AFFECTIVE AGENTS PERCEIVING EACH OTHER'S ACTIONS

Matthijs Pontier
VU University, Amsterdam
Center for Advanced Media Research Amsterdam
De Boelelaan 1083, 1081HV, The Netherlands
mpontier@few.vu.nl
www.few.vu.nl/~mpontier

Ghazanfar Farooq Siddiqui
VU University, Amsterdam
Center for Advanced Media Research Amsterdam
De Boelelaan 1083, 1081HV, The Netherlands
ghazanfa@few.vu.nl
www.few.vu.nl/~ghazanfa

KEYWORDS

Agent-Based Modelling, Affective Decision-Making, Perception

ABSTRACT

In this paper, an extension of a formalized model of affective decision making is presented, based on the informally described I-PEFiC model. This extension manages that the actions agents undertake have an effect on other agents. The agents change their perceptions and beliefs about other agents if actions are taken. Further, the anger level of the agents is simulated. Simulation experiments show that the actions of agents can change the beliefs and the perceptions of another agent so much that the other agent changes its mind and chooses to perform another action than it was currently doing.

INTRODUCTION

From the last decade, a lot of research has been dedicated to developing Intelligent Virtual Agents (IVAs) with more realistic graphical representations. However, these agents often do not show very realistic human-like emotional behavior. For example, many IVAs can show emotions by the means of facial expressions or the tone in their voice, but most of them still struggle to show the right emotions at the right moment (e.g., emotion regulation (Marsella and Gratch 2003) and moods (Beck 1987)). Let alone actually understanding and reacting empathically to the emotional state of other agents, or human users. Previous research has shown that closely mimicking humans is important for an agent to increase human involvement in a virtual environment (e.g., Van Vugt et al. 2006).
From Frijda's point of view, emotions are goal driven (Frijda 1986). The emotional system examines the surroundings for related stimuli that are either beneficial or harmful for the concerns, motives, and goals of the individual (Frijda 1986, p. 494, p. 463). According to broaden-and-build theory, positive emotions are vehicles for individual growth and social connections: by building people's personal and social resources, positive emotions transform people for the better, giving them better lives in the future (Fredrickson 2001, p. 224). Previous research also showed that human beings usually make unconscious rather than conscious decisions (Bargh and Chartrand 1999).
We created agents for imitating human behavior, who can recognize each other as a personal friend as well as means to an end. The Interactive model of Perceiving and Experiencing Fictional Characters (I-PEFiC), which is based on the theory of Frijda (Frijda 1986), was taken as a foundation with regard to recognizing each other as a personal friend (Van Vugt 2008). Within this model an agent A can compute the trade-off between how involved it is with another agent (e. g., Agent B is good) and what keeps the agent at a distance (e. g., Agent B is bad) (Bosse et al. 2008). This involvement-distance trade-off is the outcome of assessing the features of an agent on several dimensions. Use intentions are calculated additionally that prompt the agent to carry out actions towards another agent.
These actions are based on goals, which play a role in the judgment formation of the agent about the other agent, but also more affective influences are taken into account. The previous decision making model (Hoorn et al. 2008) only describes the affective decision making process itself. Simulation experiments with this model showed that in situations where this can be considered human-like, agents make affective decisions rather than decisions that would be the best rational decision (i. e., the decision option with the highest expected utility). However, it did not explain the effects of the performed actions on the agents' perceptions, beliefs and levels of anger. Neither did it explain how these changed beliefs and perceptions influenced the following decision making processes.
In this paper we improved the affective decision making model (Hoorn et al. 2008), so that the agents update their beliefs and perceptions of the ethics and affordances of other agents when actions are being performed. Further, the effects of these actions on the emotions of the agents are simulated. These changes on their turn influence the decision making process in the agents. This enables the agents to change their mind and decide to change the action they want to perform. The simulation experiments described in this paper will show what kind of effects the actions have on the agents' beliefs, perceptions, and emotions and how this affects their decision making.

Proceedings 23rd European Conference on Modelling and Simulation ©ECMS Javier Otamendi, Andrzej Bargiela, José Luis Montes, Luis Miguel Doncel Pedrera (Editors)
ISBN: 978-0-9553018-8-9 / ISBN: 978-0-9553018-9-6 (CD)

MODELLING APPROACH

Modelling the various aspects involved in affective decision making in an integrated manner poses some challenges. On the one hand, qualitative aspects have to be addressed, such as performing an action. On the other hand, quantitative aspects have to be addressed, such as levels of anger. The modelling approach based on the modelling language LEADSTO (Bosse et al. 2007) fulfils these needs. It integrates qualitative, logical aspects such as used in approaches based on temporal logic (e. g., Barringer et al. 1996) with quantitative, numerical aspects such as used in Dynamical Systems Theory (e. g., (Ashby 1960), (Port and Gelder 1995)).

In LEADSTO, direct temporal dependencies between state properties in two successive states are modelled by executable dynamic properties defined as follows. Let a and b be state properties of the form "conjunction of literals" (where a literal is an atom or the negation of an atom), and e, f, g, h non-negative real numbers. Then in the leads to language $a \twoheadrightarrow_{e, f, g, h} b$, means:

> *If state property* a *holds for a certain time interval with duration* g*, then after some delay (between* e *and* f*) state property* b *will hold for a certain time interval of length* h.

Here, atomic state properties can have a qualitative, logical format, such as an expression desire(d), expressing that desire d occurs, or a quantitative, numerical format such as has_value(x, v) expressing that variable x has value v.

IMPLEMENTATION

I-PEFiC is a model (Figure 1) that is empirically well validated (Van Vugt et al. 2006, Van Vugt et al. 2007, Van Vugt 2008). The I-PEFiC model has three phases: encoding, comparison, and response (Van Vugt et al. 2007).

During encoding, the user appraises an agent's features for their level of ethics (good or bad), aesthetics (beautiful or ugly), and epistemics (realistic or unrealistic). During the encoding, moreover, the user evaluates in how far the agent system has affordances (aids or obstacles), which make the agent useful as a computer tool or not.

In the comparison phase, the features are judged for similarity (similar or dissimilar) (e.g., "I am not like the agent"), relevance of features to user goals (relevant or irrelevant) and valence to goals (positive or negative outcome expectancies). The measures in the encode phase - moderated by the factors in the comparison phase - determine the responses, that is, the levels of involvement with and distance towards the embodied agent. Moreover, the intention to use the agent as a tool indicates actual use and together with involvement and distance, this determines the overall satisfaction of the user with the agent; in our case of Agent A with Agent B. The I-PEFiC model has been formalized in (Bosse et al. 2008).

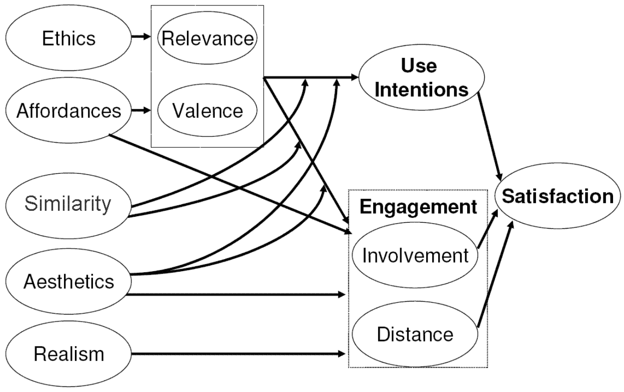

Figure 1: Graphical Representation of I-PEFiC

The model presented in this paper is an extension of a model of affective decision making (Hoorn et al. 2008) based on a formalization of the I-PEFiC model (Bosse et al. 2008) in the LEADSTO environment (Bosse et al. 2007). In this decision making model, decisions are made based on rational as well as affective processing. In the model an agent has desired and undesired goal-states. The agent perceives affordances of the other agents by means of beliefs that the agents facilitate or inhibit reaching certain goal-states. These perceived affordances of other agents are compared to the goal-states it wants to achieve or avoid. While doing this, it can reason about the outcome expectancies of using the other agent for a certain action (e. g., comforting, fighting, or criticizing). In humans, such outcome expectancies lead to certain quick and mostly subconsciously generated action tendencies. In our agents, as in humans, action tendencies influence the experienced involvement and distance towards the other agent. The involvement and distance towards another agent are combined with the use intentions of that agent and the expected utilities of the possible actions to calculate the expected levels of satisfaction of these actions. These expected levels of satisfaction are compared to reach a final decision that is based on rationality as well as affective influences.

We created a library of actions the agents can perform. In this action library, the type of each action is specified. Actions can be specified as

(1) Positive approach
(2) Negative approach
(3) Change
(4) Avoid

In this paper, the action library consists of one action for each type. Comfort is an action of the type positive approach, fight is an action of the type negative approach, criticize is an action of the type change, and avoid obviously is an action of the type avoid. If an agent tries to perform an action of type 1, 2, or 3 towards another agent, while the other agent is avoiding the agent, it will not succeed in performing this action. If an agent performs an action towards another agent, this affects the agent that is the object of the action, as well as the agent that is performing the action itself. The formulas and values not described in this paper can be found in http://www.few.vu.nl/~mpontier/ECMS-2009-Appendices.pdf.

Adjusting the perceived ethics

If an agent performs an action towards another agent, the agent that is the object of the action can change its perception of the goodness and badness of the agent performing the action. For example, if an agent fights another agent, the agent that is the object of this action will probably decrease its perception of the goodness of the fighting agent, and increase its perception of the badness of the fighting agent. To establish this change in perception, the bias (in the range [0, 2]) changes according to the actions that are being performed. A bias > 1 means overestimation, and a bias < 1 means underestimation. To calculate the effect of an action on the bias for perceiving the goodness of the agent performing the action, we have developed the following formula:

$$new\_bias(good) = p_{good} * old\_bias(good) + (1-p_{good}) * v_{(agent,\ action,\ good)}$$

In this formula, new_bias(good) is the new value of the bias, old_bias(good) is the old value of the bias, and the persistency factor p_{good} is the proportion of the old bias that is taken into account to determine the new bias for perceiving goodness of the agent that is performing the action. In this paper, for clarity in the simulation experiments this persistency factor is set to 0.85 for all agents, but this could just as easy be personalized per agent. The new contribution to the bias is $v_{(agent,\ action,\ good)}$, a value that an agent attaches to the goodness of being the object of the performed action. In practice this means that if a certain action is performed towards an agent, the bias of good will move towards the value the agent attaches to being the object of that action. In this paper, all the agents attach the same, arbitrarily chosen values to being the object of actions, but also this could easily be personalized per agent. As the biases are in the range [0, 2], these values are also in the range [0, 2]. The values can be found in Table 1.

Table 1: The Values of Goodness the Agents Attach to Being the Object of an Action

Action	Value
Comfort	1.5
Fight	0.5
Criticize	1.25
Avoid	0.75

The actions for calculating the effect of an action on the bias for perceiving badness are calculated in a similar way as the effect on the bias for perceiving goodness. The only difference is that the values used in this formula are the values the agents attach to the badness of being the object to the performed action. These values can be found in Table 2.

$$new\_bias(bad) = p_{bad} * old\_bias(bad) + (1- p_{bad}) * v_{agent,\ action,\ bad}$$

Table 2: The Values of Badness the Agents Attach to Being the Object of an Action

Action	Value
Comfort	0.5
Fight	1.5
Criticize	1.25
Avoid	1.25

Adjusting the perceived affordances

If an agent performs an action towards another agent, the agent that is the object of the action can also change its beliefs about the affordances of the agent performing the action. Beliefs have a value in the domain [-1, 1]. A belief of 1 represents a strong belief that the statement the belief is about is true, and a belief of -1 represents a strong belief that it is not true. For example, if an agent fights another agent, the belief that avoiding the other agent helps to reduce your anger might increase. To calculate the effect of an action on beliefs about the agent performing the action, we have developed the following formulas:

if beliefchange ≥ 0:
$$new\_belief = old\_belief + belief\_adaptation * beliefchange * ((1 - old\_belief) / 2)$$

if beliefchange < 0:
$$new\_belief = old\_belief + belief\_adaptation * beliefchange * ((1 + old\_belief) / 2)$$

In these formulas, new_belief is the new value of the belief, old_belief is the old value of the belief, and belief_adaptation is a variable, set at 0.1, that determines the speed with which the beliefs are changed when being the object of an action. The beliefchange is a variable in the range [-1, 1] that determines the change of a belief about the performing agent when an agent is performing an action, or is the object of an action performed by another agent. Multiplying with ((1 - old_belief) / 2) when beliefchange is positive, and with ((1 + old_belief) / 2) when beliefchange is negative manages that the values of the beliefs change less if they approach their boundaries, and prevents them from going out of the domain [-1, 1]. In this paper, the values for beliefchange are the same for all agents,

but this could easily be personalized per agent. All beliefchange values can be found in Table 3.
In this table, in the columns the actions that are the cause of the belief change are shown. In the rows the affected beliefs are shown. These beliefs are about an action facilitating a certain goal. For example, if an agent A fights another agent B, agent A will change its belief that comforting agent B will help to reduce his own anger with a beliefchange value of -0.75, as can be seen in Table 3.

Table 3: The beliefchange Values when Actions Are Performed

Affected Belief		Actions causing the belief change			
Action	**Goal**	Comfort	Criticize	Fight	Avoid
Comfort	Self	0.2	-0.2	-0.75	-0.6
Criticize	Self	-0.5	0.9	-0.5	0.6
Fight	Self	-0.9	0.3	0.75	0.4
Avoid	Self	-0.25	0.8	0.9	0.7
Comfort	Others	0.4	0.2	-0.5	-0.3
Criticize	Others	-0.25	-0.25	0.25	-0.4
Fight	Others	-0.8	-0.8	0.25	-0.7
Avoid	Others	-0.35	-0.6	0.75	0.8

Adjusting the emotions of the agents

The actions that an agent performs, and the actions that are performed to an agent, affect the emotions of that agent. The emotion simulated in this paper is the level of anger, but also other or even multiple emotions could be simulated in a similar manner. To calculate the effect of actions performed on the anger level, we have developed the following formula:

$$\text{new\_anger} = p_{anger} / n * \text{old\_anger} + (1 - p_{anger} / n) * \Sigma(\text{changed\_anger}) / n$$

In this formula, new_anger is the new anger level, and old_anger is the old anger level. The persistency factor p_{anger} is the proportion of the old anger level that is taken into account to determine the new anger level. In this paper, for clarity in the simulation experiments, this persistency factor is set to 0.95 for all agents, but this could easily be personalized per agent. The number of actions that is taken into account for calculating the new anger level is represented by n. The persistency factor is divided by n, so that there will be less persistency in the anger level if multiple actions are taken into account.
The new contribution to the anger level is the mean of all the changed_anger variables that are attached to the actions taken into account. This changed_anger is a variable that indicates which value the anger level approaches given a certain action. In practice this means that if a certain action is performed towards an agent, the anger level will move towards the value of the changed_anger attached to that action. For example, if an agent fights another agent, this will make its anger level approach 0.7, because if the anger level is very high, fighting another agent might help to release this anger, although it will never help to decrease it to a low anger level. On the other hand, if an agent has a low anger level, fighting another agent will probably increase the level of anger.
In this paper, all the agents attach the same changed_anger values to actions, but it would be just as easy to let each agent have its own personal values. As the anger levels are in the range [0, 1], the changed_anger values are also in the range [0, 1]. The values used for this paper can be found in Table 4.

Table 4: The changed_anger Values the Agents Attach to Specific Actions

Action	Value if subject	Value if object
Comfort	0.2	0.2
Fight	0.7	0.9
Criticize	0.2	0.6
Avoid	0.3	0.7

SIMULATION EXPERIMENTS

The simulation model introduced in the previous section was used to perform a number of experiments under different parameter settings. In the experiments, the three agents Harry, Barry, and Gary followed a (fictitious) anger management therapy. An infinite number of actions could be inserted in the system, but for clarity, nonetheless, for each action type we inserted only one instance of an action. The action to comfort another agent is a type of positive approach whereas the action for negative approach was to fight another agent. Criticizing another agent was the action associated with change, and the action for avoiding the agent was to simply move away from another agent. In the simulation experiments, a calculation step takes one timepoint and an action takes five timepoints. After these five timepoints the action taken can be reconsidered and changed. The results of the experiment are described below. Notice that in each graph in this paper, the scale on the y-axis can differ.
We expected that the system can generate simulations in which actions lead to changes in perception in such a way, that agents change their mind and perform another action than they would have done before their perception was changed.

Experiment 1: Baseline

To start, an initial experiment was performed that served as a control condition for the rest of the experiments. In this experiment, the designed features for beautiful and ugly (aesthetics), good and bad (ethics), and realistic and unrealistic (epistemics) are all set to 0.5 (see Figure 1). All biases for perceiving these features are set to the neutral value of 1. The agents have no beliefs about other agents (all beliefs are set to 0), and the only ambition they have is to reduce their own anger with an ambition level of 0.5.

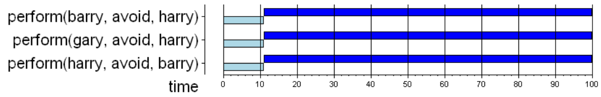

Figure 2: Simulation Results of Experiment 1

In Figure 2, along the X-axis the timepoints are shown, and along the Y-axis several statements are shown. A dark blue bar means the statement holds at that timepoint, and a light blue bar means the statement is false at that timepoint. As can be seen in Figure 2, these settings result in Barry and Harry avoiding Gary, and Harry avoiding Barry after timepoint 11.

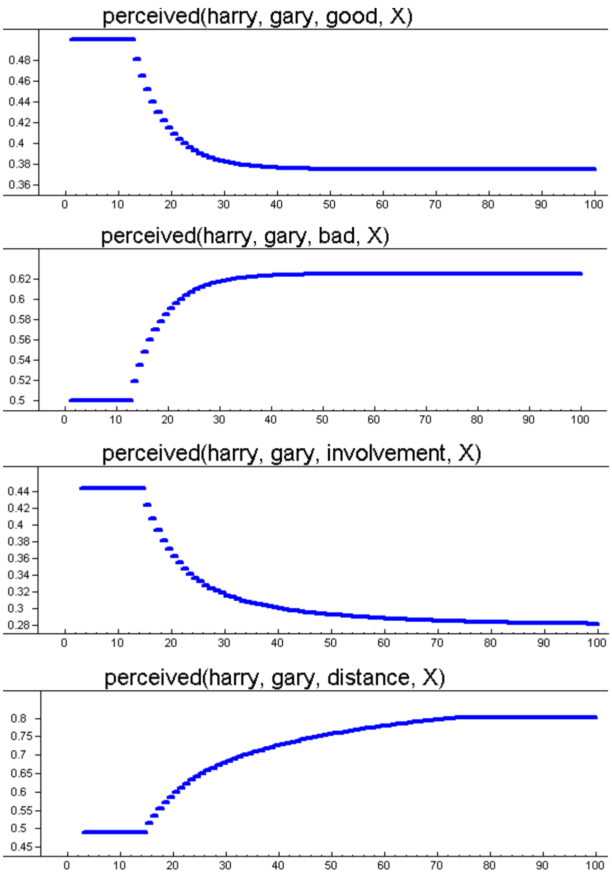

Figure 3: Harry's Involvement, Distance, Perceived Goodness and Perceived Badness of Gary

As can be seen in Figure 3, the agents who are being avoided increase their distance from 0.49 to 0.82, and decrease their involvement from 0.44 to 0.28 towards the avoiding agent. This happens because the perceived goodness of these agents decreases from 0.5 to 0.38 and badness increases from 0.5 to 0.62. Being avoided also causes changes in beliefs.

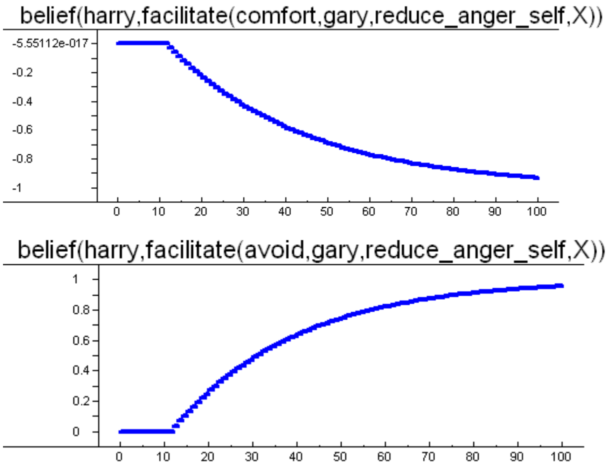

Figure 4: The Beliefs of Barry about Harry during Experiment 1.

As can be seen in Figure 4, the agents who are being avoided start to think that avoiding, criticizing, or fighting the avoiding agent will help them to reduce their own anger, and that comforting the avoiding agent will inhibit their goal of reducing their own anger. As initially the agents have no beliefs about each other, the agents have no intentions to use each other at the start of the simulation.

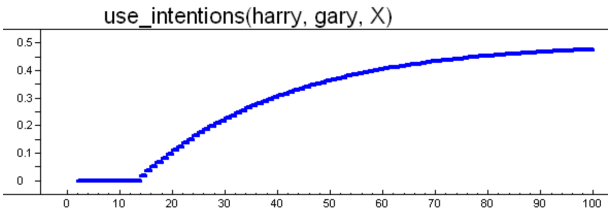

Figure 5: Harry's Intentions to Use Gary during Experiment 1.

Because the agents that are being avoided start to have beliefs about the avoiding agents, their intentions to use that agent also increase from 0 to 0.46, as can be seen in Figure 5.

Being avoided also changes the anger level. As can be seen in Figure 6, the anger level of Barry decreases from 0.60 to 0.50, and the anger level of Harry decreases only from 0.60 to 0.57 (notice the differences in scale on the y-axis). Harry's anger level decreases less because he is avoided by both Barry and Gary. Gary's anger level reduces to 0.30, which is even more than that of Barry, because he is not being avoided by any agent. This shows that being avoided by multiple agents has a greater impact than being avoided by only one agent.

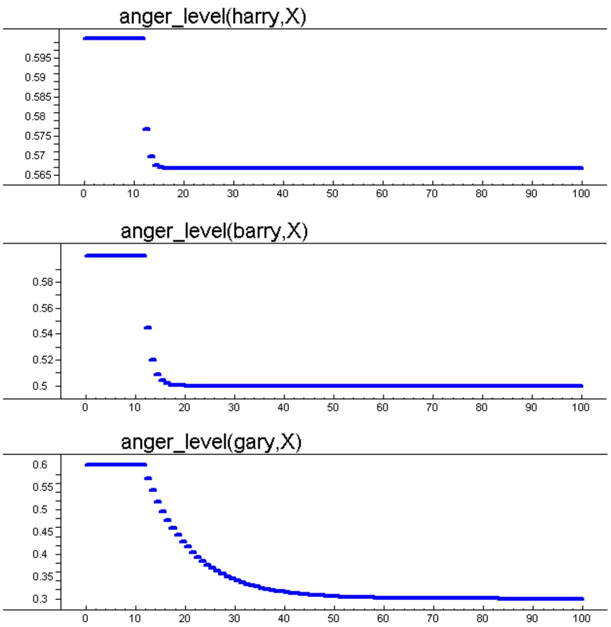

Figure 6: Anger Levels of the Agents in Experiment 1

Experiment 2: Harry beliefs he should not avoid Barry

We performed an experiment in which Harry has a strong belief (value = 1) that avoiding Barry will inhibit his goal of reducing his own anger. Harry also has a stronger ambition to reduce his own anger, with a value of 1 instead of 0.5. The remaining variables have the same values as in the baseline condition.

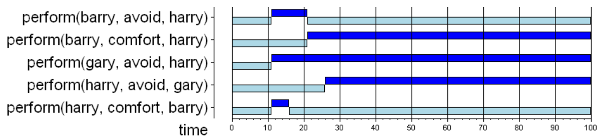

Figure 7. Simulation Results of Experiment 2.

Due to this, instead of avoiding Barry, Harry now tries to comfort Barry at time-point 11, as can be seen in Figure 7. However, he does not succeed in comforting Barry, as Barry is avoiding Harry, just like in the baseline experiment. This causes Harry to stop avoiding Barry at time point 16 and increases his anger level to from 0.53 to 0.70, as can be seen in Figure 8.

Not being comforted anymore also slightly increases Barry's anger level from 0.25 to 0.27 in the five following time steps. In the mean time, Barry has observed that Harry tried to comfort him, which decreases his anger level from 0.60 to 0.25. It also changes his perceptions of the ethics of Harry. As can be seen in Figure 9, Barry starts to see Harry as a good guy, with a perceived goodness of 0.64 and a perceived badness of 0.36.

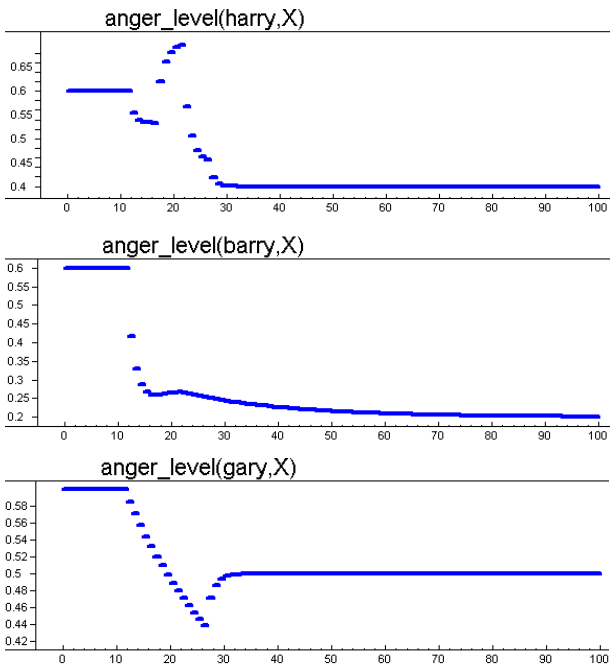

Figure 8: Anger Levels of the Agents in Experiment 2

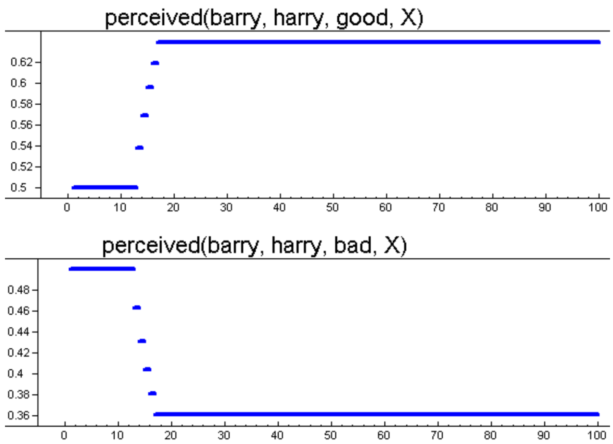

Figure 9: Barry's Perception of the Ethics of Harry during Experiment 2

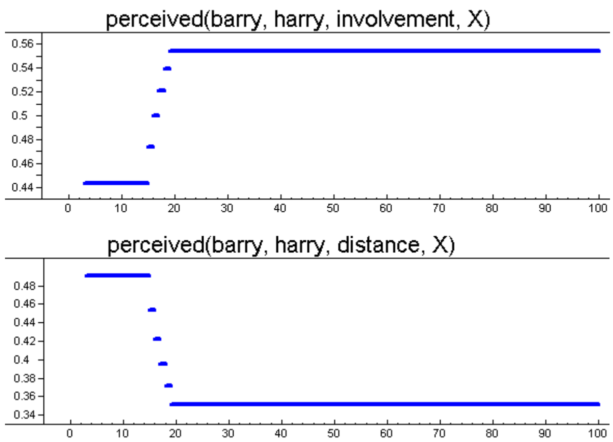

Figure 10: Barry's Involvement and Distance towards Harry during Experiment 2

This also slightly increases Barry's belief that comforting Harry will help him to reduce his anger, while his beliefs that avoiding, criticizing or fighting Barry will help him to reduce his own anger are slightly reduced.
These changes in beliefs and perceptions cause the involvement of Barry towards Harry to increase from 0.44 to 0.56 and the distance from Barry to Harry to decrease from 0.50 to 0.34, as can be seen in Figure 10. This causes him to start comforting Harry instead of avoiding him at timepoint 21.

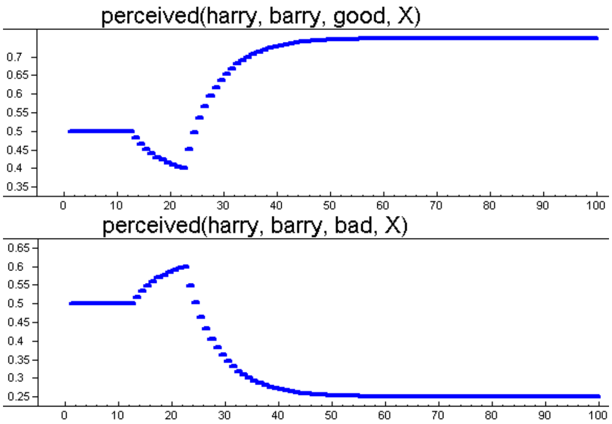

Figure 11: Harry's Perception of the Ethics of Barry during Experiment 2

However, in the mean time, being avoided by Barry has changed Harry's opinion about him. The perception of his goodness has decreased from 0.50 to 0.39, and the perception of his badness has increased from 0.50 to 0.61, as can be seen in Figure 11.

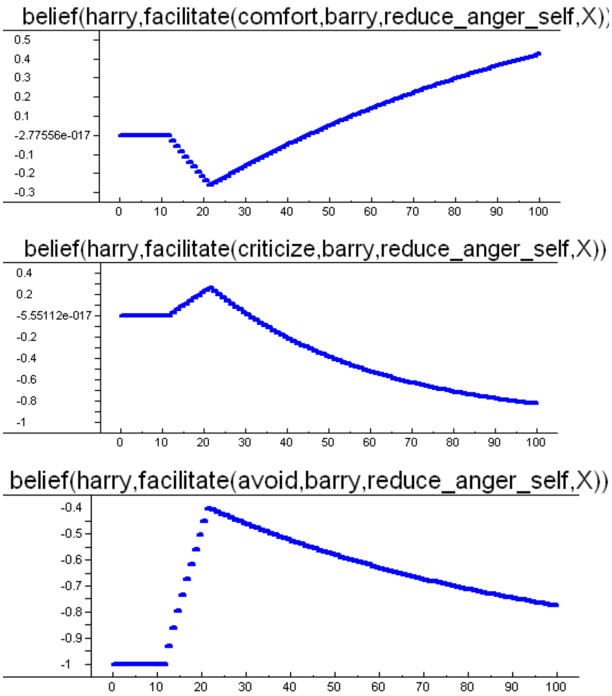

Figure 12: Harry's Beliefs about Barry during Experiment 2

As can be seen in Figure 12, also Harry's beliefs that fighting, criticizing, and especially avoiding Barry will help to reduce his own anger increase, whereas his belief that comforting Barry will help to reduce his own anger decreases. In the mean time, being avoided by Gary has decreased Harry's involvement towards Gary, increased his distance towards Gary, and has increased his beliefs that criticizing, fighting, or avoiding Harry will help him to reduce his own anger in a similar way as in experiment 1.

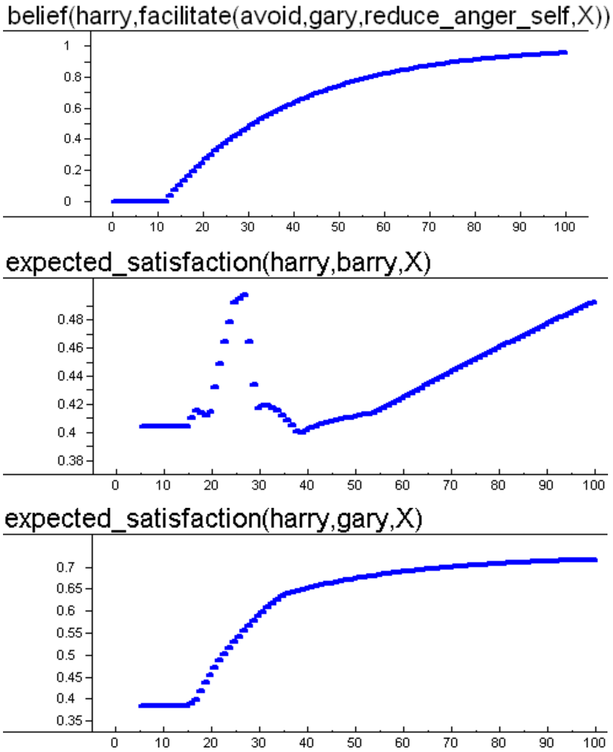

Figure 13: Harry's Belief that Avoiding Gary Will Help Him to Reduce his own Anger, and his Expected Satisfaction of Performing an Action towards Barry and Gary during Experiment 2

As can be seen in Figure 13, this increases Harry's expected satisfaction of performing an action towards Gary, namely avoiding him, and this expected satisfaction exceeds the expected satisfaction of performing an action towards Barry. This causes Harry to change his mind and start avoiding Gary at timepoint 26. In the mean time, after Barry started to comfort Harry at timepoint 21, Harry's anger level decreases from 0.70 to 0.40, and Barry's anger level from 0.28 to 0.20.
Also the distance from Harry to Barry decreases from 0.66 to 0.09 and his involvement increases from 0.37 to 0.67, as can be seen in Figure 14. This happens due to an increase in Barry's perception of Harry's goodness from 0.39 to 0.75 and a decrease in perceived badness from 0.61 to 0.25, and because his increasing belief that comforting Harry will help him to reduce his own anger, and fighting, criticizing or avoiding Harry will inhibit his goal to reduce his own anger.

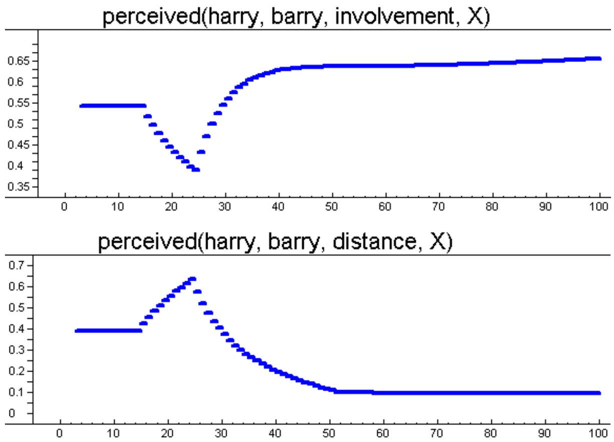

Figure 14: Harry's Involvement and Distance towards Barry during Experiment 2

DISCUSSION

In this paper, we presented an extension of a formalized model of affective decision making (Hoorn et al. 2008), based on the informally described I-PEFiC model (Van Vugt et al. 2006, Van Vugt et al. 2007, Van Vugt 2008). This extension manages that the actions the agents undertake have an effect on the agents. The agents change their perceptions and beliefs about other agents if actions are taken. Further, the anger level of the agents is being simulated. Simulation experiments have been performed to show how the actions affect the agents. Experiment 1 showed that if multiple agents perform an action towards another agent, this has a bigger effect on its anger level than if only one agent would perform that action. Experiment 2 showed that if an agent performs an action towards another agent, this can change the beliefs and the perceptions of the other agents so much that the other agent changes its mind and chooses to perform another action than it was currently doing, leading to a completely different situation than in experiment 1, confirming our expecations as mentioned in section 4. These results are as would have been expected from the theory (Hoorn et al. 2007, Van Vugt et al. 2006, Van Vugt 2008).

In this paper, the simulation experiments are performed in the domain of anger management therapy. The only simulated emotion is anger, and for each type of action there is only one possible action to perform. However, this model could be used for any type of domain, with other, or multiple emotions simulated at the same time. Also as many actions as desired could be added to the action library. This way, the model could be used to perform simulations involving decision-making, emotions, and changing perceptions for any domain.

Of course, a lot of additions could still be made to our model. For instance, the persistency factors for changing the beliefs and perceptions of other agents could be made dependent on the period of time the agents know each other (or on the number of interactions).

Existing models of decision-making usually assume this process to be rational, which would exclude the possibility of emotions playing a role other than disturbing the process (Gutnik et al. 2006). However, humans often make irrational decisions. A good example for this is the Ultimatum game (Thaler 1988), for which behavioral research showed that low offers (20% of total amount) have a 50% chance of being rejected. Participants reported that they found low offers unfair, and therefore out of anger they selected the irrational option (Gutnik et al. 2006).

Models of decision making usually have a hedonic bias, and generally try to find the action with the highest expected utility. Some decision theoretic models, such as (Gmytrasiewycz and Lisetti 2001), take emotions into account, but in those models, emotions merely confirm good rational decisions – emotional states as modes of decision making. However, these models cannot explain irrational behavior, where actions with a (relatively) low expected utility are chosen. Our balancing model takes the expected utility as well as involvement-distance trade-offs into account. This way, situations in which emotions overwhelm rationality can be explained and simulated.

There have been a number of approaches to model decision-making based on emotions in autonomous agents. However, none of these studies uses a detailed model of perception of others to explain how these affective influences in the decision making process are generated. Usually, these models somehow assume that emotions are there. For instance, (Velasquez 1998) presents a model of emotion-based decision-making, which is an extension of a previous model (Velasquez 1997). The model that is presented in this paper assumes a perceptual system, but how this perceptual system actually works is not considered in the paper.

Ahn & Picard (Ahn & Picard 2005) present a computational framework of affective-cognitive learning and decision making for affective agents. This model integrates affect and cognition, where 'internal rewards from cognition and emotion' and 'external rewards from the external world' serve as motivations for learning and decision making. In this model emotions are generated based on these rewards, but perceiving others in the world is left out of consideration.

In future research, we plan to combine the model with an existing computational model for emotion regulation (Bosse et al. 2007). Whereas the current model focuses on the elicitation of emotion, that model addresses the regulation of emotion. Further, we intend to explore where these models and the EMA model (Gratch and Marsella 2004) complement each other, and use this to further refine the models.

Finally, in a later stage of the project, we will confront our formalization with empirical data of human affective trade-off processes. As soon as the model is validated and adapted, we will start exploring the possibilities to build a robot that can interact with real humans. We hope to develop a robot that can communicate affectively with humans in a more natural way, that is, with a mind of its own, in pursuit of its own goals, and acting emotionally intelligent.

ACKNOWLEDGEMENTS

We kindly want to thank Tibor Bosse for his input to this paper, and for giving feedback at earlier drafts.

REFERENCES

Ahn, H., and Picard, R. W. 2005. "Affective-cognitive learning and decision making: A motivational reward framework for affective agents." In Tao, J., Tan, T., and Picard, R. W., (Eds.), *Affective Computing and Intelligent Interaction, First International Conference*, 866–873

Ashby, R. 1960. *Design for a Brain*. Second Edition. Chapman &Hall, London. First edition 1952.

Bargh, J. A., and Chartrand, T. L. 1999. "The Unbearable Automaticity of Being." *American Psychologist*, 54, 462-479.

Barringer, H., Fisher, M., Gabbay, D., Owens, R., and Reynolds, M. 1996. *The Imperative Future: Principles of Executable Temporal Logic*. John Wiley & Sons.

Beck, A. T. 1987. "Cognitive models of depression." *Journal of Cognitive Psychotherapy, An International Quarterly*, 1, 5-37.

Bosse, T., Hoorn, J. F., Pontier, M., and Siddiqui, G. F. 2008. "A Robot's Experience of Another Robot: Simulation." In Sloutsky, V., Love, B. C., and McRae, K. (Eds.), *Proceedings of the 30th International Annual Conference of the Cognitive Science Society*. CogSci'08. 2498-2503.

Bosse, T., Jonker, C. M., Meij, L., van der, and Treur, J. 2007. "A Language and Environment for Analysis of Dynamics by Simulation." *International Journal of Artificial Intelligence Tools*, 16, 435-464.

Bosse, T., Pontier, M., and Treur, J. 2007. "A dynamical system modelling approach to Gross' model of emotion regulation." In Lewis, R. L., Polk, T. A., and Laird, J. E. (Eds.), *Proceedings of the 8th International Conference on Cognitive Modeling*. ICCM'07. Taylor and Francis, 187-192.

Fredrickson, B. L. 2001. "The Role of Positive Emotions in Positive Psychology: The broaden-and-build theory of positive emotions." *American Psychologist*, 56, 218-226.

Frijda, N. H. 1986. *The emotions*. New York: Cambridge University.

Gmytrasiewycz, P. J., and Lisetti, C. L. 2001. "Emotions and Personality in Agent Design and Modeling." In Bauer, M., Gmytrasiewicz, P. J., and Vaassileva, J. (Eds.), *Lecture Notes in Computer Science*, Springer-Verlag Berlin Heidelberg, 237-239.

Gratch, J., and Marsella, S. 2004. "A Domain-independent Framework for Modeling Emotion." *Journal of Cognitive Systems Research*, 5, 4, 269-306.

Gutnik, L. A., Hakimzada, A. F., Yoskowitz, N. A., and Patel, V. L. 2006. "The role of emotion in decision-making: A cognitive neuroeconomic approach towards understanding sexual risk behavior." *Journal of Biomedical Informetics*, 39, 720-736.

Hoorn, J. F., Konijn, E. A., Van Vliet, H., and Van der Veer, G. 2007. "Requirements change: Fears dictate the must haves; desires the won't haves" In *Journal of Systems and Software*, Vol. 80, pp. 328-355.

Hoorn, J. F., Pontier, M., and Siddiqui, G. F. 2008. "When the User is Instrumental to Robot Goals: First Try – Agent Uses Agent." In *Proceedings of the 2008 IEEE/WIC/ACM International Conference on Intelligent Agent Technology*. IAT'08, 296-301.

Marsella, S., and Gratch, J. 2003. "Modeling coping behavior in virtual humans: Don't worry, be happy." In *Proceedings of Second International Joint Conference on Autonomous Agents and Multiagent Systems*. AAMAS'03. ACM Press, 313-320.

Port, R. F., and Gelder, T. van (Eds.). 1995. *Mind as Motion: Explorations in the Dynamics of Cognition*. MIT Press, Cambridge, Mass.

Thaler, R. H. 1988. "The Ultimatum Game." *Journal of Economic Perspectives*, 2, 195-206.

Van Vugt, H. C. 2008. *Embodied Agents from a User's Perspective*. Doctoral dissertation, VU University, Amsterdam.

Van Vugt, H. C., Hoorn, J. F., Konijn, E. A., and De Bie Dimitriadou, A. 2006. "Affective affordances: Improving interface character engagement through interaction." *International Journal of Human-Computer Studies*, 64, 874-888.

Van Vugt, H. C., Konijn, E. A., Hoorn, J. F., Eliëns, A., and Keur, I. 2007. "Realism is not all! User engagement with task-related interface characters." *Interacting with Computers*, 19, 267-280.

Velásquez, J. D., 1997. "Modeling Emotions and Other Motivations in Synthetic Agents." In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*, 10-15

Velásquez, J. D., 1998. "Modeling Emotion-Based Decision Making." In Cañamero, D. (Eds.): *Emotional and Intelligent: The Tangled Knot of Cognition*, 164-169

http://www.few.vu.nl/~mpontier/ECMS-2009-Appendices.pdf

AUTHOR BIOGRAPHIES

MATTHIJS PONTIER was born in Amsterdam, The Netherlands, and went to the VU University of Amsterdam. He obtained his Bachelor degree for Artificial Intelligence as well as for (Cognitive) Psychology, and got his Masters degree in the direction of Cognitive Science in 2007. Currently, he is employed as a Ph.D. student at the VU University of Amsterdam, both at the department of Artificial Intelligence and the Center for Advanced Media Research Amsterdam (CAMeRA). His e-mail address is: `mpontier@few.vu.nl` and his website can be found at `http://www.few.vu.nl/~mpontier`

GHAZANFAR FAROOQ SIDDIQUI was born in Mandi Bahaulddin, Pakistan. He did his Master Degree in Computer Science at Quaid-i-Azam University, Islamabad**,** Pakistan. He spent seven years teaching graduate and under-graduate courses at University of Agriculture, Faisalabad and Quaid-i-Azam University, Islamabad. Since 2006, he is a PhD student at VU University of Amsterdam, both at the department of Artificial Intelligence and the Center for Advanced Media Research Amsterdam (CAMeRA). His e-mail address is: `ghazanfa@few.vu.nl` and his website can be found at `http://www.few.vu.nl/~ghazanfa`