PROVIDING DELIBERATION TO EMOTIONAL AGENTS

Daniel Pérez-Pinillos, Susana Fernández and Daniel Borrajo
Departamento de Informática, Universidad Carlos III de Madrid, Avda. de la Universidad 30, Leganés, Madrid, Spain

Keywords: Planning, Agent, Emotion, Personality.

Abstract: Modelling real persons or virtual agents motivations, personality and emotions is a key feature of many user-oriented applications. Most of the previous work has defined rich cognitive models of motivations, personality and emotions, but have relied on some kind of reactive scheme of problem solving and execution. Instead, this work proposes a deliberative emotional model for virtual agents based in their basic needs, preferences and personality traits. More specifically, we advocate the integration of these comprehensive agents models within deliberative automated planning techniques, so that plans to be executed by agents to achieve their goals already incorporate reasoning at the emotional level.

1 INTRODUCTION

The work on reasoning about emotions is becoming increasingly relevant, specially in contexts such as assistive technology, user interfaces, or virtual agents (Fellous and Arbib, 2005; Bates, 1994). Emotions have been studied in psychology, neurology and physiology from a wide variety of points of view and each field focuses the attention on different aspects (Damasio, 1994). It seems there is some agreement to consider emotion as an inborn and subjective reaction to the environment, with an adaptive function, and accompanied of several organic, physiological and endocrine changes (Frijda, 1988). Another point of agreement is that emotions are an outstanding factor in humans, because they modify and adapt their usual behavior. In the development of systems that interact with persons, as human behavior simulators, emotions can not be ignored, because, on one hand, they may help on this interaction and, on the other hand, they constitute a decisive part of human reasoning and behavior. This is specially true when reasoning about sequential decision-making, as in medium-long term planning, where the sequence of decisions can be influenced by the emotional state of agents.

During the last years, several emotion-oriented systems have been developed, that normally follow Frijda's theory about emotions (Frijda, 1995). This theory is based on the hypothesis that emotions are functional most of the time. Thus, the use of emotions in artificial systems is needed to achieve individual objectives, and the design of social agents with emotional characteristics is justified. Emotions also cover the interaction of the individual with the environment. For instance, individuals try to move away objects that put in danger their survival, while they approach objects that cater for their needs (Breazeal, 2003).

Emotions are also very related to characteristics of human personality. In contemporary psychology, there are five factors or dimensions of personality, called the Big Five factors (Goldberg, 1993), which have been scientifically defined to describe human personality at the highest level of organization. The Big Five traits are also referred to as a purely descriptive model of personality called the Five Factor Model (Costa and McCrae, 1992; McCrae, 1992). The Big Five factors are: openness to experience, conscientiousness, extraversion, agreeableness and neuroticism (opposite to emotional stability).

Examples of previous work on computational models of emotions is the work of Cañamero (Cañamero, 1997; Cañamero, 2003) that proposes a homeostatic approach to the motivations model. She creates a self-regulatory system, very close to natural homeostasis, that connects each motivation to a physiological variable, which is controlled within a given range. When the value of that variable differs from the ideal one, an error signal proportional to the deviation, called drive, is sent, and activates some control mechanism that adjusts the value in the right direction. There are other architectures based on drives, as the Dorner's

Pérez-Pinillos D., Fernández S. and Borrajo D..
PROVIDING DELIBERATION TO EMOTIONAL AGENTS.
DOI: 10.5220/0003167200980105
In *Proceedings of the 3rd International Conference on Agents and Artificial Intelligence* (ICAART-2011), pages 98-105
ISBN: 978-989-8425-40-9

PSI architecture used by Bach and Vuine (Bach and Vuine, 2003) and also by Lim (Lim et al., 2005), that offer a set of drives of different type, as certainty, competence or affiliation.

Most of these works on emotional agents are based on reactive behaviors as the work by Cañamero. When a drive is detected, it triggers a reactive component that tries to compensate its deviation, taking into account only the following one or two actions. Thus, there is no inference being done on medium-long term goals and the influence of emotions on how to achieve those goals. Regarding deliberative models, there are some works on emotions based on planning, but mainly oriented to storytelling like emergent narrative in FEARNOT! (Aylett et al., 2005) and the interactive storytelling of Madame Bovary on the Holodeck (Cavazza et al., 2007). The work of Gratch and coauthors (Gratch, 1999; Gratch et al., 2002) shows a relevant application of emotional models to different research areas in artificial intelligence and autonomous agents design, endowing them with an ability to think and engage in socio-emotional interactions with human users.

In the present work, a model of long term reasoning based on emotions and factors of personality has been designed. It follows some ideas introduced in (Avradinis et al., 2003) using concepts that already appeared in Cañamero's works, like motivations and the use of drives to represent basic needs. Our model uses automated planning for providing long term deliberation on effects of actions taking into account not only the agents goals, but also the impact of those actions in the emotional state of the agent. Other models, as Rizzo's works (Rizzo et al., 1997), combine the use of emotions and personality to assign preferences to the goals of a planning domain model, but the changes in the emotional state happen in another module. Thus, they are not really used in the reasoning process. A similar integration of a deliberative and a reactive model is the one in (Blythe and Reilly, 1993) where the emotions reasoning is performed again by the reactive component.

We have defined a planning domain model that constitutes the reasoning core of a client in a virtual and multi-agent world (Fernández et al., 2008). It is a client/server game oriented towards the intensive use of Artificial Intelligence controlled Bots, and it was designed as a test environment of several Artificial Intelligence techniques. The game borrows the idea from the popular video game THE SIMS. Each agent controls a character that has autonomy, with its own drives, goals, and strategies for satisfying those goals. In this implementation, we introduce the concept of how an agent prefers some actions and objects depending on its preferences, its personality traits and its emotional state, and the influence of those actions on long term achievement of goals. Thus, agents solved problems improving the quality of the solution, achieving better emotional states.

The remainder of the paper describes the model design, the description of the domain that implements the model, the empirical results that validate the model and the conclusions derived from the work, together with future research lines.

2 MODEL DESIGN

Our aim in this work is to include emotions and human personality traits in a deliberative system, that uses automated planning in order to obtain more realistic and complex behavior of agents. These behaviors are necessary to implement a wide variety of applications such as agents that help users to change their way of life, systems related with marketing and advertising, educational programs, systems that play video games or automatically generate text. The goal is to show that the use of emotional features, with the establishment of preferences about certain actions and objects in its environment, improves the performance of a deliberative agent by generating better plans.

In the virtual world, an agent tries to cater for its needs, its motivations, through specific actions and interacting with different objects. Five basic needs have been identified for the agent, which are easily identifiable in human beings: hunger, thirst, tiredness, boredom and dirtiness. Along with the first three, widely used in many systems, we have added dirtiness and boredom, which are more domain-specific to add a wider variety of actions and get richer behaviors. These basic needs increase over time, so their values increase as time goes by. Thus, the agent always needs to carry out actions to maintain its basic needs values within reasonable limits.

To cater for each of these basic needs, the agent must perform actions. For example, it can drink to satisfy its thirst or sleep to recover from fatigue. There are different actions to cater for the same need, and the agent prefers some actions over others. Thus, the agent may choose to read a book or play a game to reduce boredom. Besides, the effects of those actions can be different depending on its emotional state. It will receive more benefit from applying more active actions when its emotional state is more aroused and more passive or relaxed actions when it is calm.

To carry out each of these actions, the agent needs to use objects of specific types. Thus, it will need food to eat, a ball to play or a book to read. There are different objects of each type in its environment and the agent has preferences over them. When an agent executes an action with an object, its emotional state is modified depending on the agent personality, and preferences and activations for this object.

We have chosen to implement a model widely-accepted in psychology that represents the emotional state of an agent as a two-dimensional space of two qualities: valence and arousal (Duffy, 1941). Valence ranges from highly positive to highly negative, whereas arousal ranges from calming or soothing to exciting or agitating. The first one is a measure of the pleasantness or hedonic value, and the second one represents the bodily activation. Other models use a set of independent emotions, which requires defining a group of basic emotions. However, not all combinations of values for these emotions are a valid emotional state (e.g. the combination of maximum values in the emotions of joy and anger is not a realistic emotional state). In general, the valence and arousal model can be shown to be equivalent to the explicit representation of the usual set of emotions of other computational cognitive simulations, though it requires a simpler representation and reasoning. For instance, an emotion such as happiness can be represented as high valence and high arousal. Both models are recognized and defended by experts in psychology, but we prefer the second alternative because it makes processing easier and prevent invalid states. In our model, the valence and the arousal are modified by the execution of actions, so both values are modified when an agent executes an action with an object, depending on the agent preference and activation for this object, the personality traits and the emotional state. Our goal is that the agent generates plans to satisfy its needs and to achieve the most positive value of valence.

3 DOMAIN DESCRIPTION

In order to use domain-independent planning techniques, we have to define a domain model described in the standard language PDDL (Fox and Long, 2003). This domain should contain all the actions that the agent can perform in order to achieve the goals. Automated planning can be described as a search for a solution on a problem space where, the states are represented using a set of predicates, functions and types, and the actions are described with a set of preconditions and effects that model the state transitions. An action is applicable only if all its preconditions hold in the current state and executing the action changes the current state by adding and deleting the action effects. A problem is specified as an initial state (true literals in the starting state) and a set of goals. Also, an optimization metric (as in our case valence and/or arousal) can be defined. Our domain has been designed based on the previous concepts of drive, emotion, preference, activation and personality traits to represent each agent of the virtual world. Now, we will define the different concepts composing the model, in automated planning terms.

3.1 Drives

As already said, we use five drives: hunger, thirst, tiredness, dirtiness and boredom. Drives are represented in the domain through functions. The ideal value for all drives is established at zero. So, when a drive has a value of zero, its need is totally satisfied. Any other value means the intensity of the need and the distance to the ideal value. The value of each drive is increased as time goes by to represent the need rise. To reduce it, the agent has to carry out some action. For instance, the agent must `eat` to reduce the drive `hunger`. Given that the drives increase with time, every time an action is executed, one or more drives will be decreased, but the rest will be increased. Thus, the planning task becomes hard if we want all drives to be fulfilled (below a given threshold).

3.2 Objects

Objects describe the different elements of the virtual world. Objects may be of two kinds: resources (or physical objects) and rooms. Resources represent objects needed to carry out the actions to cater for needs; for instance, food, balls, books, etc. Rooms describe physical spaces, where the agents may move and where resources are placed. Both kinds of objects are represented as planning types and several instances of them will be present in each problem. Also resources may be of two kinds: fungible resources and non-fungible resources.

3.3 Personality Traits

Personality traits describe the agents personality and are based on the Big Five factors model (openness to experience, conscientiousness, extraversion, agreeableness and neuroticism). Openness to experience involves active imagination, aesthetic sensitivity, preference for variety and intellectual curiosity. Openness is modeled as a higher preference for new experiences, i.e., an agent with high openness (open-minded) tends to use and prefer new objects to known objects, while an agent with low openness will tend to prefer known objects to new objects. Neuroticism represents the degree of emotional stability of the agent. The bigger the neuroticism is, the smaller the emotional stability is. So, neuroticism is implemented as the variation factor of the emotional state and is represented by the use of a PDDL function. Thus, the emotional state of a neurotic agent will vary more suddenly than a stable one when actions are applied, as described later.

Conscientiousness includes elements such as self-discipline, carefulness, thoroughness, organization, deliberation and need for recognition. We implement conscientiousness as a factor in the decrements of the drives due to action executions, representing how meticulous the agent is in carrying out the action. Thus, an agent with a high value of contientiouness gets a bigger effect when applying actions (a bigger decrease of the involved drive). But, similarly, the other drives will also increase proportionately to the contientiouness value as time passes.

The last two factors, extraversion and agreeableness, are related to social interaction. Thus, they will be used in future versions of the system that include multiple agents and interactions among them.

3.4 Emotional State

The agents emotional state is determined by two components: valence and arousal. Valence represents whether the emotional state of the individual is positive or negative and to which degree. Arousal represents the bodily activation or agitation. We represent them in the domain as PDDL functions. Since we want to obtain plans that maximize the valence, we have to define the planning problems metric accordingly. Even if PDDL allows generic functions to be defined as metrics, most current planners can only deal with metrics that are defined over minimizing an increasingly monotonous function (no action can have an effect that decreases its value), since metrics are considered in PDDL as costs and each action has an associated cost.

In our model, objects used in the actions can cause valence both to increase (when the agent likes the object) or decrease (when it does not like it). Therefore, it is not possible to use the valence directly as the problem metric. Instead, we define an increasingly monotonous function, `v-valence`, that the planner tries to minimize. Its increment depends on two factors. The first one refers to the neuroticism of the agent, which is the reverse of the emotional stability. It measures the impact of any action on the agent emotional state. Thus, the greater the agent neuroticism, the greater the changes in its emotional state. The second factor represents the type of change produced by the preferences of the agent.The greater the preference for the action and for the object, the lower the value of this factor is, and, therefore, the lower the increase in the `v-valence` function (that the planner tries to minimize). Thus, each action increases `v-valence`, with positives values between 0 and a threshold in the following amount:

$$\Delta v = (\frac{n}{n_{max}}) \times (p_{max} - \frac{(p_a + p_o)}{2})$$

where v is the value of `v-valence`, n the agent neuroticism, n_{max} the maximum possible value for neuroticism, p_{max} the maximum possible value for a preference, p_a the agent preference for the executed action and p_o the agent preference for the used object. In case the object is new to the agent, p_o=-1 and we replace p_o for the value of the agent openness.

Thus, the metric used consists on minimizing that value, so we add the following to the planning problems:

`(:metric minimize (v-valence))`.

3.5 Preferences

Preferences describe the agent personal likes for each physical object of its environment. They are represented as PDDL functions of the form:

`(= (preference apple) 5)`

These values are not modified during the planning process and they are between zero, for the detested objects, and ten, for the favourite ones. Preferences can also describe the agent personal likes for each action. They are represented as PDDL functions of the form:

`(= (read-preference) 5)`

Again, these values are not modified during the planning process and they are between zero, for the detested actions, and ten, for the favourite ones. Preferences affect the direction and degree of changes on the value of the valence, produced by the effects of actions.

3.6 Activations

Activations describe the effect over the agent arousal for each physical object of its environment. They are represented as PDDL functions of the form:

`(= (activation apple) 5)`

These values are not modified during the planning process and they are between zero, for the objects that relax, and ten, for the objects that agitate. Activations can also describe the effect over the agent arousal for

```
(:action READ
     :parameters   (?reading-object - reading-object)
     :precondition (and (taken ?reading-object)
                        (not (time)))
     :effect (and (time)
                  (decrease (boredom) (conscientiousness))
                  (when (and (< (boredom) 0))
                    (and (assign (boredom) 0)))
                  (when (and (< (preference ?reading-object) 0))
                    (and
                      (increase (valence) (* (/ (neuroticism) (max-neuroticism)) (- (/ (+ (preference ?reading-object) (read-preference))
                                           (max-preference)) 1)))
                      (increase (v-valence) (* (/ (neuroticism) (max-neuroticism))(- (max-preference) (/ (+ (preference ?reading-object)
                                           (read-preference)) 2))))))
                  (when (and (> (preference ?reading-object) 0))
                    (and
                      (increase (valence) (* (/ (neuroticism) (max-neuroticism)) (- (/ (+ (openness) (read-preference)) (max-preference)) 1)))
                      (increase (v-valence) (* (/ (neuroticism) (max-neuroticism))(- (max-preference) (/ (+ (openness) (read-preference)) 2))))))
                  (increase (arousal) (* (/ (neuroticism) (max-neuroticism)) (- (/ (+ (activation ?reading-object) (read-activation))
                                           (max-activation)) 1)))
                  (increase (v-arousal) (* (/ (neuroticism) (max-neuroticism))(- (max-activation) (/ (+ (activation ?reading-object)
                                           (read-activation)) 2))))))
```

Figure 1: Example of action (`READ`) to cater for the `boredom` need.

```
(:action EAT
     :parameters   (?food - food)
     :precondition (and (taken ?food)
                        (not (time)))
     :effect (and (time)
                  (decrease (hunger) (conscientiousness))
                  (when (and (< (hunger) 0))
                    (and (assign (hunger) 0)))
                  (when (and (< (preference ?food) 0))
                    (and
                      (increase (valence) (* (/ (neuroticism) (max-neuroticism)) (- (/ (+ (preference ?food) (eat-preference))
                                               (max-preference)) 1)))
                      (increase (v-valence) (* (/ (neuroticism) (max-neuroticism))(- (max-preference) (/ (+ (preference ?food)
                                            (eat-preference)) 2))))))
                    (when (and (> (preference ?food) 0))
                        (and
                          (increase (valence) (* (/ (neuroticism) (max-neuroticism)) (- (/ (+ (openness) (eat-preference)) (max-preference)) 1)))
                          (increase (v-valence) (* (/ (neuroticism) (max-neuroticism))(- (max-preference) (/ (+ (openness)
                                                   (eat-preference)) 2))))))
                    (increase (arousal) (* (/ (neuroticism) (max-neuroticism)) (- (/ (+ (activation ?food) (eat-activation)) (max-activation)) 1)))
                    (increase (v-arousal) (* (/ (neuroticism) (max-neuroticism))(- (max-activation) (/ (+ (activation ?food)
                                               (eat-activation)) 2))))))
```

Figure 2: Example of an action (`EAT`) to cater for a need (`hunger`).

each action. They are represented as PDDL functions of the form:

`(= (read-activation) 5)`

Again, these values are not modified during the planning process and they are between zero, for the actions that relax, and ten, for the actions that agitate.

3.7 Actions

Actions defined in the domain describe activities that the agent may carry out. There are five types of actions:

- `Actions to cater for its needs`: Each one of these actions needs one object of a specific type to decrease in one unit its corresponding drive value. In this group of actions, we have defined: `eat`, `drink`, `sleep`, `bath`, `shower`, `play`, `read`, `watch` and `listen`. Some of these actions require that the agent has taken the object used, like `eat`, `drink` or `read`. Others, however, only require that the object is located in the same room of the agent, like `bath` or `sleep`. In addition, some actions such as `eat` and `drink` decrease the available amount of the object used.

  In Figures 1 and 2, we show two examples of this type of action. We can see that the changes in the agent emotional state (valence and arousal) depend on the preferences, activations and personality traits, so we have an integrated model of these concepts, that can affect how actions are combined in order to solve the agents problems.

- `TAKE and LEAVE actions`: the agent uses them to take and leave objects required to perform some actions, like `eat` or `drink`.
- `BUY action`: the agent uses it to purchase new resources. Agents must be in a shop and the resource must be available to be bought.
- `GO action`: allows the agents to move as Figure 3 shows.
- `TIME action`: It is a fictitious action (Figure 4) that represents the influence of the course of time over the value of the drives. Its execution produces an increase on all drives, so that it simulates the passing of time. The increment depends

```
(:action GO
    :parameters   (?place-from - place ?place-to - place)
    :precondition (and (in ?place-from)(not (time)))
    :effect (and (time)
                 (increase (valence) (* (/ (neuroticism) (max-neuroticism)) (- (/ (* (go-preference) 2) (max-preference)) 1)))
                 (increase (arousal) (* (/ (neuroticism) (max-neuroticism)) (- (/ (* (go-activation) 2) (max-activation)) 1)))
                 (increase (v-valence) (* (/ (neuroticism) (max-neuroticism)) (- (max-preference) (go-preference))))
                 (increase (v-arousal) (* (/ (neuroticism) (max-neuroticism)) (- (max-activation) (go-activation))))
                 (not (in ?place-from))
                 (in ?place-to)))
```

Figure 3: GO action.

on the agent's conscientiousness. We also force the planner to be executed after every other action application (through the `time` predicate).

```
(:action TIME
     :parameters   ()
     :precondition (and (time))
     :effect       (and
                    (increase (boredom) (* 0.1 (conscientiousness)))
                    (assign (boredom) (min (max-drive) (boredom)))
                    (increase (dirtiness) (* 0.1 (conscientiousness)))
                    (assign (dirtiness) (min (max-drive) (dirtiness)))
                    (increase (hunger) (* 0.1 (conscientiousness)))
                    (assign (hunger) (min (max-drive) (hunger)))
                    (increase (thirst) (* 0.1 (conscientiousness)))
                    (assign (thirst) (min (max-drive) (thirst)))
                    (increase (tiredness) (* 0.1 (conscientiousness)))
                    (assign (tiredness) (min (max-drive) (tiredness)))
                    (not (time))))
```

Figure 4: TIME action.

All actions (except for TIME) modify (in their effects) the emotional state that depend on the agent preferences, activations and personality traits. Along with the metric of the problem, this allows us to model the agents behaviour. So, there are no hard constraints on our model. All agents can perform all actions, but they prefer (soft constraints) the ones that better suit their preferences, personality and current emotional state.

3.8 Goals

The agent motivation is to satisfy its basic needs, so goals consist of a set of drives values that the agent has to achieve. As an example, goals may consist of the achievement of need values that are under a given threshold. They could be very easily combined with other kinds of standard planning goals, creating other kinds of domains. For instance, we could define strategy games where agents should accomplish some tasks, taking into account also their needs.

4 EXPERIMENTS

We report here the results obtained with the proposed model comparing its performance to a reactive model. In the case of the deliberative model, we have used an A* search technique with the well-known domain-independent heuristic of FF (Hoffmann, 2001). This heuristic is not admissible, but even if it does not ensure optimality, it is good enough for our current experimentation. In the case of the reactive model, we have used a function to choose the best action at each step (to cover the drive with the higher value, i.e. the worse drive). These search techniques have been implemented in an FF-like planner, SAYPHI (De la Rosa et al., 2007).

4.1 Experimental Setup

We have defined several kinds of problems for this domain. In each problem, we have established a specific initial need in one of the drives, which are called dominant drives. Each of these dominant drives will have a initial value higher than the rest of drives. Also, we have defined a problem where all five drives are dominant drives. The goal is to fulfill all the agent needs, so we have defined it as having a value below a threshold for all drives. Furthermore, for each action, the agent has three objects to choose from, with varying degrees of preference: preferred, indifferent and hated, and a new object (the agents do not have an "a priori" preference for this object) for testing openness.

The experiments were performed with four different personality models: (1) a standard personality (average values in all traits), (2) a neurotic personality (high value of neuroticism and average values for the rest), (3) an open-minded personality (high value of openness and average values for the rest) and (4) a meticulous personality (high value of conscientiousness and average values for the rest).

4.2 Results

Figures 5 to 8 show the end value of the (`valence`) for each problem. In all cases, the value obtained by the proposed deliberative model is significantly better than the reactive one. This is due to a better employment of the `buy` action and the reduction on `go` actions of the deliberative model. The reactive model always tries to satisfy the need associated to the most dominant drive at each time. So, for instance, if reducing the current dominant drive requires drinking, and there is no drink in the current agent room, then the agent will move to another room where the drinking action can be accomplished. However, the deliberative model reasons on a medium-long term, so if the need in another drive, not being the dominant one, can be satisfied in the current room, the plan will prefer to reduce it now, even if the dominant drive increases a bit. Most previous work on emotional agents would mimic the reactive model, while our model is able to take into account future recompenses in an integrated way with other agents goals. We also see that if the personality tends to be more neurotic, then the deliberative model is even better than the reactive one, since actions effects are increased, and drives increase more acutely.

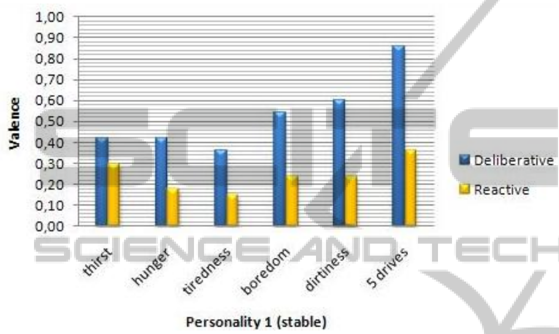

Figure 5: Quality of the plans for the stable agent.

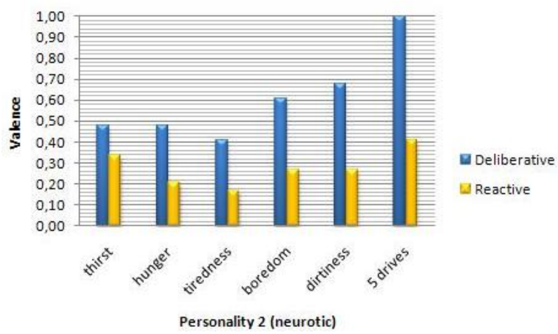

Figure 6: Quality of the plans for the neurotic agent.

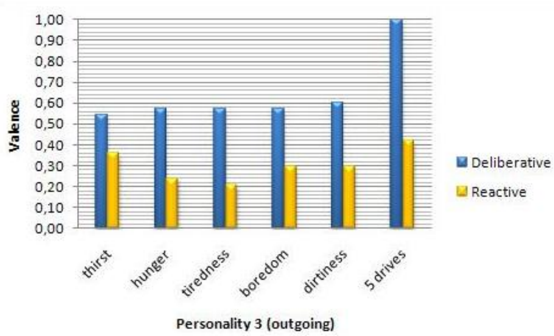

Figure 7: Quality of the plans for the open-minded agent.

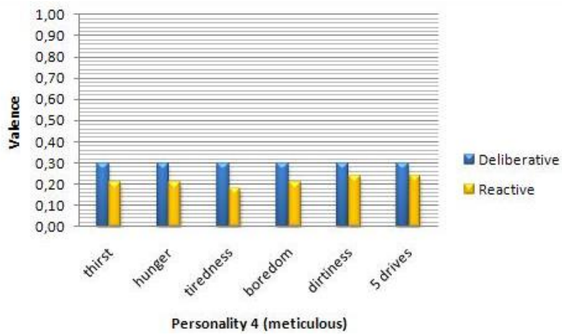

Figure 8: Quality of the plans for the meticulous agent.

5 CONCLUSIONS AND FUTURE WORK

This work proposes a model of long term reasoning integrating emotions, drives, preferences and personality traits in autonomous agents, based on AI planning. The emotional state is modeled as two functions: valence and arousal. This two-dimensional model has been chosen because it is simpler and offers the same representation capabilities as the rest of emotional models. Anyhow, it is not difficult now to integrate any other emotional model. Thus, actions produce variations in the valence depending on the agent personality and agent preferences. The goal is to generate plans that maximize the valence, while satisfying the agent needs or drives. Given that current planners only deal with monotonous functions as metric functions, we converted the non-monotonous valence into a monotonous one, `v-valence`. The results of the experiments show that the quality of the solutions (measured as the value of the `valence`) improves when the deliberative model is used compared to the reactive one. Thus, the increase in the quality of the solutions implies a more realistic behavior of the agent.

The proposed model is the first step in the development of a richer and more complex architecture. In the next future, we would like to include new actions in the domain, especially those related to the processes of social interaction, by including some component that reasons about multi-agent interaction and collaboration. Another future work is to model the idea of well-being, which will focus the agent to keep all its needs below a certain level along time. The physiological well-being of the agent will influence its emotional state altering the value of valence. This idea is very related to the idea of continuous planning to control the behaviour of virtual agents (Avradinis et al., 2003).

REFERENCES

Avradinis, N., Aylett, R., and Panayiotopoulos, T. (2003). Using Motivation-Driven Continuous Planning to Control the Behaviour of Virtual Agents. In *International Conference on Virtual Storytelling*, pages 159–162.

Aylett, R. S., Louchart, S., Dias, J., Paiva, A., and Vala, M. (2005). Fearnot!: an experiment in emergent narrative. *Intelligent Virtual Agents*, pages 305–316.

Bach, J. and Vuine, R. (2003). The AEP Toolkit for Agent Design and Simulation. In *MATES 2003, LNAI 2831*, pages 38–49. Springer.

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

Blythe, J. and Reilly, W. S. (1993). Integrating Reactive and Deliberative Planning for Agents. Technical report, Carnegie Mellon University, Pittsburgh, PA, USA.

Breazeal, C. (2003). *Biological Inspired Intelligent Robots.* SPIE Press.

Cañamero, D. (1997). Modeling motivations and emotions as a basis for intelligent behavior. In *First International Symposium on Autonomous Agents (Agents'97)*, pages 148–155. New York, NY: The ACM Press.

Cañamero, D. (2003). *Designing Emotions for Activity Selection in Autonomous Agents*. MIT Press.

Cavazza, M., Lugrin, J., Pizzi, D., and Charles, F. (2007). Madame bovary on the holodeck: Immersive interactive storytelling. In *Proceedings of the ACM Multimedia 2007*. Augsburg, Germany: The ACM Press.

Costa, P. and McCrae, R. (1992). *Revised NEO Personality Inventory (NEO-PI-R) and NEO Five-Factor Inventory (NEO-FFI) manual.* Odessa, FL: Psychological Assessment Resources.

Damasio, A. R. (1994). *Descartes' Error.* New York: Gosset; Putnam Press.

De la Rosa, T., Olaya, A. G., and Borrajo, D. (2007). Using Cases Utility for Heuristic Planning Improvement. In Weber, R. and Richter, M., editors, *Case-Based Reasoning Research and Development: Proceedings of the 7th International Conference on Case-Based Reasoning*, pages 137–148, Belfast, Northern Ireland, UK. Springer Verlag.

Duffy, E. (1941). An explanation of emotional phenomena without the use of the concept of emotion. *Journal of General Psychology*, 25:283–293.

Fellous, J.-M. and Arbib, M. A. (2005). *Who needs emotions? : the brain meets the robot*, volume Series in. Oxford University Press, Oxford ; New York. edited by Jean-Marc Fellous and Michael A. Arbibill. ; 25 cm.

Fernández, S., Asensio, J., Jiménez, M., and Borrajo, D. (2008). A Social and Emotional Model for Obtaining Believable Emergent Behavior. In Traverso, P. and Pistore, M., editors, *Proceedings of The 13th International Conference on Artificial Intelligence: Methodology, Systems, and Applications (AIMSA-08)*, volume 5253/2008 of *Lecture Notes in Computer Science*, pages 395–399, Varna, Bulgaria. Springer Verlag.

Fox, M. and Long, D. (2003). Pddl2.1: An extension to pddl for expressing temporal planning domains. *Journal of Artificial Intelligence Research*, 20:61–124.

Frijda, N. H. (1988). The laws of emotion. *American Psychologist*, 43:349–358.

Frijda, N. H. (1995). Emotions in robots. In *Comparative Approaches to Cognitive Science.*, pages 501–516. In H.L. Roitblat and J.-A. Meyer (Eds.), Cambridge, MA: The MIT Press.

Goldberg, L. (1993). The structure of phenotypic personality traits. *American Psychologist.*, 48:26–34.

Gratch, J. (1999). Why you should buy an emotional planner. In *Proceedings of the Agents 1999 Workshop on Emotion-based Agent Architectures (EBAA 1999) and ISI Research Report*, pages 99–465.

Gratch, J., Rickel, J., André, E., Badler, N., Cassell, J., and Petajan, E. (2002). Creating interactive virtual humans: Some assembly required. *IEEE INTELLIGENT SYSTEMS*, 17:54–63.

Hoffmann, J. (2001). FF: The fast-forward planning system. *AI magazine*, 22:57–62.

Lim, M. Y., Aylett, R., and Jones, C. M. (2005). Emergent affective and personality model. In *In The 5th International Working Conference on Intelligent Virtual Agents*, pages 371–380. Kos, Greece. September 2005. Springer.

McCrae, R.R.and John, O. (1992). An introduction to the five-factor model and its applications. *Journal of Personality*, 2:175–215.

Rizzo, P., Veloso, M., Miceli, M., and Cesta, A. (1997). Personality-driven social behaviors in believable agents. In *In Proceedings of the AAAI Fall Symposium on Socially Intelligent Agents*, pages 109–114. AAAI Press.