When the User is Instrumental to Robot Goals

Johan F. Hoorn, Matthijs Pontier, and Ghazanfar F. Siddiqui
VU University Amsterdam, Department of Artificial Intelligence
Center for Advanced Media Research Amsterdam (CAMeRA)
De Boelelaan 1081a, 1081HV Amsterdam, The Netherlands
jf.hoorn@camera.vu.nl; {mpontier, ghazanfa}@few.vu.nl
http://www.few.vu.nl/~{mpontier, ghazanfa, jfhoorn}

Abstract

To create a robot with a mind of its own, we extended a formalized version of a model that explains human-robot interaction with mechanisms for goal-directed behavior. By running simulation experiments, we found that robots could perceive affordances in other agents to achieve their goals and suppress rational decisions in favor of affective decisions, given baseline involvement or distancing tendencies. Limitations are that models of situation selection are still wanted and empirical validation is needed. However, our good-bad balancing approach explains more complex phenomena in (affective) decision making than hedonic-bias models do.

1. Introduction

Recently, much 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 [13], stress and workload [6], and moods [2]). Let alone actually understanding and reacting empathically to the emotional state of other agents, or human users. Since a requirement of virtual agents is to closely mimic human affective behavior, this is a problem that should be solved. Previous research has shown that closely mimicking humans is important for an agent to increase human involvement in a virtual environment (e.g., [16]). Existing systems based on IVAs typically lack abilities to show emotions (not only by means of facial expressions, but also by behavior), and to interpret those of others. This means that existing systems based on IVAs can be made more effective.

An important view from emotion psychology is that emotions are goal-driven. The emotional system scans the environment for relevant stimuli that are either potentially beneficial or harmful for the concerns, motives, and goals of the individual ([7], p. 494, p. 463).

In the robot world, the user can be seen as a personal friend as well as a means to an end [11]. With regard to being a personal friend, the Interactive model of Perceiving and Experiencing Fictional Characters (I-PEFiC) serves as a starting point [15]. Within this framework, the robot can calculate a trade-off between what involves the robot with the user (e.g., user is skilled) and what keeps the robot at a distance (e.g., user mistreats my hardware) [3]. In addition, use intentions are calculated that prompt the robot to undertake action in favor or against the user (ibid.)

These actions are based on robot goals, which play a role in the judgment formation of the robot about its user. There are eight (2^3) possible types of judgments a robot can have about how the features of a user afford the achievement of different robot goals or not (Table 1) (cf. [12]). A judgment consists of an ontological statement about the user plus a measure of agreement. Each constituent in the judgment evokes a positive (p) or negative (n) covert response. During the weighing, mixed emotions occur, which is a somewhat confusing experience. That is why ① (Table 1) is the preferred mode of conversation and the easiest statement to respond to. Because affordances have predictive power for user engagement [16], all the n-responses that occur during weighing feed into distance; all the p-responses into involvement. The action tendencies that are connected to positive or negative valence ([7], p. 207) will feed into the intentions to 'make use' of the user and above threshold, the robot shows overt behavior (e.g., to converse with the user, kick, hug).

We came to the following general hypothesis, which depends on the eight possibilities tabulated in Table 1 for its outcome.

Table 1. Robot judgments on user as means to an end, resulting in valencies that precede action tendencies

Judgment	Means	Affords	Goal	Agreement[1]	Weighing	Valence to means[3]
①	Skilled user ^{p}	facilitates ^{p}	Robot efficiency ^{p}	Agree ^{p}	p = (p · p) · p = (p) · p =	p
				Disagree[2] ^{n}	p = (p · p) · n = (p) · n =	n
②	Skilled user ^{p}	inhibits ^{n}	Robot efficiency ^{p}	Agree ^{p}	p = (n · p) · p = (n) · p =	n
				Disagree ^{n}	p = (n · p) · n = (n) · n =	p
③	Skilled user ^{p}	facilitates ^{p}	Robot inefficiency ^{n}	Agree ^{p}	p = (p · n) · p = (n) · p =	n
				Disagree ^{n}	p = (p · n) · n = (n) · n =	p
④	Skilled user ^{p}	inhibits ^{n}	Robot inefficiency ^{n}	Agree ^{p}	p = (n · n) · p = (p) · p =	p
				Disagree ^{n}	p = (n · n) · n = (p) · n =	n
⑤	Unskilled user ^{n}	facilitates ^{p}	Robot efficiency ^{p}	Agree ^{p}	n = (p · p) · p = (p) · p =	p
				Disagree ^{n}	n = (p · p) · n = (p) · n =	n
⑥	Unskilled user ^{n}	inhibits ^{n}	Robot efficiency ^{p}	Agree ^{p}	n = (n · p) · p = (n) · p =	n
				Disagree ^{n}	n = (n · p) · n = (n) · n =	p
⑦	Unskilled user ^{n}	facilitates ^{p}	Robot inefficiency ^{n}	Agree ^{p}	n = (p · n) · p = (n) · p =	n
				Disagree ^{n}	n = (p · n) · n = (n) · n =	p
⑧	Unskilled user ^{n}	inhibits ^{n}	Robot inefficiency ^{n}	Agree ^{p}	n = (n · n) · p = (p) · p =	p
				Disagree ^{n}	n = (n · n) · n = (p) · n =	n

[1] Attribution of truth according to robot's world view or 'belief system'
[2] Gray cells indicate an unconventional, counter-intuitive, belief system that urges to adapt conventional theory
[3] If valence is positive, an action tendency to approach the means (here, the user) occurs. If valence is negative, depending on situation, context, or personality, an action tendency to avoid, attack, remove, or change the means occurs (e.g., the robot starts educating the user)

H1: Features of a user or other agent are means to afford robot goals, and through weighing lead to a measure of valence toward that means, propelling action tendencies to approach, avoid, attack, change, or do nothing with the means.

We will test this hypothesis by performing simulation experiments on the formalized model, under various parameter settings.

2. Implementation

I-PEFiC is a model (Figure 1) that is empirically well validated [15]. In the formalization of the I-PEFiC model [3], the robot bases its involvement-distance trade-off towards the user or other agents on ethics, aesthetics, affordances, realism and similarity (cf. [15]). In the formalization of [3], the value for affordances was based on the expected possibilities to communicate with another agent. This is a simplification, as perceived affordances should always relate to goals [16].

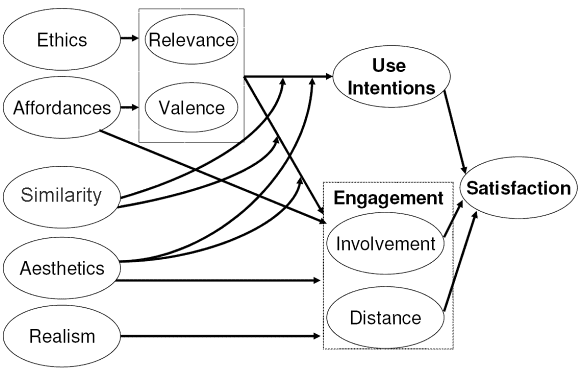

Figure 1. Graphical Representation of I-PEFiC

In this paper, the formalized I-PEFiC model [3] is extended with goal-directed, rational behavior, and implemented in the LEADSTO modeling environment [4]. The model in this paper, however, bases the affordances and their effects on use intentions, involvement, and distance on the goals of the robot, and how the affordances of the agent that are perceived by the robot relate to those goals. Further, we somewhat modified the formalization, so that it better fits the I-PEFiC model as described in [15]. Space limits a full description of the formalization, therefore see [17].

To relate affordances to goals, the robot has goals it wants to achieve: desired goals. For example, a robot wants maintenance, power supply, and a knowledge base. The robot also has goals it wants to avoid: undesired goals. For example, a robot does not want to hurt its owner nor does it want to be destroyed.

The goals of the robot have a certain value for *valence* [-1, 1] and *relevance* [0, 1]. If a goal is desired, it will have a positive valence. If a goal is undesired, it will have a negative valence. If one goal is more relevant for the robot than another, the relevance of that one goal is higher. By multiplying valence with relevance, the level of ambition for a goal is calculated.

Within the system, robots perceive that agents, users, or other things have features. These features can afford a robot to perform a certain action, which affects achieving a certain goal. In other words, these features are means to an end (Table 1).

In the system, robots perceive the affordances of each other, and of other things. These perceived affordances are not necessarily the same as the affordances meant by the designer. For instance, the designer can design a chair to sit on (designed affordances), but an agent could also use this chair to beat up another agent (perceived affordances).

The robots in the system can compare their perceived affordances of other agents or things to the goals they want to achieve or avoid. While doing this, they can reason about the outcome expectancies of using the other agents for a certain action (e.g., speech acts, kicking, or hugging). For example, if a robot has the desired goal to go inside a house, it could believe that the action to open the door of that house could be an action that facilitates this goal.

In humans, such outcome expectancies lead to certain quick and mostly subconsciously generated action tendencies. In our robots, action tendencies influence the experienced involvement and distance towards the other agent. The four possible action tendencies to perform towards another agent in our system are:

(1) Positive approach (e.g., to hug)
(2) Negative approach (e.g., to attack)
(3) Change (e.g., to teach; mix of positive / negative approach)
(4) Avoidance

2.1. Calculating Expected Utilities

In the system, the robots can perform actions to reach their goals. The system contains a library of goals, and each robot has a level of ambition for each goal. There are goals the robot wants to reach, and goals that the robot wants to avoid, all with several levels of importance. The levels of ambition the robot attaches to the goals are represented by a real value between [-1, 1], where a negative value means that the goal is undesired and a positive value means that the goal is desired. A bigger value means the goal is more important for the robot.

The robots can perform actions to reach their goals. The system contains a library of actions from which the robots can choose. The robot has a belief about each action that it will inhibit or facilitate a certain goal. Its estimation of the facilitation of the goal by the action is represented by a real value between [-1, 1], -1 being full inhibition, 1 being full facilitation. The following formulas are used to calculate the expected utilities of actions.

$$\text{ExpectedUtility}_{(\text{Action, Agent, Goal})} = \text{Belief}_{(\text{facilitates(Action, Agent, Goal)})} * \text{Ambition}_{(\text{Goal})}$$

$$\text{ExpectedUtility}_{(\text{Action, Agent})} = \Sigma(\text{ExpectedUtility}_{(\text{Action, Agent, Goal})})$$

Given the level of ambition for a goal and the believed facilitation of a goal by an action towards another agent, the robot calculates the expected utility of performing that action towards that agent regarding that goal by multiplying the believed facilitation of the goal with the level of ambition for the goal.

Because an action usually affects several goals that might be conflicting, the 'general' expected utility of performing a certain action towards an agent is calculated by summing all expected utilities regarding all goals in the system that are related to the action.

Because a robot usually performs only one action at a time with respect to another agent, the use intentions of the robot are calculated by taking the maximum expected utility of all actions the robot can perform to the agent. Agents that facilitate desired or inhibit undesired goals raise positive use intentions with the robots, while agents that facilitate undesired or inhibit desired goals will raise negative use intentions.

$$\text{UseIntentions}_{(\text{Robot,Agent})} = \max(\text{ExpectedUtility}_{(\text{Robot,Action,Agent})})$$

2.2. Effects on involvement and distance

In the action library, the type of each action is specified. Actions can be specified as (1) Positive approach, (2) Negative approach, (3) Change, or (4) Avoid.

The heuristic to calculate the expected utilities of actions, as described in the previous paragraphs, is also used to generate action tendencies. So if an agent has a high expected utility for a certain action, it will also generate a strong action tendency for that specific action.

$$\text{ActionTendency}_{(\text{Action, Agent})} = \text{ExpectedUtility}_{(\text{Action, Agent})}$$

The generated action tendencies are used to calculate the effect of the affordances of an agent on the robot's involvement with and distance towards that agent. To calculate this effect, a weighed sum of all the action tendencies is taken, as can be seen in the formulas below. In these formulas, the β's represent the weights of the action tendencies on involvement and distance.

$$\text{Effect of Affordances on Involvement} = \beta_{i\leftarrow na}*AT_{neg\_appr} + \beta_{i\leftarrow pa}*AT_{pos\_appr} + \beta_{i\leftarrow ch}*AT_{change} + \beta_{i\leftarrow av}*AT_{avoid}$$

$$\text{Effect of Affordances on Distance} = \beta_{d\leftarrow na}*AT_{neg\_appr} + \beta_{d\leftarrow pa}*AT_{pos\_appr} + \beta_{d\leftarrow ch}*AT_{change} + \beta_{d\leftarrow av}*AT_{avoid}$$

Table 2. Weights of action tendencies on robot's involvement and distance

Weight	Value	Weight	Value
$\beta_{i\leftarrow pa}$	0.75	$\beta_{d\leftarrow pa}$	-0.75
$\beta_{i\leftarrow na}$	0.25	$\beta_{d\leftarrow na}$	0.75
$\beta_{i\leftarrow ch}$	0.50	$\beta_{d\leftarrow ch}$	0.50
$\beta_{i\leftarrow av}$	-0.50	$\beta_{d\leftarrow av}$	0.50

As can be seen in Table 2, the generated action tendencies classified as negative approach increase the robot's involvement a little and increase distance a lot. If the robot feels the tendency to negatively approach the user or another agent, this will slightly increase its involvement with that agent, as involvement represents a tendency to approach [van Vugt, 2008], but simultaneously will increase its distance toward the user or agent, as negative approach implies quite some distance. The generated action tendencies classified as “positive approach” increase involvement and decrease distance. The generated action tendencies classified as “change” increase involvement as well as distance. Finally, the generated action tendencies classified as “avoid” decrease involvement and increase distance. The effects of affordances on involvement and distance are computed as can be found in [17].

2.3. Making a decision

All possible actions in the system are related to other agents. In the decision process, the robot first selects an agent to perform the action on. To do this, for all possible agents it meets, the robot calculates the expected satisfaction (Figure 1) of interacting with that agent, using the following formulas:

$$\text{Involvement-Distance-Tradeoff} = \gamma * \max(I, D) + (1-\gamma)*(I+D)/2$$
$$\text{Expected\_Satisfaction}_{(\text{Robot, Agent})} = \beta_{es \leftarrow idt} * IDT + \beta_{es \leftarrow ui} * UI$$

The expected satisfaction is calculated by trading involvement (I) for distance (D) as described in [3], and taking a weighed mean of the involvement-distance tradeoff (IDT) and the use intentions (UI). In this paper, the weight of the involvement-distance-tradeoff for expected satisfaction is taken as 0.8 and the weight of use intentions for expected satisfaction is taken as 0.2, which is arbitrary and needs future empirical validation. Thus, the robot bases its choice which agent to interact with on the rationally generated use intentions, as well as on the more affectively generated trade-off between involvement and distance. The robot approaches the agent that promises the highest expected satisfaction during interaction.

Once the robot has selected an agent to interact with, it decides which action to take. For each possible action, it calculates the expected satisfaction, according to the following formulas:

Expected Satisfaction Positive Approach =
$$\beta_{espa \leftarrow i} * I + \beta_{espa \leftarrow d} * (1-D) + \beta_{espa \leftarrow eu} * EU_{act}$$

Expected Satisfaction Negative Approach =
$$\beta_{esna \leftarrow i} * (1-I) + \beta_{esna \leftarrow d} * D + \beta_{esna \leftarrow eu} * EU_{act}$$

Expected Satisfaction Change =
$$\beta_{esch \leftarrow i} * I + \beta_{esch \leftarrow d} * D + \beta_{esch \leftarrow eu} * EU_{act}$$

Expected Satisfaction Avoid =
$$\beta_{esav \leftarrow i} * (1-I) + \beta_{esav \leftarrow d} * D + \beta_{esav \leftarrow eu} * EU_{act}$$

The expected satisfaction of doing a specific action with a certain agent is calculated by taking a weighed sum of the robot’s involvement and distance, and the expected utility of the particular action. In the paper, these weights are taken as can be seen in Table 3, but they can differ per robot, according to its ‘personality’. A robot might have a high threshold for negative approach, while another robot

Table 3. Values for weights of involvement, distance, and expected utility on the expected satisfaction of performing a type of action

Weight	Value	Weight	Value
$\beta_{espa \leftarrow i}$	0.4	$\beta_{esch \leftarrow i}$	0.4
$\beta_{espa \leftarrow d}$	0.4	$\beta_{esch \leftarrow d}$	0.3
$\beta_{espa \leftarrow eu}$	0.2	$\beta_{esch \leftarrow eu}$	0.3
$\beta_{esna \leftarrow i}$	0.4	$\beta_{esav \leftarrow i}$	0.5
$\beta_{esna \leftarrow d}$	0.4	$\beta_{esav \leftarrow d}$	0.3
$\beta_{esna \leftarrow eu}$	0.2	$\beta_{esav \leftarrow eu}$	0.2

easily approaches other agents negatively. If the robot has a high level of involvement with and a low level of distance towards an agent, it will approach the agent positively. If the agent has a low level of involvement and a high level of distance, it approaches the agent negatively or avoids it. If the agent evokes a high level of involvement as well as a considerable level of distance, the robot is most likely to try to change the agent, for example, to teach it. The robot will pick the action with the highest expected satisfaction and perform it.

If the effects of the robot’s actions on user or agent are captured and analyzed, this model could be used to let robots interact with each other in a meaningful way, based not only on rationality, but also on affective tendencies.

3. Simulation Results

To test our hypothesis H1, the simulation model introduced in the previous section was used to perform a number of experiments under different parameter settings. In each experiment, three robots (Harry, Barry, and Gary) followed a (fictitious) anger management therapy. In this setting, an infinite number of actions can be inserted in the system, but for clarity, for each action type we inserted only one instance of an action. The action related with positive approach was to comfort the other agent, whereas the action for negative approach was to hit the other agent. Criticizing another agent was the action associated with change, and the action for avoiding the agent was to simply move away from it. The results of the experiments are described below.

Baseline Condition.

To start, an initial experiment was performed that served as a control condition for the remaining experiments. In this condition, the designed features for beautiful and ugly (aesthetics), good and bad (ethics), and realistic and unrealistic (epistemics) were set to 0 (see Figure 1). All beliefs of the robots about actions facilitating goals as well as the ambition levels for those goals were set to 0. As can be seen in Table 4, this parameter setting led all robots to have a level of involvement of 0.12, and a level of distance 0.1 towards each other. Because the robots did not have any goals or beliefs about goals, the expected utilities of all possible actions were 0, and therefore their use intentions towards each other were also 0. Because all robots were exactly the same, and had a very low involvement, distance, and use intentions with respect to each other, they all had the same low (0.09) expected satisfaction of interacting with each other. The expected satisfaction of the actions to perform towards the other robots was 0.39 for fighting, 0.41 for comforting, 0.47 for avoiding, and 0.08 for criticizing. This resulted in all robots avoiding each other, as they had the highest expected satisfaction for performing this action.

Table 4. Simulation results of the baseline condition and the meaning of all abbreviations in the tables in this paper

	All Other Agents
All Agents	I= 0.12 D = 0.1 UI = 0 ES = 0.09 ES of PA = 0.41 ES of NA = 0.39 ES of CH = 0.08 ES of AV = 0.47 Action = avoid

Meaning of Abbreviations
I = Involvement D = Distance UI = Use Intentions ES = Expected Satisfaction PA = Positive Approach NA = Negative Approach CH = Change AV = Avoid

Experiment 1: The effect of having a goal

In this experiment, the parameter settings are the same as in the baseline condition, except that now Harry had a strong ambition for the goal to reduce his anger (level of ambition with value = 1) and thought he could do this by fighting with Gary (belief with value=1). Because of this, Harry had an expected utility of 1 for fighting Gary, and generated an action tendency of 1 for this action, which caused Harry to have use intentions of 1 for Gary. The generated action tendency to fight Gary had a small increasing effect on his involvement with (0.12 → 0.17) and a bigger increasing effect on his distance (0.1 →0.25) towards Gary. Harry's expected satisfaction for fighting Gary increased greatly (0.39→0.63), while there were only minor changes in the expected satisfaction of the other possible actions. Although Harry did not feel very involved with or at a distance towards Gary, he primarily rationally chose to fight Gary to reach his goal of reducing his own anger. Note that in the Table 5, Harry only calculates expected satisfactions for actions to perform to Gary. This is, because Harry's expected satisfaction of interacting with Gary is higher than that for Barry, and therefore Barry is left out of consideration.

Table 5. Simulation results of Experiment 1

	Harry	Barry	Gary
Harry		I = 0.12 D = 0.10 UI = 0 ES = 0.09 EU actions = 0	I = 0.17 D = 0.25 UI = 1 ES = 0.38 EU of fight =1 ES PA = 0.37 ES NA = 0.63 ES CH = 0.14 ES AV = 0.49 Action = fight
Barry	I = 0.12 D = 0.10 UI = 0 ES = 0.09 EU actions = 0 ES PA = 0.41 ES NA = 0.39 ES CH = 0.08 ES AV = 0.47 Action = avoid		I = 0.12 D = 0.10 UI = 0 ES = 0.09 EU actions = 0 ES PA = 0.41 ES NA = 0.39 ES CH = 0.08 ES AV = 0.47 Action = avoid
Gary	I = 0.12 D = 0.10 UI = 0 ES = 0.09 EU actions = 0 ES PA = 0.41 ES NA = 0.39 ES CH = 0.08 ES AV = 0.47 Action = avoid	I = 0.12 D = 0.10 UI = 0 ES = 0.09 EU actions = 0 ES PA = 0.41 ES NA = 0.39 ES CH =0.08 ES AV = 0.47 Action = avoid	

Experiment 2: The effect of being involved

In this experiment, compared to the baseline, the robot Gary was designed to be a beautiful, good, and realistic character (the designed values for these three parameters are set to 1). Because of this, the other agents had a much higher involvement with (0.12→0.49), and a somewhat lower distance (0.10→0.07) towards Gary. The expected satisfaction of the actions for the other agents to perform to Gary were influenced by these changes in involvement and distance towards him. It had a facilitating effect on the expected satisfaction of comforting

Table 6. Simulation results of Experiment 2

	Harry	Barry	Gary
Harry		I = 0.12 D = 0.10 UI = 0 ES = 0.09 EU actions = 0	I = 0.49 D = 0 UI = 0 ES = 0.29 EU actions = 0 ES PA = 0.60 ES NA = 0.20 ES CH = 0.20 ES AV = 0.25 Action = Comfort
Barry	I = 0.12 D = 0.10 UI = 0 ES = 0.09 EU actions = 0		I = 0.49 D = 0. UI = 0 ES = 0.29 EU actions = 0 ES PA = 0.60 ES NA = 0.20 ES CH = 0.20 ES AV = 0.25 Action = Comfort
Gary	I = 0.07 D = 015 UI = 0 ES = 0.10 EU actions = 0 ES PA = 0.37 ES NA = 0.43 ES CH = 0.07 ES AV = 0.51 Action = Avoid	I = 0.07 D = 0.15 UI = 0 ES = 0.10 EU actions = 0 ES PA = 0.37 ES NA = 0.43 ES CH = 0.07 ES AV = 0.51 Action = Avoid	

Gary (0.41→0.60) and criticizing him (0.08→0.20). It had an inhibiting effect on fighting Gary (0.39→0.20) and avoiding him (0.47→0.25). This resulted in Harry and Barry comforting Gary instead of avoiding him.

Experiment 3: Having a goal and being involved

In Experiment 1, Harry wanted to reduce his anger and thought he could do this by releasing his anger on Gary and fight him. In experiment 3, however, Gary was designed to be beautiful, good, and realistic (the designed values for these three parameters are set to 1), which made Harry be very involved (0.17→0.54) with Gary and less distant (0.25→0.15). This decreased his expected satisfaction of fighting (0.63→0.44) and avoiding (0.49→0.27) Gary and increased his expected satisfaction of comforting (0.37→0.56) and criticizing (0.14→0.26) him. Because Harry was too involved with Gary and had too little distance to fight him, he chose to comfort him instead, although he did not believe this would help him achieve his goals. The expected utility for Harry to fight Barry was 1, whereas all other expected utilities were 0, so that rationally Harry would choose to fight Barry. However, due to other factors, Harry was involved with Barry, which caused him to make an affective decision and comfort Barry.

Table 7. Simulation results of Experiment 3

	Harry	Barry	Gary
Harry		I = 0.12 D = 0.10 UI = 0 ES = 0.09 EU actions = 0	I = 0.54 D = 0.15 UI = 1 ES = 0.55 EU fight = 1 ES of PA = 0.56 ES of NA = 0.44 ES of CH = 0.26 ES of AV = 0.27 Action = Comfort
Barry	I = 0.12 D = 0.10 UI = 0 ES = 0.09 EU actions = 0		I = 0.49 D = 0 UI = 0 ES = 0.29 EU actions = 0 ES PA = 0.60 ES NA = 0.20 ES CH = 0.20 ES AV = 0.25 Action = Comfort
Gary	I = 0.07 D = 0.15 UI = 0 ES = 0.10 EU actions = 0 ES PA = 0.37 ES NA = 0.43 ES CH = 0.07 ES AV = 0.51 Action = Avoid	I = 0.07 D = 0.15 UI = 0 ES = 0.10 EU actions = 0 ES PA = 0.37 ES NA = 0.43 ES CH = 0.07 ES AV = 0.51 Action = Avoid	

4. Discussion

In this paper, the computational I-PEFiC model [3] was extended with goal-directed judgment formation [12] and overt actions. This way, the robots had the availability over goals to establish beliefs about how certain actions would sustain or obstruct them.

Simulation experiments were performed to test the behavior of this model. The hypothesis was that users or agents were means to achieve robot goals and that the calculated valence to those means would make the robot select one of possibly five actions: positive approach, negative approach, avoid, change, or do nothing. Our experiments showed that indeed this was the case and that robots could combine rational decisions with affective decisions.

In experiment 1, Harry primarily chose rationally to fight Gary, because he believed this would reduce his anger. In experiment 3, Harry held the same belief and by looking at the expected utilities of his actions alone, he should have chosen to fight Gary, but he did not. Because Harry was strongly involved with Gary, Harry suppressed the aggressive tendency of kicking Gary and comforted him instead, an affective choice.

There has been a continuous debate over the conflict between rational and emotional decision making. Most models of decision-making assume the process to be rational, which would exclude the possibility of emotion playing a role, other than of hindrance [10].

However, humans often make irrational decisions. A good example for this is the Ultimatum game [14]. In this game the players have to split a sum of money. The first player makes an offer how to split the money, and the second player can either accept or reject this offer. When the second player rejects, none of the players receives any money. In this game, the rational decision for the second player would be to accept any offer, as some money is always better than no money. In reality, behavioral research has shown that low offers (20% of total amount) have a 50% chance of being rejected. Based on participant reports, they rejected low offers because of anger (negative emotion), felt due to unfairness [10]. This nicely shows the existing conflicts between rational and emotional influences.

The somatic marker hypothesis [1] indicates that emotions have a strong influence over our decision-making abilities. According to this hypothesis, feelings (somatic markers) mark response options to real or simulated decisions. These somatic markers are then used in an automatic process to quickly select good decision options.

Existing models of decision making, such as [e, f], usually have a hedonic bias, and generally try to find the action with the highest expected utility. Some decision theoretic models take emotions into account, but in those models emotions only have the function to assist in making good rational decisions; e.g., emotional states are viewed as modes of decision making [8]. However, these models cannot explain irrational behavior, where actions with a (relatively) small expected utility are chosen. The model presented in this paper takes the expected utility, as well as involvement-distance-tradeoffs into account, according to the I-PEFiC model [15]. This way, the model takes rational as well as emotional influences into account, and not only rational behavior, but also behavior where emotional influences override the rational choice can be explained and simulated.

In future research, we will combine the model presented in this paper with an existing computational model for emotion regulation [5]. Whereas the current model focuses on the *elicitation* of emotions and affective decision making, that model addresses the *regulation* of emotion. We expect that both models will smoothly fit together, since the affective decision making as described in this paper could also be applied to emotion regulation strategies such as situation selection, situation modification, and attentional deployment. Finally, in a later stage of the project, the formalization will be validated against empirical data of human affective trade-off processes. As soon as the model has been validated and adapted, we will start exploring the possibilities to apply it to real humans instead of agents; i.e., 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.

Acknowledgements

We kindly want to thank Jonathan Gratch and Tibor Bosse for their input to this paper.

References

[1] Bechara, A., and Damasio, A.R. "The somatic marker hypothesis: a neural theory of economic decision" In: *Games and Economic Behavior*, volume 52, 2005, pp. 336-372.

[2] Beck, A.T., "Cognitive models of depression" In: *Journal of Cognitive Psychotherapy, an International Quarterly*, volume 1, 1987, pp. 5-37.

[3] Bosse, T., Hoorn, J.F., Pontier, M., and Siddiqui, G.F., "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*, 2008. To appear.

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

[5] Bosse, T., Pontier, M., and Treur, J. "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, 2007, pp. 187-192.

[6] Endsley, M. "Toward a Theory of Situation Awareness in Dynamic Systems" In: *Human Factors*, vol. 37, 1995, pp. 32-64.

[7] Frijda, N.H. "The emotions" New York, Cambridge University, 1986.

[8] Gmytrasiewycz, P.J., and Lisetti, C.L. "Emotions and Personality in Agent Design and Modeling" In: *M. Bauer, P.J.Gmytrasiewicz, and J. Vaassileva(Eds.), Springer-Verlag Berlin Heidelberg* , 2001, PP. 237-239.

[9] Gratch, J., and Marsella, S. "Evaluating the modeling and use of emotions in virtual humans" In: *Jennings, N.R., Sierra, C., Sonenberg, L., and Tambe, M. (eds.), Proceedings of the Third International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS 2004)*, ACM Press, New York, 2004, pp. 320-327.

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

[11] Hoorn, J.F. "A Robot's Experience of its User: Theory" In: *Sloutsky, V., Love, B.C., and McRae, K. (eds.), Proceedings of the 30th International Annual Conference of the Cognitive Science Society, CogSci'08*, 2008. To appear

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

[13] Marsella, S., and Gratch, J., "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, 2003, pp. 313-320.

[14] Thaler, R.H. "The Ultimatum Game." In: *Journal of Economic Perspectives*, volume 2, 1988, pp. 195-206.

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

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

[17] http://www.few.vu.nl/~ghazanfa/IAT-2008.pdf.