated, the agent should get habituated to them and does not react in the same way.
- Learning mechanisms seem to be pretty necessary also. They will be related to the emotional module, with the somatic marker hypothesis (Damasio 1994) as a key point.

The goal of all this modules is to allow the designer to create agents with higher behavior complexity in an attempt to make them more believable.

Acknowledgments

This work has been partially funded by a personal grant from the University of Las Palmas de Gran Canaria to the first author. Additional funds come from research project UNI2002/16 of the same university and from research project PI2003/165 of the Canary Islands Government.

The authors would also like to thanks the reviewer of the initially submitted long abstract due to the insightful comments and very useful suggestions.

References

Alvarado, N.; Adams, S. S.; and Burbeck, S. 2002. The role of emotion in an architecture of mind.

Breazeal, C. 2002. *Designing Sociable Robots*. MIT Press.

Canamero, D. 2000. Designing emotions for activity selection. Technical report, University of Aarhus, Denmark.

Castelfranchi, C. 2000. *Affective Appraisal vs Cognitive Evaluation in Social Emotions and Interactions*. Springer. 76–106.

Damasio, A. R. 1994. *Descartes' Error: Emotion, Reason, and the Human Brain*. Putnam Pub Group.

Déniz, O., et al. 2002. Casimiro: A robot head for human-computer interaction. In *Proceedings of 11th IEEE International Workshop on Robot and Human Interactive Communication*.

Ekman, P. 1999. *Handbook of Cognition and Emotion*. Sussex, U.K.: John Wiley and Sons, Ltd. chapter 3, 45–60.

Gratch, J. 1999. Why you should buy an emotional planner. In *Proceedings of the Agents'99 Workshop on Emotion-based Agent Architectures (EBAA'99)*.

Kinny, D.; Rao, A. S.; and Georgeff, M. P. 1996. A methodology and modeling technique for systems of BDI agents. In *Agents Breaking Away: Proceedings of the Seventh European Workshop on Modeling Autonomous Agents in a Multi-Agent World, MAAMAW'96*.

Loewenstein, G. 1996. Out of control: Visceral influences on behavior. *Organizational Behavior and Human Desicion Making* 65(3):272–292.

Maes, P. 1989. How to do the right thing. *Connection Science Journal* 1(3):291–323.

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

Rao, A. S., and Georgeff, M. P. 1995. BDI agents - from theory to practice. In *Proceedings of the First Intl. Conference on Multiagent Systems*.

Sloman, A. 2000. How many separately evolved emotional beasties live within us? Technical report, The University of Birmingham. Revised version of Invited Talk: at workshop on Emotions in Humans and Artifacts Vienna, August 1999.

Velásquez, J. D. 1997. Modeling emotions and other motivations in synthetic agents. In *Proceedings of AAAI-97*. 545 Technology Square, NE43-812, Cambridge, MA 02139: MIT AI Lab.

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

Velásquez, J. D. 1998b. When robots weep: Emotional memories and decision-making. In *Proceedings of AAAI-98*, 70–75. 545 Technology Square, NE43-935, Cambridge, MA 02139: MIT AI Lab.