Building embodied agents that experience and express emotions: A football supporter as an example

The Duy Bui Dirk Heylen Anton Nijholt
Department of Computer Science
University of Twente
{theduy,heylen,anijholt}@cs.utwente.nl

Abstract

In this paper we presented Obie, an embodied agent that experiences and expresses emotions. Obie has an adaptive, quantitative and domain-independent emotion component which appraises events to trigger emotions. Obie's emotions are expressed via his utterances or his facial expressions. The expression via utterances is done by a simple mapping from emotions to text fragments. The mapping from emotions to facial expressions is done by a fuzzy rule-based system. Obie's utterances and facial expressions are presented in his 3D talking head. In the research described in this paper, Obie was implemented as a football supporter agent. We show how Obie experiences different emotions during a football match. We also indicate how Obie with different personalities experiences emotions differently.

Keywords: social and conversational agents, emotions and personality

1 Introduction

Recent scientific findings show that emotions play an important role in human cognitive functions. As summarized by Picard [1] in her "Affective Computing", emotions affect creativity, evaluative judgement, rational decision making, communication, and other cognitive processes. This has also been supported by many other scientists including Gelernter [2], Forgas and Moylan [3], and Damasio [4].

Recognizing the importance of emotions to human cognitive functions, Picard [1] concluded that if we want computers to be genuinely intelligent, to adapt to us, and to interact naturally with us, then they will need the ability to recognize and express emotions, to model emotions, and to show what has come to be called "emotional intelligence".

As early as the 1930s, traditional character animators, in particular from Disney, have incorporated emotion into animated characters to make audiences "believe in characters, whose adventures and misfortunes make people laugh - and even cry" [5]. The animators believe that emotion, appropriately timed and clearly expressed, is one of the keys to creating the quality of animated films. In the areas of computational synthetic agents, emotions have received much attention for their influences in creating believable characters, e.g. [6, 7].

In this paper we present Obie, an embodied agent that experiences and expresses emotions. Obie is built based on several systems which we have presented before. These systems are: the 3D face model, which is presented in [8] and is developed into a talking head in [9]; ParleE, the implementation of emotions [10]; and the fuzzy rule-based system [11], which converts emotions into facial expressions. The core of Obie is ParleE, a quantitative, domain-independent and adaptive computational implementation of emotions for an embodied agent situated in a multi-agent environment. ParleE also allows Obie with different personalities to experience emotions differently.

In the research described in this paper, Obie was implemented as a football supporter agent. He is watching a football match in which a team which he supports is playing. Obie can experience different emotions by appraising events based on his goals, standards, and preferences. Obie can also shows his emotions on his 3D talking head. We consider three types of football match. The first type consists of the real football matches that take place in stadiums. The second type is the robot football matches where the human players are replaced by physical robots. The third type is the simulation matches where we have 2D or 3D virtual robots and the field only exists on the computer screen. Typical events that occur in a football match are: kick-off, penalty, goal, free-kick, etc. These events can be obtained in various ways. For a real football match, the events may be extracted directly by translation from visual to verbal representations or translation from a news stream produced by a mediator (e.g., a human commentator) to a textual representation [12]. For a robot cup match, the events can be extracted from a team's vision system [13]. The events in a simulation match can be extracted directly from the data of the match.

There are several reasons why we choose the domain of a football supporter to implement. First of all, football is an emotional game. There are many events in the the game that trigger emotions of not only players but also of coaches, supporters, etc. A last-minute goal triggers happiness or relief in some people whereas it triggers sadness, anger or disappointment in other people. Implementing the football supporter's domain gives us the chance to test many emotions as well as blends of emotions. Secondly, because the actions in a football match happen fast, the emotional state of a football supporter also changes fast during the match. This gives the implemented agent a chance to experience many emotions in a short period. Thirdly, a supporter in a football match can experience extreme emotions, which allows us to observe the expressed emotions more easily. Finally, the utterances of a supporter are usually short and

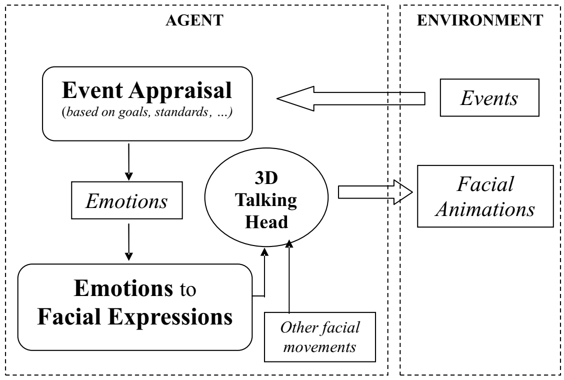

Figure 1: Obie's architecture.

simple. Therefore, we do not have to implement a very sophisticated text generation component in order to test the agent's ability to express emotions during speech. Our motivation to implement a football supporter agent is also inspired by two existing projects in our department: the robot soccer project [14] and the football commentary generation project [12].

The paper is organized as follows. First, we present Obie's architecture. Then, the football supporter's domain is discussed. Finally, an illustration of Obie's behavior is given.

2 Obie's architecture

Obie was built based on several systems that we have presented before [8, 9, 10, 11]. An overview of his architecture can be seen in Figure 1. Obie takes events as input. Then he appraises these events according to his goals, standards and preferences to generate emotions. Some basic emotions are mapped to emotional facial expressions. Next, facial expressions are displayed on Obie's 3D talking head. Obie's internal state is also shown in the form of charts and graphs as well as the content of the speech which Obie utters. The emotional expressions via the content of the speech is done by a simple mapping between emotions and text, e.g., the emotion fear is mapped with the sentence "Oh no!".

The event appraisal component Events are appraised by ParleE, a quantitative, flexible and adaptive computational implementation of emotions for an embodied agent situated in a multi-agent environment [10]. ParleE has been inspired by various other implementations of emotions such as Elliott's Affective Reasoner [6], Velásquez's Cathexis [15], El-Nasr et al's FLAME [16] and, in particular, by Gratch's Émile [7]. Like some of these and many other implementations, ParleE generates emotions mostly based on Ortony et al.'s theory of appraisal [17] (the emotion surprise is based on Roseman's model [18]). Nevertheless, ParleE possesses some significant properties of its own. The main novel differences with other systems are: (i) the way it uses forward-chaining search within a finite depth to obtain the probability of achieving a goal; (ii) the way it uses models of other agents' plans and goals to predict their behavior and set up expectations about the likelihood of events; and (iii) the way it incorporates personality, individual characteristics and motivational states in the implementation.

In ParleE, we consider agents in a [Markov Decision Processes]-style framework, in which the state of the world is known but actions are probabilistic. When there is more than one agent, the actions which other agents are going to perform are also unknown. Using the information from an agent's planner and a model of the other agent, ParleE calculates intensities of emotion based on the change in probability of achieving the agent's goal. This is derived from the view of Oatley and Johnson-Laird [19] and Reilly [20]. When exporting to a new application, ParleE reuses the agent's existing planner and information about the new world. Like Émile [7], it provides a domain-independent and flexible way of generating emotions. Therefore, ParleE just fits in any application with the same framework. However, different from Émile's threat detection approach, ParleE uses forward-chaining search within a finite depth to obtain the probability of achieving a goal. By doing this, Parle solves the problem of mistreating the event that is both establisher and threat to the agent's goal (cf. [10]). ParleE also provides a generic way of incorporating personality, individual characteristics and motivational states. In ParleE, personality traits are categorized based on different processes that an intelligent agent could perform such as perceiving, reasoning, feeling emotion and so on [21]. For example, if the agent's reasoning process focuses on undesirable effects, the agent is considered pessimistic; if the agent's reasoning process focuses on desirable effects, the agent is considered optimistic. Individual characteristics determine how easily the agent experience certain emotions.

The center of ParleE is a probabilistic planner, the domain of which is represented in STRIPS [22]. Therefore, the events which the agent takes as input are represented in STRIPS. An event could be: the outcome of the agent's actions; the occurrence of other agents' actions and the outcome of their actions.

Each time an event happens, ParleE calculates the agent's probability of obtaining the goals. The difference between this probability before and after the event happens is taken as the impact of the event on the agent's goals. Then, an Emotion Impulse Vector (EIV) is generated by appraising the event using the rules proposed by the OCC appraisal theory [17] based on the agent's goals, plans and standards. An EIV contains the values of the event's impact on emotions. The EIV is then used to update the Emotion State Vector (ESV), which contains values representing intensities of emotion.

Because ParleE is domain-independent, it is ready to use as long as the domain for the planner is well designed.

From emotions to facial expressions Six basic emotions are converted to facial expressions by a fuzzy rule-based system, which is discussed in [11].

Following Ekman and Friesen [23], we consider the following six emotions: Sadness, Happiness, Anger, Fear, Disgust and Surprise. These are said to be universal in the sense that they are associated consistently with the same facial expressions across different cultures. Ekman and Friesen also describe in detail what the expressions for these emotions and certain blends look like. Emotion feelings may differ in intensity. In [23] it is pointed out how for each of the basic emotions the expression can differ depending on the intensity of the emotion. It is therefore important for us to build our system on a representation that takes intensities into account. We have used their descriptions as the basis for our fuzzy rules.

The 3D Talking Head A muscle-based 3D face model was presented in [8]. This face model can realize the following objective: it is able to produce both realistic facial expressions and real-time animation for standard personal computers. The 3D face model contains: (i) a face mesh that allows high quality and realistic facial expressions, which is sufficiently simple in order to keep the animation real-time and is able to assist the muscle model to control the deformations; (ii) a muscle model that produces realistic deformation of the facial surface, handles multiple muscle interaction correctly and produces bulges and wrinkles in real-time.

Different facial movements on this face model are combined temporally resulting in a 3D talking head, which is described in [9].

3 The football supporter domain

In our application, we consider the situation in which a **"for"** and an **"against"** team are playing against each other. Obie is a supporter agent, while other agents are the **"for"** and **"against"** teams. For simplicity, we consider each team as a single agent. The possible actions Obie can do are: watching, cheering, etc. These actions do not affect the match being played. In order to predict the likelihood of events and to appraise events, Obie has knowledge about possible actions that other agents could take to model their plan. Examples of these actions are: tackling when the team does not have the ball, shooting when the team has the ball and the ball is in the other team's half-field or penalty area. Obie's objective is the same as the objective of the **"for"** team, which is the **"for"** team winning, whereas the objective of the **"against"** team is the **"against"** team winning.

We implement Obie's emotions focusing on only one other agent (either **"for"** or **"against"** team). Obie and the other agent perform their actions turn by turn. Obie's action is selected based on the planning algorithm, which is given in [10]. The other agent's action is extracted from a script file that describes the match or is extracted in real-time from a real soccer match or a robot soccer match.

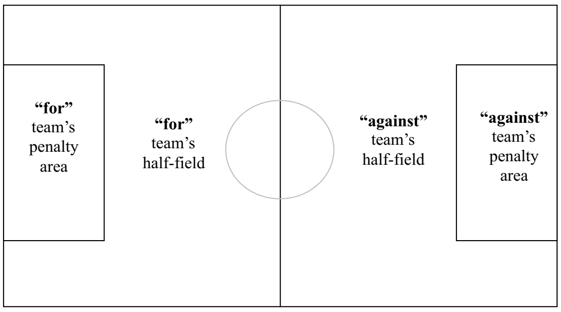

Figure 2: The view of a soccer field.

3.1 Designing the domain

In our application, an event can inform us about what action is taken, what facts are deleted and what facts are added. Facts in the football supporter domain contains the following information: what the current score is, where the ball is at this moment, which team is controlling the ball, whether a kick-off, a free-kick or a penalty is being taken, whether a free-kick is being taken.

To make the information where the ball is during the match usable, we simply divide the soccer field into four regions, which are displayed in Figure 2: **"for"** team's penalty area, **"for"** team's half-field (except the penalty area), **"against"** team's half-field area (except the penalty area), **"against"** team's penalty area.

To support a multi-agent environment, a predicate indicating which agent is acting at a certain moment in time is added to the domain actions. The main advantage of this implementation is that we can still use the planning algorithm without modifying the hard code much.

As the STRIPS format does not support quantity, we choose a simple way of representing the possible score rather than extending the STRIPS for quantity. The consequence of using a quantitative STRIPS format would require a big change in the planning algorithm. Our simple way of representing the score is to represent the current difference in the score between the **"for"** team and the **"against"** team:

$$-10, \cdots, -2, -1, 0, +1, +2, \cdots, +10$$

This way, we would not be able to represent all the numbers, but we think that 10 is reasonable maximum difference between the two teams.

3.2 *Representing the domain*

Possible facts in the domain are represented in STRIPS as follows: position of the ball: e.g., (ball-pa team) — the ball is in the penalty area of the **"for"** team; who is controlling the ball: e.g., (ball-control team) — the **"for"** team is controlling the ball; who is acting (for supporting multi-agent planning as described above): e.g., (acting team) — it is the team's turn; match status: e.g., (no-penalty team) — it is not the penalty for the **"for"** team; score: e.g., (score minus-one), (score zero), etc.

After defining the possible facts, the possible actions for a team are: attacking actions (the team is controlling the ball): e.g., *long-shot* — this action is available when the ball is in the other team's half-field; passing actions (the team is controlling the ball): e.g., *pass-from-pa* — pass when the ball is in the team's penalty area; defending actions (the team is not controlling the ball): e.g., *tackle-in-pa* — tackle when the ball is in the team's penalty area; penalty, free-kick, and kick-off related actions: e.g., *defend-freekick-in-hf* — defend free-kick in the team's half-field.

Each action requires the presence of certain facts and produces certain facts. For example, the **"for"** team's action *take-penalty* requires the presence of the fact (penalty team). Each action can have several different outputs. The action *take-penalty* may result in a goal or not. However, this action leads to a goal with a very high probability.

Each time an event happens, Obie's knowledge about the state of the world is updated. Then Obie's planner recalculates the probability of obtaining the goal, which is the **"for"** team winning. The desirability of the event is assessed based on this probability and is used to trigger Obie's emotions. For example, when the score difference is 0, the event: "the **'for'** team scores a goal from a free-kick" changes the score difference to $+1$. This event increases the likelihood that the **"for"** team will win. Therefore, Obie's happiness is triggered.

4 Illustration

The interface of our football supporter agent, Obie, is shown in Figure 3. In the top, there are two windows to show the current state of the world (current facts) and what has happened. In the bottom, there is a chart and a graph to show Obie's emotion state. There is also a 3D face model representing Obie's body.

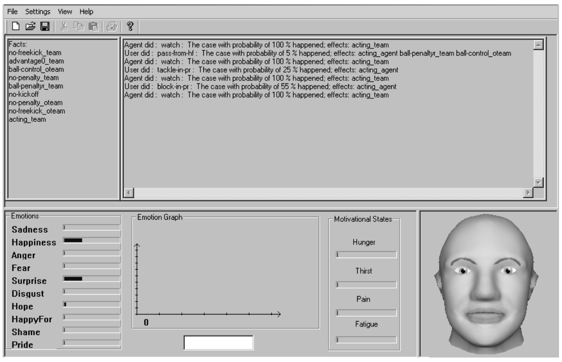

Figure 3: A snapshot of the football supporter application.

The football match starts with the score difference zero, the ball is controlled by the **"for"** team in the **"for"** team's half-field. The match lasts for ninety minutes. Each minute there is an event happening due to the actions of the two teams. The events are read from a script file. Obie's purpose is the **"for"** team winning.

As mentioned before, we consider, for simplicity's sake, each team as a single agent and we only deal with one team rather than both teams. We model the actions of the two teams through the action of one single team. For example, the action "defend free-kick" of one team means that the other team is taking the free-kick. Obie has a model of the other agent to derive the likelihood of what will happen next. Every minute, Obie appraises the occurred event based on his goals and standards which triggers some emotions that Obie will experience.

Now we will show how the match progresses and how Obie experiences and expresses his emotions during the match. There are several main events in the match. The **"against"** team opens the score at the 12th minute. The **"for"** team starts attacking back and levels the score at the 30th minute. Ten minutes later, the **"for"** team has a free-kick from their half-field. Amazingly, the free-kick reaches the **"against"** team's penalty area and still in the control of the **"for"** team. Without any hesitation, the **"for"** team makes a shot and leads by one goal. The match continues without any special event until the 80th minute when the **“against”** team nets a goal, which brings the score back to equal, i.e., the score difference back to zero.

The emotions of Obie occur in mixture. However, we will analyze individual emotions to see clearly how each emotion changes during the match. Figure 4 shows the graph of Obie’s fear during the match and Obie’s expression at the beginning of the match. Recall that Obie’s purpose is the **“for”** team winning and the starting score difference is zero. Therefore, Obie has reason to worry at the beginning of the match that the **“for”** team may not win. This fear state ends at the 40th minute when the **“for”** team scores the second goal to lead by one goal. The fear emotion is triggered again at the 80th minute when the **“against”** team levels off the score. Figure 5 shows the graph of Obie’s happiness from the 28th minute to the 46th minute, when there are two goals scored by the **“for”** team. The figure also shows Obie’s expression at the 30th minute when the first goal of the **“for”** team is scored. As can be seen from the graph, Obie’s happiness is increasing when the **“for”** team is attacking and especially after the two goals. After the 46th minute, there are no special events. As a consequence, Obie’s happiness decays to return to the neutral state. Figure 6 shows the graph of Obie’s surprise and his expression of surprise. As can be seen from the graph, surprise is triggered seldom. For example, at the 40th minute, the **“for”** team takes a long free-kick from their half-field and still receives and controls the ball in the **“against”** team’s penalty area. This situation usually happens with very low probability. Obviously, it triggers Obie’s surprise. After being triggered, Obie’s surprise decays very fast as well. Figure 7 shows the graph of Obie’s sadness during the occurrence of the **“against”** team’s second goal and Obie’s expression at the 80th minute. The graph also shows that sadness decays very slowly after it has been triggered. Blends of emotions are shown in Obie’s face as well. For example, at the 40th minute, both surprise and happiness are triggered. They both are expressed in Obie’s face, which can be seen in Figure 8.

We have also tested how Obie with different personalities experiences emotions (cf. [10]). Figure 9 shows the happiness and sadness of Obie with four different personalities: neutral, pessimistic, realistic and sensitive. As can been seen from the figure, compared to the neutral Obie, the emotion state of pessimistic Obie is biased toward negative emotions. Realistic Obie concentrates more on what has happened than what is expected. Sensitive Obie experiences every emotion more intensively than the neutral Obie.

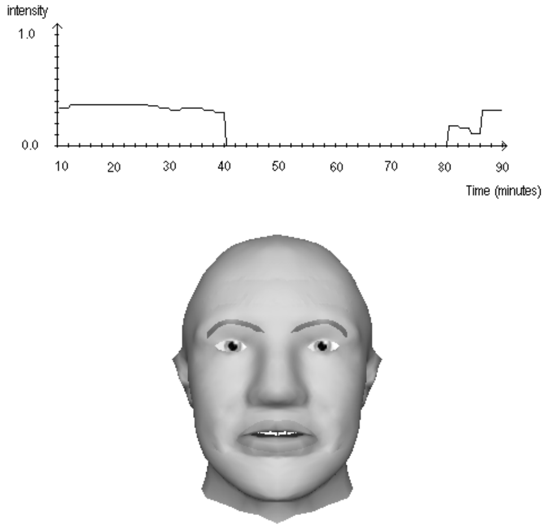

Figure 4: The graph of Obie’s fear emotion and a snapshot of his expression of fear.

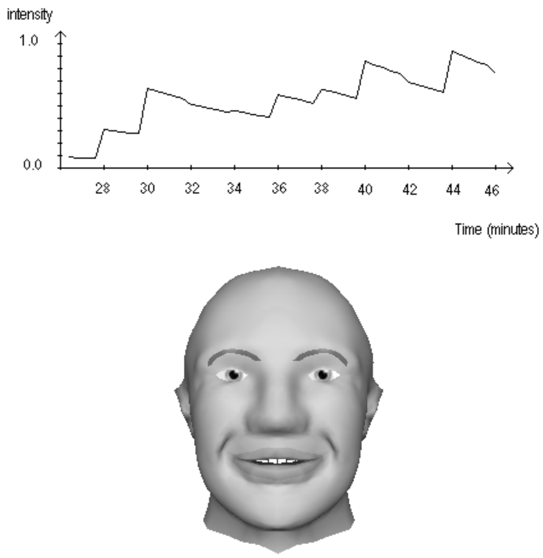

Figure 5: The graph of Obie’s happiness emotion and a snapshot of his expression of happiness.

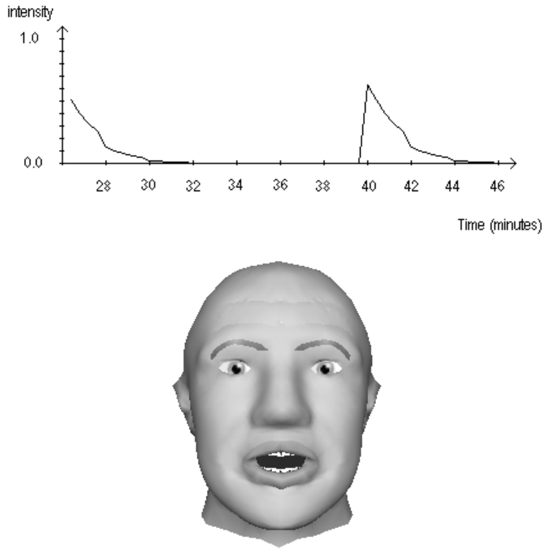

Figure 6: The graph of Obie's surprise emotion and a snapshot of his expression of surprise.

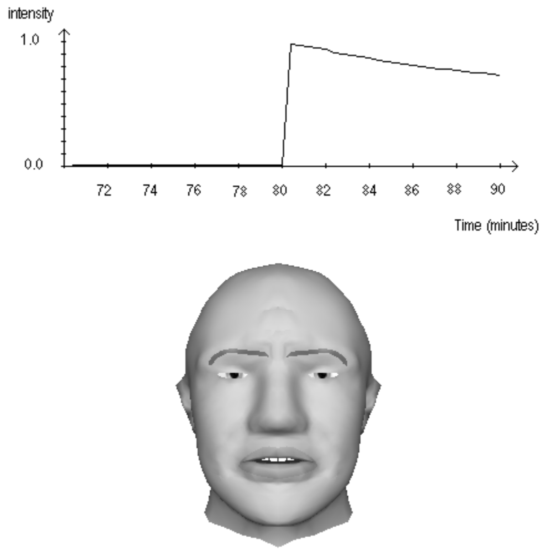

Figure 7: The graph of Obie's sadness emotion at the last goal and a snapshot of his expression of sadness.

5 Conclusion

In this paper we presented Obie, an embodied agent that experiences and expresses emotions. Obie has an emotion component which appraises events to trigger emotions [10]. Obie's emotions are expressed via his utterance or his facial expressions. The expression via utterance is done by a simple mapping from emo-

Figure 8: A snapshot of Obie's expression of happiness and surprise.

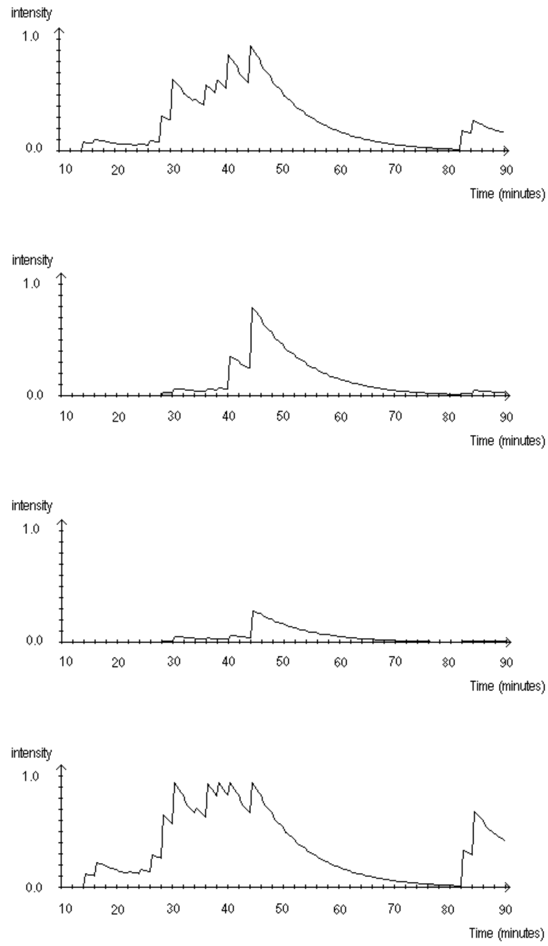

Figure 9: The graph of Obie's happiness emotion with four different personalities: neutral, pessimistic, realistic and sensitive (from top to bottom).

tions to text fragments. The mapping from emotions to facial expressions was described in [11]. Obie's utterance and facial expressions are presented in his 3D talking head [8, 9]. In the research described in this paper, Obie was implemented as a football supporter agent. We have shown how Obie experiences different emotions during a football match. We have also indicated how Obie with different personalities experiences emotions differently.

References

[1] R. Picard. *Affective Computing*. MIT Press, 1997.

[2] D.H. Galernter. *The muse in the machine*. Free Press, New York, 1994.

[3] J. P. Forgas and S. Moylan. After the movies: the effects of transient mood states on social judgments. *Personality and Social Psychology Bulletin*, 13, 1987.

[4] A.R. Damasio. *Descartes' error: Emotion, reason, and the human brain*. G.P. Putnam, New York, 1994.

[5] F. Thomas and O. Johnston. *The Illusion of Life*. Abbeville Press, New York, 1981.

[6] C. Elliott. *The Affective Reasoner: A Process model of emotions in a multi-agent system*. PhD thesis, Northwestern University, Evanston, IL, 1992.

[7] J. Gratch. Émile: Marshalling passions in training and education. In *Proc. of the Fourth International Conference on Autonomous Agents*. ACM Press, 2000.

[8] T. D. Bui, D. Heylen, and A. Nijholt. Improvements on a simple muscle-based 3d face for realistic facial expressions. In *Proc. CASA-2003*, pages 33–40, 2003.

[9] T. D. Bui, D. Heylen, and A. Nijholt. Combination of facial movements on a 3d talking head. In *Proc. CGI 2004*, 2004. to appear.

[10] T. D. Bui, D. Heylen, M. Poel, and A. Nijholt. Parlee: An adaptive plan-based event appraisal model of emotions. In *Proc. KI 2002*, pages 129–143, 2002.

[11] T. D. Bui, D. Heylen, M. Poel, and A. Nijholt. Generation of facial expressions from emotion using a fuzzy rule based system. In *Proc. AI 2001*, pages 83–95, 2001.

[12] A. Nijholt, H. J. A. op den Akker, and F. M. G. de Jong. Language interpretation and generation for football commentary. In *ACTAS-1: VIII Symposio Social*, 2003.

[13] N. S. Kooij. The development of a vision system for robotic soccer. Master's thesis, Dept. of Comp. Sci., Uni. of Twente, 2003.

[14] R. Seesink, W. Dierssen, N. Kooij, A. Schoute, M. Poel, E. Schepers, and T. Verschoor. Fast data sharing within a distributed multithreaded control framework for robot teams. In *FIRA World Congress*, 2003.

[15] J. D. Velásquez. Modeling emotions and other motivations in synthetic agents. In *Proceedings of AAAI-97/IAAI-97*, pages 10–15, Menlo Park, 1997. AAAI Press.

[16] M. S. El-Nasr, J. Y., and T. R. Ioerger. FLAME-fuzzy logic adaptive model of emotions. *Autonomous Agents and Multi-Agent Systems*, 3(3):219–257, 2000.

[17] A. Ortony, G. L. Clore, and A. Collins. *The Cognitive Structure of Emotions*. Cambridge University Press, 1988.

[18] I.J. Roseman. Cognitive determinants of emotions: A structural theory. *Review of Personality and Social Psy.*, 5, 1984.

[19] K. Oatley and P. N. Johnson-Laird. Towards a cognitive theory of emotions. *Cognition and Emotion*, 1(1):29–50, 1987.

[20] W. S. Reilly. *Believable Social and Emotional Agents*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, USA, 1996.

[21] D. Rousseau. Personality in computer characters. In *Proc. of the 1996 AAAI Workshop on Entertainment and AI A-Life*. AAAI Press, 1996.

[22] R.E. Fikes and N.J. Nilsson. Strips: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2(3-4), 1971.

[23] P. Ekman and W. V. Friesen. *Unmasking the Face: A Guide To Recognizing Emotions From Facial Clues*. Prentice-Hall, Englewood Cliffs, New Jersey, 1975.