

An Affective Agent Playing Tic-Tac-Toe as Part of a Healing Environment

Matthijs Pontier^{1,2} and Ghazanfar Farooq Siddiqui^{1,2,3}

¹ VU University Amsterdam, Center for Advanced Media Research Amsterdam, De Boelelaan 1083, 1081HV Amsterdam, The Netherlands

² VU University Amsterdam, Department of Artificial Intelligence, De Boelelaan 1083, 1081HV Amsterdam, The Netherlands

³ Quaid-i-Azam University Islamabad, 45320, Pakistan.
{mpr210, ghazanfa}@few.vu.nl
<http://www.few.vu.nl/~{mpr210, ghazanfa}>

Abstract. There is a growing belief that the environment plays an important role in the healing process of patients, supported by empirical findings. Previous research showed that psychological stress caused by loneliness can be reduced by artificial companions. As a pilot application for this purpose, this paper presents an affective agent playing tic-tac-toe with the user. Experimenting with a number of agents under different parameter settings shows the agent is able to show human-like emotional behavior, and can make decisions based on rationality as well as on affective influences. After discussing the application with clinical experts and making improvements where needed, the application can be tested in a clinical setting in future research.

Keywords: Cognitive Modeling, Emotion Modeling, Healing Environment

1 Introduction

Many people do not like the atmosphere in hospitals. Since two decades, there is a growing belief that not only the health care itself, but also the environment affects the healing process of the patients. This has increased the interest in healing environments. The role of the environment in the healing process is a growing concern among health care providers, environmental psychologists, consultants and architects. Among them the consensus is growing that not only the level of care, but also the design of the health care facility affects the wellness of its patients [7].

Researchers are finding that making changes and additions to the physical and social environment of the health care facility, thereby taking the patient into account, can positively influence patients' outcomes (e.g., [2], [6], [18], [22], [25]). Moreover, health care professionals are finding that changes in design can enhance recovery in patients, and reduce the length of their stay in the hospital [15]. On the other hand, researchers are also finding that unfamiliar environments in clinics, hospitals, and

nursing homes can produce psychological stress that can negatively affect healing and wellness. Poor design has even been linked to negative effects on the patient, such as anxiety, delirium, elevated blood pressure levels, and an increased intake of pain drugs [23].

One factor that can be reduced by a healing environment is psychosocial stress. An important predictor of psychosocial stress is loneliness [12]. Loneliness is a common problem frequently encountered in the elderly in long-term care facilities. Many people that are staying in a long-term care facility lack social interaction. Artificial toys can be used to reduce loneliness. Previous research showed that animal-shaped toys can be useful as a tool for occupational therapy (e.g., [18], [26], [27]). Robot animal therapy has been widely investigated. For example, Dautenhahn and Robins [20], [28] used mobile robots and dolls respectively to treat autistic children. Wada and Shibata developed Paro [27], a robot shaped like a baby-seal that interacts with users to encourage positive mental effects. Interaction with Paro has been shown to improve users' moods, making them more active and communicative with each other and caregivers. Research groups have used Paro for therapy at eldercare facilities and with those having Alzheimer's disease [14], [17]. Banks et al. [2] showed that animal-assisted therapy with an AIBO dog helped just as good for reducing loneliness as therapy with a living dog. In their paper they indicate that AIBO was not used to its full capacity and that if more options were used, its effects might be further enhanced. Over the past decade, a lot of novel work on computational models of emotion in virtual agents can be observed. Nevertheless, compared to human affective complexity, current emotion models of virtual agents are still quite simple. If an artificial companion demonstrates human-like emotional behavior, this might increase its ability to reduce loneliness of patients in a long-term care facility, as part of a healing environment.

In our paper, we present a virtual agent that could be seen as a pilot application for this purpose. The artificial companion is an affective virtual agent that can play tic-tac-toe, equipped with Silicon Coppélia [19], an integration of three affect-related models as proposed in [3]. Because it is equipped with these affect-related models, it can show human-like emotional behavior. Therefore, it might be a useful to serve as an artificial interaction partner for patients in a long-term care facility.

2 The Application

The application presented in this paper is an affective virtual agent that can play tic-tac-toe against the user. The object of tic-tac-toe is to get three in a row on a three by three game board. You play on a three by three game board. Players alternate placing X's and O's on the game board until one of the players has three in a row, or all nine squares on the board are filled, which means the game ends in a tie. For creating the virtual agent, we used Haptik's peopleputty software [11]. Through this program we created the face of the virtual agent. The agent simulates 5 emotions: joy, distress, hope, fear and surprise, which can be expressed with either a low or a high intensity.

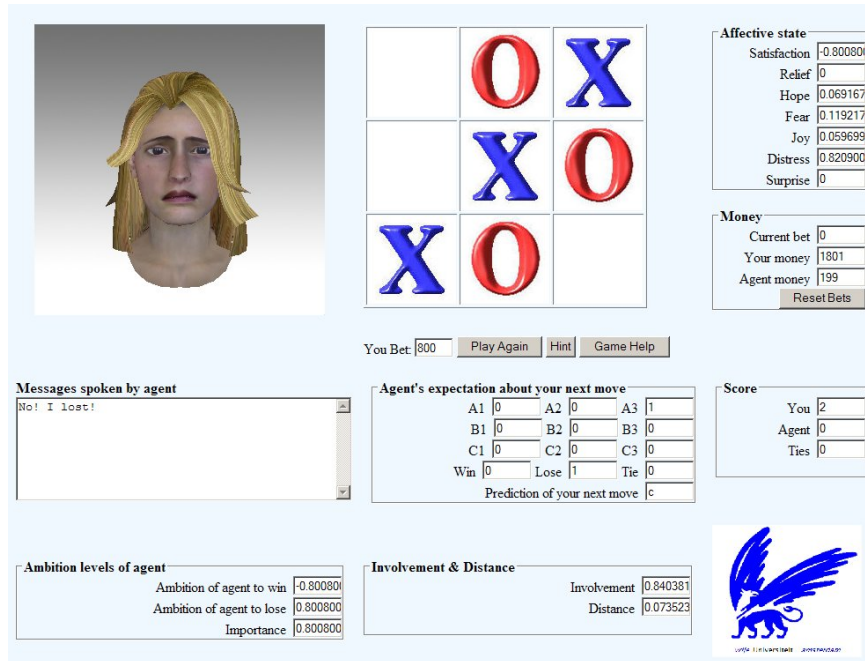


Fig. 1. The website with the tic-tac-toe application.

We created 32 (2^5) different emotional states using peopleputty; one for each possible combination of two levels of intensity of the five emotions simulated by the agent. We created a webpage for the application, on which the virtual agent was embedded as a Haptik player. We used JavaScript [1], a scripting language, in combination with scripting commands provided by the Haptik software [11], to control the Haptik player within a web browser.

Figure 1 shows the resulting website. In Figure 1, the agent, playing O's, just lost a game, and therefore looks sad. The website shows in the top left the agent which the user plays against. If the agent speaks a message, this is additionally shown in a textbox that is shown below the agent. Below this textbox, the ambition levels of the agent for winning and losing, and the importance of the current game for the agent are shown. Right next to this, the level of involvement and distance from the agent towards the user are shown. Right next to the textbox, the predictions of the agent about the expected next move and the outcome of the game are shown. Above this, just below the tic-tac-toe board, the user can enter its bet, and there is a 'play again' button which the user can click on to play a new game with the inserted amount of money as bet. Additionally, there is a hint button, which makes the agent give a hint to the user about in which square to make a move. Further, there is a 'game help' button. If the user clicks this button, the agent will explain the rules of the game.

Because the purpose of this paper is to show how the application works, the affective state is not only shown by means of a facial expression of the agent, but the emotion variables are also shown numerically on the top right. If the application would be

designed to be used by human users, these numerical values would not be shown, and only the facial expression would be visible to the user. Below the affective state in numbers, the amount of money that is currently played for (current bet), the amount of money of the user (your money), and the amount of money of the agent (agent money) are shown. There is also a ‘reset bets’ button, which resets the bets to the starting values. Below this, the number of games won by the user, the number of games won by the agent, and the number of ties is shown.

The tic-tac-toe board is on the top center of the website, right next to the agent. The user can make a move by clicking on one of the squares, on which the agent will react by performing its own move. After each move of the human user, the agent speaks a message, which is additionally displayed in the text-area below its face, depending on the emotional state of the agent. If the game has finished, the amount of money bet for will be added to the winner, and subtracted from the loser. The agent speaks a message, depending on the outcome of the game, and its emotional state. The user can enter a new bet and click the ‘play again’ button to play another game.

2.1 The models incorporated in the agent

This virtual agent presented in this paper was constructed by incorporating Silicon Coppélia [19], an integration of three affect-related models into an existing virtual agent that can play tic-tac-toe [29]. The three models that were integrated into Silicon Coppélia as suggested in [3] were:

1. EMA [9], [16], a model to create agents that exhibit and cope with (negative) affect based on Smith & Lazarus’ theory of emotion [21]. A graphical representation of EMA is shown in Figure 2.

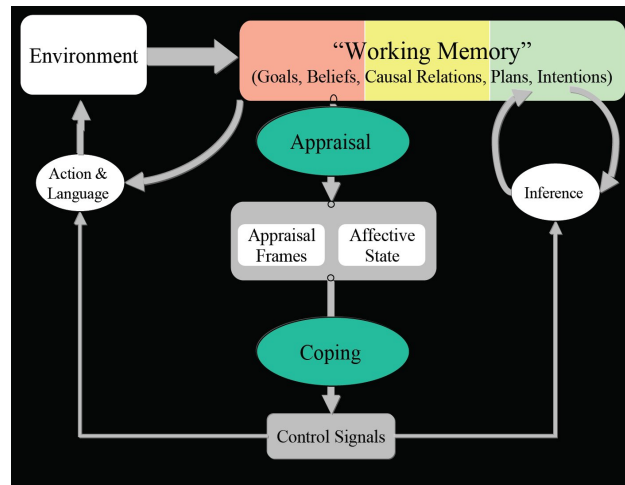


Fig. 2. A graphical representation of EMA

- CoMERG [4] (the Cognitive Model for Emotion Regulation based on Gross), which can simulate different emotion regulation strategies explained by Gross [10] using a set of logical rules and difference equations. Figure 3 shows a graphical representation of the emotion regulation model by Gross.

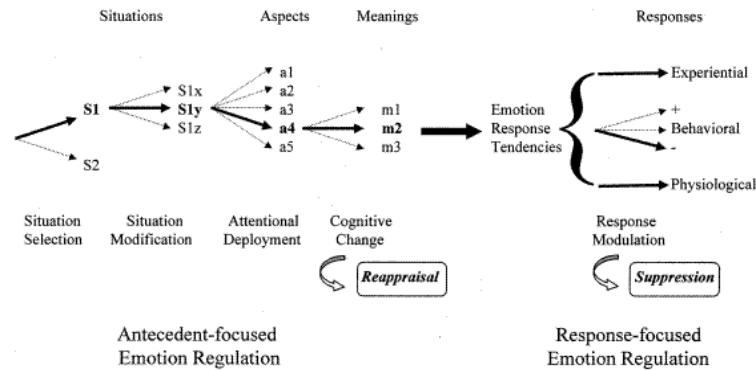


Fig. 3. A graphical representation of the emotion regulation model by Gross

- I-PEFiC^{ADM} [13], a model for building agents that can trade rational for affective choices based on the concern-driven theory of Frijda [8]. A graphical representation of I-PEFiC^{ADM} is shown in Figure 4.

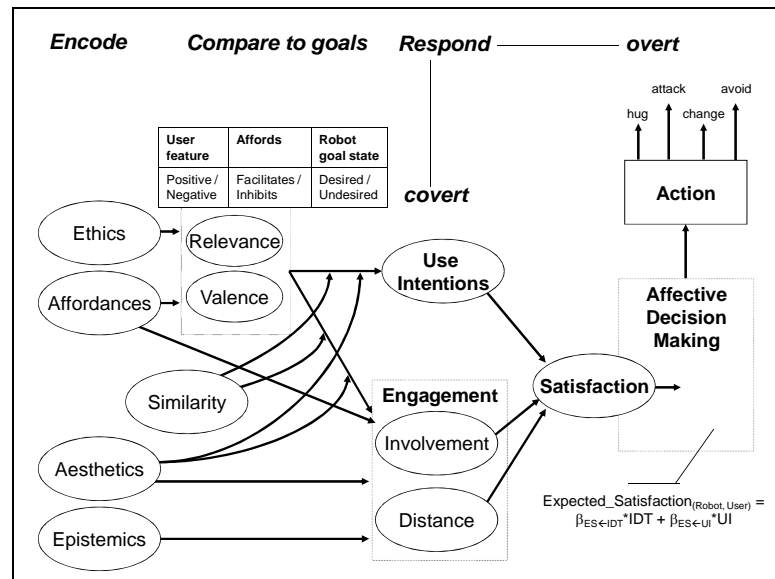


Fig. 4. A graphical representation of I-PEFiC^{ADM}. Curved arrows indicate interaction effects.

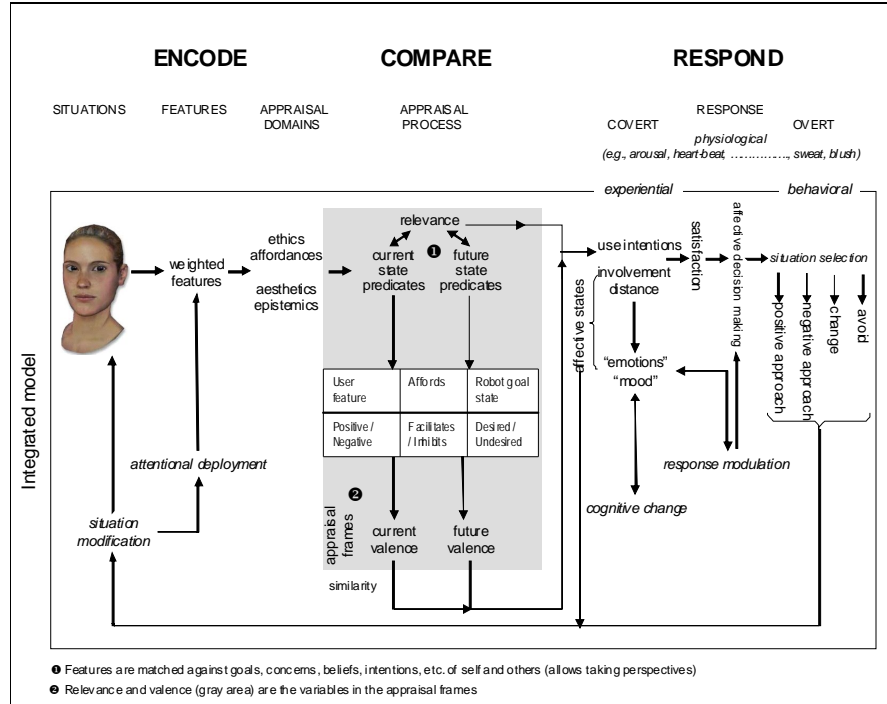


Fig. 5. A graphical representation of Silicon Coppélia.

Integrating these models enabled agents in simulation experiments to show richer interaction than they could with any of the models alone. Using the combined model, they could simulate emotions based on beliefs about states in the world, and how these states affect their goals. The model was also used to simulate affective decision making processes, in which decisions are made not only based on rationality, but also on affective influences, enabling the agent to make irrational decisions where appropriate. Further, emotion regulation strategies can be applied by the agent, to regulate its (simulated) emotions [19].

Figure 5 shows a graphical representation of Silicon Coppélia. On the far left of this figure, we see a virtual agent. The agent develops state predicates about her opponent. The agent acquires personal meaning for her opponent because she compares the features of her opponent with her own personal goals, beliefs, and concerns. This establishes her relevance and valence to her opponent. While relevance determines the intensity of affect, valence governs its direction. The agent can also take perspectives and look at others through the eyes of another person or agent.

When the initial appraisal process is completed, the agent is ready to affectively respond to her opponent. Relevance and valence form an appraisal frame that feeds into her involvement and distance towards her opponent. Inside, the agent will ‘experience’ several (perhaps ambiguous) emotions.

During affective decision making, the agent selects the move in the game that promises the highest expected satisfaction. The performed action leads to a new situation, and after her opponent also made his/her move, the model loops until the game has finished.

2.2 Determining which action to take

In Silicon Coppélia [19], the agents perceive each other's features according to I-PEFiC [24], by multiplying a designed value (a value the designer expects to raise in another agent), and a bias for perceiving this feature. For this application, as starting values, all biases were set to the neutral value of 1 for the agent. The designed values of the human user were set as can be found in Table 1. These values were chosen arbitrarily, and in an application with real users they should be reconsidered together with an expert. This will lead the agent to perceive the human user with the following variables:

Table 1. Designed values for perceiving the features of the human user.

Feature	Designed value of human user
Good	0.6
Bad	0.2
Beautiful	0.5
Ugly	0.3
Realistic	0.9
Unrealistic	0.1
Intended Aid	0.7
Intended Obstacle	0.4

In Silicon Coppélia, the agents also have beliefs that features of other agents affect certain goal-states in the world. For this application, the possible goal-states are 'agent wins' and 'human wins'. The beliefs about features facilitating these goal-states were set as can be seen in Table 2:

Table 2. Designed values for perceiving the features of the human user.

Feature or Action	Facilitates 'agent wins'	Facilitates 'human wins'
Good	0.5	-0.5
Bad	-0.5	0.5
Beautiful	0.1	0.1
Ugly	0.1	0.1
Realistic	0.1	0.1
Unrealistic	0.1	0.1
Intended Aid	0.7	-0.7
Intended Obstacle	-0.7	0.7

The agent also has beliefs about actions facilitating goal-states. As there are nine squares on the board, there are nine possible actions for the agent: putting an O in

each possible place. In Silicon Coppélia, the agents have beliefs that action facilitate goal-states in the range [-1, 1], where -1 means the agent believes the action strongly inhibits the goal-state, and 1 means the agent believes the action strongly facilitates the goal-state.

Each turn, the beliefs that actions facilitate goal-states are calculated based on a heuristic that estimates the chances of winning when a certain move is made. If the agent can make three O's in a row, the belief that performing this action facilitates 'agent wins' is set to 1, and the belief that this facilitates 'human wins' is set to -1. If the human user has two X's in a row, and the third square in this row is still empty, the agent will believe that putting an O on this square facilitates 'agent wins' with a value of 0.9, and inhibits 'human wins' with a value of -0.9. If a square is already occupied, the belief that putting an O on this square facilitates a goal is set to 0 for both goals, as it is an illegal move. If none of these rules apply for a square, the belief that putting an O on the middle square (B2) facilitates 'agent wins' is set to 0.5, and that it facilitates 'human wins' to -0.5. Similarly, the belief that putting an O on corner squares (A1, A3, C1 and C3) facilitates 'agent wins' is set to 0.3 and for 'human wins' this is set to -0.3. Finally, the belief that putting an O on the remaining squares (A2, B1, B3 and C2) facilitates 'agent wins' is set to 0.1 and for 'human wins' this is set to -0.1. For actions which facilitate desired goals, and inhibit undesired goals (i.e., actions with a high expected utility, generalized over multiple goals), strong action tendencies are calculated [19].

In Silicon Coppélia, each action has a level of positivity and a level of negativity. In this application, the level of negativity of each action is defined as the belief of the agent that the action facilitates winning the game. The level of negativity of the action is defined as 1 minus this belief.

Using these variables, the affective decision-making of Silicon Coppélia [19] is used to determine the action of the agent. The expected satisfaction is calculated using the following formula, and the action with the highest level of expected satisfaction is picked.

$$\begin{aligned} \text{ExpectedSatisfaction}(\text{Action}) = & \\ & w_{at} * \text{Action_Tendency} + \\ & w_{pos} * (1 - \text{abs}(\text{positivity} - \text{bias}_i * \text{Involvement})) + \\ & w_{neg} * (1 - \text{abs}(\text{negativity} - \text{bias}_o * \text{Distance})) \end{aligned}$$

The agent will search for the action with the level of positivity that is closest to the level of (biased) involvement towards the user, the level of negativity closest to the level of (biased) distance towards the user, and the strongest action tendency. The importance of positivity, negativity and action tendency for selecting an action can be adjusted by changing the weights (w_{pos} , w_{neg} , and w_{at} respectively). If an agent wants to perform more positive actions, it can, for example, increase its bias for involvement, and decrease its bias for distance. This way, the agent will prefer more positive and less negative actions.

Note that this way, the agent can also deliberately lose by setting a high ambition level for the goal 'human wins', or because the agent is too involved with the user to try to win each game. The agent can also determine its ambition level for winning on the outcomes of the previous games, and the amount of money that is played for. For example, it can try to win the game if it has less money than the human user, and

deliberately try to lose if it has more money than the human user. The agent can also determine the importance of each game, by dividing the amount that is played for by the total amount of money the agent has left. This importance can then be the deviation from 0 in the ambition level.

2.3 Calculating the emotions of the agent

The agent simulates some emotions while playing the game, based on the actions that are being performed by the user, the perceived likelihood of winning and losing, and the outcome of the game. Hope and fear are calculated each time the agent has made its move and the human user is on turn. The hope and fear of the agent are based on the perceived likelihood it will win or lose the game. If the human user is on turn and can make a winning move, it estimates the likelihood for losing 0.8, the likelihood for a tie 0.1, and the likelihood for winning 0.1. If the agent could make a winning move if it would be on turn (but it cannot, because the human user is), it will estimate its likelihood for winning 0.5, the likelihood for a tie 0.4 and its likelihood for losing 0.1. Otherwise, the likelihood for winning and losing are both estimated 0.3 by the agent and the likelihood for a tie is estimated 0.4.

The found likelihoods are used in the following function to calculate the hope for a goal. This function is similar to the function described in [5].

IF $f \geq \text{likelihood}$ THEN $\text{hope_for_goal} =$
 $-0.25 * (\cos(1 / f * \pi * \text{likelihood}(\text{goal})) - 1.5) * \text{ambition_level}(\text{goal})$

IF $f < \text{likelihood}$ THEN $\text{hope_for_goal} =$
 $-0.25 * (\cos(1 / (1-f) * \pi * (1-\text{likelihood}(\text{goal}))) - 1.5) * \text{ambition_level}(\text{goal})$

These functions differ from most approaches present in the literature, since their top is not situated at the point where the likelihood is 0. In these functions, f is a shaping parameter (in the domain $[0, 1]$) that can be used to manipulate the location of the top of the hope curve. The value of this parameter may differ per individual, and represents 'fatalism' (or 'pessimism'): the top of the likelihood/hope-curve is always situated at the point where $\text{likelihood} = f$. Thus, for an f close to 1, the top of the curve is situated to the extreme right (representing persons that only 'dare' to hope for events with high probabilities). Similarly, for an f close to 0, the top of the curve is situated to the extreme left (representing persons that already start hoping for events with low probabilities). In this paper, f is set at 0.5. We chose a smooth function instead of a linear function, because this function best matches the emotion curves found in humans. Furthermore, a higher ambition level simply leads to a higher hope (which is standard in the literature). If the ambition level is negative (i.e., the goal is undesired), the outcome of hope_for_goal will be a negative value.

The following algorithm is performed to the found values for hope_for_goal

1. Sort the values in two lists: $[0 \rightarrow 1]$ and $[0 \rightarrow -1]$
2. Start with 0 and take the mean of the value you have and the next value in the list. Continue until the list is finished. Do this for both the negative and the positive list.
3. Hope = Outcome positive list. Fear = $\text{abs}(\text{Outcome negative list})$.

The values are sorted in a list with positive `hope_for_goal`'s (i.e., hope for desired goals), and negative `hope_for_goal`'s (i.e., fear for undesired goals). For both the lists, 0 is the starting point, and the mean of the value you have and the next value in the list (where the next value is the value closest to 0 that is left in the list) is picked until the end of the list is reached. The new level of hope for the agent is the outcome of the positive list, and the new level of fear for the agent is the absolute value of the outcome of the negative list.

The joy and distress of the agent are based on reaching or not reaching desired or undesired goal-states. If a goal-state becomes true (i.e., the agent wins or the human user wins), the levels of joy and distress are calculated by performing the following formulas:

```
IF ambition_level(goal) >= 0 THEN:
    new_joy      = old_joy + mf_joy * ambition_level(goal) * (1-old_joy)
    new_distress = old_distress + mf_distress * -ambition_level(goal) * old_distress
```

```
IF ambition_level(goal) < 0 THEN:
    new_joy      = old_joy + mf_joy * ambition_level(goal) * old_joy
    new_distress = old_distress + mf_distress * -ambition_level(goal) * (1-old_distress)
```

In this formula, `mf_joy` and `mf_distress` are modification factors that determine how quickly joy and distress change if the agent wins or loses the game. In this paper, the values were both set at 1. These modification factors are multiplied with the impact value, which is `ambition_level(goal)` for joy and `-ambition_level(goal)` for distress. This way, if a desired goal is reached, this will increase joy and decrease distress, and reaching an undesired goal will decrease joy and increase distress. Multiplying with limiter `(1-old_joy)` for joy and `old_distress` for distress if the goal is desired manages the formula does not go out of range. Further, it manages that if an agent's level of joy or distress approaches an extreme value, it will be harder to push it further to the extreme, and easier to get it back to a less extreme value. If the reached goal is undesired, `old_joy` is used as limiter for joy and `(1-old_distress)` as limiter for distress, because the values of joy and distress will move in the opposite direction as when the goal is desired.

The level of surprise is calculated in a similar manner as in [5], [29]. To calculate the level of surprise during the game, the agent generates expectations about which move the user will make. If a square is free, that square gets a point. If the user can make three in a row, or prevent the agent from making three in a row the next turn, the square of that move gets 1 extra point. If the user has one X in a row, and the remaining squares of that row are free, those squares get 0.5 point. After all the squares have got their points, the sum of all points of the squares is normalized to 1. The resulting values for each square are the predicted likelihoods of the human making a move on that square. If the user makes a move on a certain square, the level of surprise for the agent is $1 - \text{likelihood}(\text{move})$. If the game finishes, the level of surprise for the agent is 1 minus the perceived likelihood that the game would end that way.

After each move of the human user, the agent speaks a message, depending on the level of surprise. In the system, there is a small database of messages, labeled with

certain emotion intensities. If the level of surprise is very low, it will show that it expected this move, with a message like 'I thought you would do that' or 'A predictable move'. On the other hand, if the agent is very surprised by the move of the user, it will speak a more surprised message, like 'That move surprised me!' or 'You are full of surprises.'

All five emotions inserted in the system (joy, distress, hope, fear and surprise) are simulated in parallel. If the level of joy, distress or surprise is below 0.5, a low intensity of the emotion is shown by the agent. If the level of joy, distress, or surprise is greater or equal than 0.5, a high intensity of the emotion is shown by the agent. Because playing the tic-tac-toe game rarely leads to extreme values of hope and fear in the agent, for hope and fear this boundary is set to 0.25.

After each game has ended, the level of satisfaction for the agent is calculated in the range [-1, 1]. If the agent wins, the level of satisfaction will be the ambition level for 'agent wins'. Similarly, if the human user wins, the level of satisfaction will be the ambition level for 'human wins'. If the game ends in a tie, the importance of the game is calculated by dividing the amount that is played for by the total amount of money of the agent. The satisfaction of the agent after a tied game is then calculated by multiplying the importance of the game with 0.5.

Also a level of relief is calculated for the agent after each game, in the range [-1, 1], by multiplying the level of satisfaction with the level of surprise. Further, a message is displayed, based on the outcome of the game and the level of satisfaction and relief, the agent speaks a message in a similar way as after each move. If the agent wins, and the level of relief is low, and the level of satisfaction is low, it will display a neutral message like 'I win'. If the level of satisfaction is higher, it will display a more enthusiastic message like 'Superb match for me!' If the level of relief is higher, it speaks this relief with a message like 'I won, that is such a huge relief'. If the agent loses, and has a relatively neutral level of satisfaction close to 0, it will display a message like 'You won'. If the agent is very dissatisfied, with a value close to -1, it will speak a more dramatic message, like 'No! I lost everything!' If the game ends in a tie, the agent speaks a neutral message like 'We tied'. If the agent's ambition was to win, and it wins, it will show a happy facial expression. If it loses, it will look sad. If the agent's ambition was to lose, and it wins, it shows a sad facial expression. If its ambition was to lose and it loses, the facial expression will be happy.

3 Results

To test the application, the behavior of the agent has been tested under various parameter settings. All agents experimented with can be found at [30].

Agent 1: The agent tries to win

The ambition level for winning of agent 1 is set to 1, and its ambition level for losing is -1. The weight of the affective influences in the decision making process is set to 0. Under these parameter settings, the agent will always try to win. Because in tic-tac-toe it is impossible to lose if you play it right and you do not want to lose, it is

impossible to win of the agent. The best game outcome that can be achieved is a draw.

If the agent wins, it will increase its joy and decrease its distress. If the agent loses, it will decrease its joy and increase its distress. If you make a move the agent does not expect, or the game ends otherwise than expected, it will be surprised. The expectations of the agent can be seen on the website. If the agent thinks it is likely that it will win, it will have a relatively high level of hope, and if it thinks it is likely that it will lose, it will have a relatively high level of fear.

Agent 2: The agent deliberately tries to lose

Agent 2 has an ambition level for winning of -1, and an ambition level for losing of 1. This means for the agent, winning is an undesired goal, and losing is a desired goal. The weight of the affective influences in the decision making process have set to 0, so the agent will always try to lose. The only way to let the agent win is to make sure you don't make three in a row, and with the last move of the agent, it can do nothing else than make the winning move. Because the agent wants to lose, it will increase its level of joy and decrease its level of distress when it loses. If it wins, it will decrease its level of joy, and increase its level of distress.

Agent 3: The agent decides whether it wants to win based on its money

For agent 3, the ambition level for winning and losing is dependent on its amount of money compared to the amount of money of the human. If the agent has more money than the human, it will have an ambition level for winning of $-1 \cdot \text{importance}$ and an ambition level for losing of $1 \cdot \text{importance}$. If the amount of money of the agent is less or equal than that of the human, the agent will have an ambition level for winning of $1 \cdot \text{importance}$ and an ambition level for losing of $-1 \cdot \text{importance}$. The weight of the affective influences of the agent is set to 0. This causes the agent to try to win, unless it already has more money than the human user. If the agent has the ambition to win and it does, it will decrease its joy, and decrease its distress, and if it loses it will decrease its joy and increase its distress. However, if the agent has the ambition to lose, and it wins, it will decrease its joy and increase its distress, and vice versa if it loses. How big the increases and decreases of joy and distress are, depends on the importance of the game.

Agent 4: The agent is too involved with the user to win

Agent 4 is very involved with its user. It is programmed to perceive its user as good, beautiful, realistic, and intending to aid (designed values set to 1). It is also programmed to perceive its user as not bad, not ugly, not unrealistic, and not intending to obstruct. This causes the agent to be very involved with the user with a value of 0.85, and not much at a distance towards the user, with a value of 0.08 at the start of the simulation. The ambition level of the agent to win is defined as the

importance of the game, and the ambition level to lose as the negation of this importance. The weight of rational influences in the decision-making process is set to 0. This makes the agent want to perform actions towards the user with a high level of positivity, and a low level of negativity. Because actions to win the game have a relatively low level of positivity and a high level of negativity, the agent will perform actions that facilitate losing the game. Because this agent always has a positive ambition to win and a negative ambition to lose, winning will always increase its level of joy, and decrease its level of distress, and losing will always decrease its level of joy, and increase its level of distress. How big this increase or decrease is, is dependent on the importance of the game.

Agent 5: The agent is balanced, and wins sometimes, and loses sometimes

Agent 5 perceives the user with values as can be seen in Table 1. This leads the agent to be involved with the user with a value of 0.63, and to be at a distance with the user with a value of 0.26 at the start of the simulation. The ambition level of the agent to win is defined as the importance of the game, and the ambition level to lose as the negation of this importance. The rational influences in the decision-making process are set to 0.8, and the influences of positivity and negativity of action are both set to 0.1. Despite having a higher ambition to win than to lose under all circumstances, the agent will be too involved with the user to try to win in a game for a small amount of money. However, if the agent's money is almost gone, or the game is about a lot of money, the agent finds the game so important that it will do its best to try to win. Similarly to agent 4, if agent 5 wins, this will increase joy and decrease distress, and vice versa if it loses.

4 Discussion

This study presents an affective virtual agent that can play tic-tac-toe. Because it is equipped with Silicon Coppélia [19], an integration of three affect-related models as suggested in [3], it can show human-like emotional behavior.

We created five different agents, each with different parameter settings, to test the behavior under various conditions. We manipulated the ambition levels of the agent, and thereby created agent 1, an agent that always tries to win, and agent 2, an agent that always deliberately tries to lose. Agent 3 determines its ambition level for winning and losing on whether it has more money than the user or not. If the user has more money than the agent, it will deliberately try to lose, but otherwise it will try to win. Agent 4 bases its decisions in the game on emotions, and because it is designed to be very involved with the user, it will perform actions that facilitate the user winning the game. Agent 5 bases its decisions partly on emotions, and partly on rationality. Agent 5 always has more ambition to win than to lose. How big this difference in ambition is, is dependent on the amount of money that is played for. This results in the agent trying to win if the agent plays for a big amount of money, or when its money is almost gone. However, if the game is only about a small amount of money, the agent will be too involved with the user to try to win. Based on whether

the agent reaches its goals (winning and losing when the agent has ambitions to win or lose), the likelihood of these goals, and the expectedness of the move of the user and the outcome of a game, the emotions joy, distress, hope, fear and surprise are simulated and shown by the agent by means of facial expressions.

This virtual agent presented in this paper should be seen as a pilot application. Many improvements can still be made, such as giving feedback in a more sophisticated manner. Before it can be tested in a clinical setting, we should first discuss with experts where the application could be improved. We should also discuss with them what type of behavior the agent should show under which conditions, and adjust the parameter settings to meet these requirements. After that, user studies should indicate under which parameter settings participants find the agent most human-like.

However, experimenting with a number of agents, each with different parameter settings indicates that a realistic affective agent playing tic-tac-toe can be created.

Previous research already showed that interacting with a robot pet could decrease loneliness in patients staying in a long-term care facility. An artificial interaction partner that can show human-like emotional behavior might even have a greater beneficial effect on decreasing loneliness in patients. In future research, we intend to perform user studies to show whether this really is the case.

Acknowledgments

We would like to thank Tibor Bosse and Edwin Zwanenburg for their input to the paper, and Edwin Zwanenburg for implementing the game environment.

References

1. About JavaScript – MDC, http://developer.mozilla.org/en/docs/About_JavaScript
2. Banks, M.R., Willoughby, L.M., Banks, W.A.: Animal-Assisted Therapy and Loneliness in Nursing Homes: Use of Robotic versus Living Dogs. *Journal of the American Medical Directors Association*, vol. 9, 173--177 (2008)
3. Bosse, T., Gratch, J., Hoorn, J.F., Pontier, M.A., and Siddiqui, G.F. (submitted): Coppélius' concoction: Similarity and complementarity among three affect-related agent models. *International Conference on Agents and Artificial Intelligence, ICAART 2010* (2010)
4. Bosse, T., Pontier, M.A., 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*, 187--192, Taylor and Francis (2007)
5. Bosse, T., Zwanenburg, E.: There's Always Hope: Enhancing Agent Believability through Expectation-Based Emotions. In: Pantic, M., Nijholt, A., Cohn, J. (eds.), *Proceedings of the 2009 International Conference on Affective Computing and Intelligent Interaction, ACII'09*, 111—118, IEEE Computer Society Press (2009)
6. Davidson, A.W.: Banking on the Environment to Promote Human Well-being. In: Seidel, A.D. (Ed.), *Banking on design? Proceedings of the 25th annual conference of the Environmental Design Research Association*, 62--66. Oklahoma City, OK: EDRA. (1994)
7. Devlin, A.S., Arneill, A.B.: Health Care Environments and Patient Outcomes: A Review of the Literature. *Environment and Behavior*, vol. 35, no. 3, 665--694 (2003)

8. Frijda, N.H., *The Emotions*. Cambridge University, New York (1986)
9. Gratch, J., Marsella, S.: Evaluating a computational model of emotion. *Journal of Autonomous Agents and Multiagent Systems* (Special issue on the best of AAMAS 2004), vol. 11, no. 1, 23–43 (2006)
10. Gross, J.J.: Emotion Regulation in Adulthood: Timing is Everything. *Current directions in psychological science*, Vol. 10, No. 6, 214–219 (2001)
11. Hapttek, Inc., <http://www.hapttek.com>
12. Hawkli, L.C., Masi, C.M., Berry, J.D., Cacioppo, J.T.: Loneliness Is a Unique Predictor of Age-Related Differences in Systolic Blood Pressure. *Psychology and Aging*, vol. 21, no. 1, 152–164. DOI: 10.1037/0882-7974.21.1.152 (2006)
13. Hoorn, J.F., Pontier, M.A., Siddiqui, G.F.: When the user is instrumental to robot goals. First try: Agent uses agent. *Proceedings of IEEE/WIC/ACM Web Intelligence and Intelligent Agent Technology 2008, WI-IAT '08, IEEE/WIC/ACM, Sydney AU*, 296–301. ISBN: 978-0-7695-3496-1. DOI: 10.1109/WIIAT.2008.113 (2008)
14. Kidd, C., Taggart, W., and Turkle, S.: A Social Robot to Encourage Social Interaction among the Elderly. *Proceedings of IEEE ICRA*. pp. 3972–3976 (2006)
15. Lemprecht, B.: The gap between design and healing. *Metropolis*, vol. 77, 123 (1996)
16. Marsella, S., Gratch, J.: EMA: A Model of Emotional Dynamics. *Cognitive Systems Research*, vol. 10, no. 1, 70–90 (2009)
17. Marti, P. Bacigalupo, M., Giusti, L., Mennecozzi, C.: Socially Assistive Robotics in the Treatment of Behavioral and Psychological Symptoms of Dementia. *Proceedings of BioRob*, 438–488 (2006)
18. Nakajima, K., Nakamura, K., Yonemitsu, S., Oikawa, D., Ito, A., Higashi, Y., Fujimoto, T., Nambu, A., Tamura, T.: Animal-shaped toys as therapeutic tools for patients with severe dementia. *Engineering in Medicine and Biology Society, 2001: Proceedings of the 23rd Annual International Conference of the IEEE*, vol. 4, 3796–3798 (2001)
19. Pontier, M.A., Siddiqui, G.F.: Silicon Coppélia: Integrating three affect-related models for establishing richer agent interaction. *Proceedings of the International Conference on Intelligent Agent Technology, IAT'09*. To appear. (2009)
20. Robins, B., Dautenhahn, K., Boekhorst, R.T., Billard, A.: Robotic Assistants in Therapy and Education of Children with Autism: Can a Small Humanoid Robot Help Encourage Social Interaction Skills?. *Journal of Universal Access in the Information Society*. 4, 105–120 (2005)
21. Smith, C.A., Lazarus, R.S.: Emotion and Adaptation. In: L.A. Pervin (ed.), *Handbook of Personality: theory & research*, Guilford Press, NY, 609–637 (1990)
22. Ulrich, R.S.: View through a Window May Influence Recovery from Surgery. *Science*, vol. 224, 420–421 (1984)
23. Ulrich, R.S.: Effects of Interior Design on Wellness: Theory and Recent Scientific Research. *Journal of Health Care Interior Design*, Vol. 3, 97–109 (1991)
24. Van Vugt, H.C., Hoorn, J.F., Konijn, E.A.: Interactive engagement with embodied agents: An empirically validated framework. *Computer Animation and Virtual Worlds*, Vol. 20, 195–204 (2009)
25. Verderber, S., Reuman, D.: Windows, Views, and Health Status in Hospital Therapeutic Environments. *Journal of Architectural and Planning Research*, vol. 4, 120–133 (1987)
26. Wada, K., Shibata, T.: Living with Seal Robots in a Care House - Evaluations of Social and Physiological Influences. In: *Intelligent Robots and Systems, IEEE/RSJ International Conference on*, pp. 4940–4945 (2006)
27. Wada, K., Shibata, T.: Social Effects of Robot Therapy in a Care House -Change of Social Network of the Residents for One Year-. *JACIII*. 13, 386–392 (2009)
28. Werry, I. Dautenhahn, K.: Applying Mobile Robot Technology to the Rehabilitation of Autistic Children. In: *Proc. of 7th Int. Symp. on Intelligent Robotic Systems*, pp. 265–272 (1999)

29. Zwanenburg, E.: Enhancing Believability through Expectation-Based Emotions in a BDI Agent. Master Thesis, Vrije Universiteit Amsterdam, (2009)
30. <http://www.few.vu.nl/~ghazanfa/IVA2009/tictactoe.html>