Enhancing Embodied Conversational Agents with Social and Emotional Capabilities

Bart van Straalen, Dirk Heylen, Mariët Theune, Anton Nijholt

Human Media Interaction, University of Twente
P.O. Box 217, 7500 AE, Enschede, The Netherlands
(straalenb, d.k.j.heylen, m.theune,a.nijholt)@ewi.utwente.nl

Abstract. In this paper we present our current work on an embodied conversational agent for training medical bad news conversations and discuss the inspiration gained from previous work of our own and others. Central in this research is the influence of emotional and social features on the selection and realization of conversational behavior.

Keywords: Embodied Conversational Agents, Social Agents, Bad News Conversations, Tutoring, Empathy

1 Introduction

Over the last few decades there has been an important shift in much work on dialogue systems. Traditional spoken dialogue systems were created to fulfill a very specific task, using spoken or written language. By combining spoken dialogue systems with a graphical representation of a human (or human-like entity), so-called Embodied Conversational Agents (ECAs) are able to communicate not only verbally but also nonverbally. In addition, because of their embodiment ECAs are perceived more like intelligent agents endowed with a personality. The work on ECAs thus incorporates more than natural language understanding; it also includes other aspects of cognitive modeling, such as models of emotion and social skills. The last few years have seen a large increase in the research that has been done on modeling emotions [1-5] and on modeling social skills [6-8].

In this paper we present our current work on embodied conversational agents and discuss the inspiration and insight gained from previous work of our own and others. Two types of ECAs are described in order to take stock of part of the state-of-the-art and to point out some of the difficulties the scientific community is dealing with at the moment. We selected tutoring agents and agents in interactive pedagogical drama to discuss, because their behaviors typically contain emotion and social features, which are of great importance in these types of tasks [15]. As the design of our agent, which will hold bad news conversations, also includes these features, we look at these types of agents to gain insight and inspiration about how to implement such features. In addition we look at which developments of traditional types of conversational systems have improved interactions with virtual humans in various ways during the last

decade. Components from some of the described systems are taken into account in the design of the agent architecture. Additionally, the design of the virtual agent is based on theoretical models of human cognition to allow virtual humans to behave more as real human beings, both in physical behavior as in mental processes. This approach is founded on the belief that "the best representation of an object is itself." By making virtual humans to be more humanlike, the quality of interacting with such agents will also most likely be improved, opening the door for new improvements and applications.

2   Tutoring Agents

In this section we take a look at virtual tutor systems to see which features play a role in selecting and expressing conversational behavior (in tutoring). In ECA research, tutoring and coaching have been a popular choice of tasks as they display a lot of different aspects of conversational interaction [9-11]. Typically the actions of a tutor includes giving instructions, asking and answering questions, providing explanations, giving examples, setting specific tasks and objectives, motivating the student, providing feedback (both positive and negative) throughout the training session and afterwards, providing support and evaluating the performance of the student. Based on this list the actions of a tutor can be divided into two broad categories: 1) providing information regarding the task at an appropriate level towards the learner and 2) engaging and directing the learner through the learning process. The challenge is to build correct models of the cognitive components that lead to the selection and realization of these actions, as it is often unclear which components lead to a certain behavior, how they function and how they interact with each other.

For current day ECA research the second category is particularly interesting, as it deals with the social and emotional skills that are of great importance when performing the role of tutor. In order to keep a learner motivated and challenged a tutor may need to praise or blame the learner (emotion directed behavior), adopt the role of a study-buddy (altering the social relation between tutor and learner) and keep track of what the learner is thinking and feeling (abducting the mental state of the user). By adapting the virtual tutor's social and emotional skills to each individual learner, the outcome of a tutoring session will most likely be improved significantly. Furthermore it will contribute to the effort to make conversational agent to behave in a more humanlike fashion.

As an example the remainder of this section will describe an intelligent tutoring system we developed called INES (Intelligent Nursing Education Software) [12-14]. The INES system is designed with the purpose of helping students practice nursing tasks using a haptic device within a virtual environment. A virtual human in the INES system provides the role of tutor with which the learner can interact (see figure 1). The virtual tutor is capable of performing acts from both categories mentioned above, but focuses on affective control of the mental state of the learner in the tutoring dialogues [12]. It does so by selecting the appropriate feedback to give to the learner after he or she has performed an action. In order to select the appropriate feedback the tutor makes an assumption about the learner's mental state and consequently adapts

the selection of its type of action, the affective language it uses and the overall tutoring strategy. For example the tutor might say "It was quite a difficult task. Try again, but put the needle in more slowly." instead of saying "You put the needle in too fast. Try again." if the learner comes across as being hesitant. The assumption is based on the observed behaviors of the learner. This includes, amongst others, the learner's confidence level and an appraisal of the learner's actions while he or she is performing the task: Did the learner make many mistakes? How grave were those mistakes? How is the overall performance so far? How (pro-) active is the student? Furthermore, the tutor system also takes into account the difficulty of the task and the emotional effect previous feedback had on the system itself. All these aspects are used to estimate the affective and motivational state of the user (anxious-confident, dispirited-enthusiastic), as well as the performance of the task [15]. As a result the socio-emotional aspects of the interaction between the learner and the virtual tutor do not only influence the learning strategies the tutor adopts but also the manner in which the conversational actions are expressed.

Fig. 1. The INES system. The student is using a haptic device that transforms to an injection needle in the virtual environment (displayed on the *left screen*). In this environment the student can interact with a *Virtual Patient* and perform simple nursing tasks. The performance is monitored by the *Virtual Tutor* (displayed on the *right screen*).

The INES system contains the variables `happy-for' and 'sorry-for' in its mental model that are updated depending on the student's success. The emotions the agent experiences are thus related to the behavior of the learner. The extent to which it responds naturally to the situation is restricted to influencing the learning process in a positive way. These variables are used to adjust the type of feedback.

It is apparent that in general the focus within a tutoring system lies on optimizing the learning achievements of the user, not on making the tutor agent behave as human-like as possible. Although the tutor agent's cognitive capabilities have been improved, e.g. adapting the complexity of the information to the level of the learner and using emotional and social conversational behavior to motivate and engage the learner, these improvements only influence the tutor agent's behaviors to the extent that they increase the performance of the learner. The improved cognitive processes do not cause the tutor agent to act more with its own interests in mind. Nevertheless,

the conversational behaviors and the cognitive models of emotions and social skill of tutoring systems described in this section provide us with a good insight on which features play a role in tutoring interactions with humans and how they influence the learner's behavior. By utilizing this knowledge and modifying the cognitive models of emotion and social skills that are described in this section, we aim to create a virtual human that can behave without restricting its behavior by a specific focus.

3 Interactive Pedagogical Agents

This section describes the second type of embodied conversational agents; agents in Interactive Pedagogical Drama [16]. Interactive Pedagogical Drama is a style of educational instruction that has the goal of teaching learners the skills that are necessary to cope with stressful and difficult situations. Within an Interactive Pedagogical Drama, learners interact with believable virtual characters in a story that is recognizable for them and elicits empathy. The goal is that by allowing the virtual characters to face and overcome difficulties, which are similar to those the learners are facing, the learners experience and learn skills that can be used to deal with their own problems. Through interaction with the system the learners can steer the story in such a way that it handles specific problems and solutions they are interested in. While the learner can influence the story, the virtual characters select their actions on their own. Interactive Pedagogical Drama differs from the tutoring systems described in the previous section in the following way. Instead of being actively encouraged by the virtual human to perform the learning task as the virtual tutor instructs, in interactive Pedagogical Drama the learner learns by observing the story. Because the conversational behavior of the virtual characters in interactive pedagogical drama is focused on the story and on other agents instead of on the learner, their behavior selections are less restricted than that of a virtual tutor. Virtual characters in interactive personal drama might also express emotions and social skills that will not surface in a tutoring system, such as anger, frustration or impolite behavior. In order to allow the learners to have a productive interaction with the drama, let them believe in the efficiency of the skills used by the virtual characters and subsequently apply those skills in their own life, it is important that the system has the following characteristics: First of all, the learners must be able to identify themselves with the characters in the story. Secondly the difficulties the virtual characters experience must both be believable and familiar to the learners. If one of these characteristics is lacking the suspension of disbelief of the entire drama fails and the learners will not benefit from the interaction. This means it is vital that the behavior the virtual humans perform is as plausible as possible. Furthermore it can be desirable that when a virtual human is asked why it performs a certain behavior it is able to give a plausible explanation.

A well-known interactive pedagogical drama system that used an agent-based approach is Carmen's Bright IDEAS. Carmen's Bright IDEAS is an interactive health intervention system designed to improve the problem solving skills of mothers of pediatric cancer patients [16]. Parents of children with a chronic disease often lack the capabilities of dealing with many demands their sick child demands and the needs of

their spouse, their healthy children and their work. The goal of the system is to teach a specific approach to social decision making and problem solving called Bright IDEAS [17] and to help parents in dealing with difficult situations. The drama of Carmen's Bright IDEAS narrates the following scenario: It relates the problems and stresses of the protagonist of the story, Carmen, who has a nine-year-old son with pediatric leukemia and a six-year-old daughter. Carmen discusses her problems with a counselor, Gina, who suggest she uses Bright IDEAS to help her to deal with difficult situations. With Gina's help Carmen goes through the initial steps of Bright IDEAS then completes the remaining steps on her own (see figure 2).

Fig. 2. Carmen's Bright IDEAS. *Carmen*, the agent on the right, discusses her problems with *Gina*.

In Carmen's Bright Ideas, the user can influence the course of the story at specific points, but does not participate directly as a story character. Instead the learner can control the actions of Carmen at an intentional level by directing the thoughts and feelings Carmen might have in a certain situation when asked about it. Subsequently, the selected thoughts and feelings are incorporated into the mental model of Carmen. This will result in the virtual agent to perform actions which are congruent with the thoughts and feeling that were selected by the learner. The thoughts and emotions the learner can choose from are formalized in such a way that the learner is able to identify them and relate herself to the situation. Where Carmen allows the learner to interact with the drama, Gina's task is to make sure the social problem solving technique is followed. The virtual counselor does so by appropriately responding to the actions of Carmen and motivating her through dialogue and gestures.

The agent architecture in Carmen's Bright IDEAS is based on a multi-layer transition-based agent model called Situation Spaces [18]. This entails that there are specific states in which the agents can find themselves, defined by the situation at that moment. The agents select their behavior based on the state they are in, instead of an event that occurs in the world. For virtual characters in Carmen's Bright IDEAS four different layers exist, each with a variety of states; problem solving, dialogue model, physical focus and emotional appraisal. For the Gina character the problem solving layer is used to give form to the dramatic structure, including the IDEAS steps and the strategies Gina uses to realize these steps. The dialogue model is used to select and

execute dialogue acts to bring these strategies about. In Carmen's case the problem solving and dialogue models are more reactive and focus on responding to the communicative actions of Gina. In both characters the physical focus layer is used to manage and execute non-verbal behavior and the emotional appraisal layer covers the agents' emotional appraisal model [16]. Although both virtual characters can make use of non-verbal communicative behavior, it has a larger impact on the learner when it is performed by Carmen.

One of the key points of Interactive Pedagogical Drama is that the learner is able to see the causal relations between a selected intention, the accompanying behavior and the effect that behavior has. As both the selectable intentions and the response behaviors are plausible, the insight in the cognitive processing helps the character to become more human-like. Also, in Interactive Pedagogical Drama the focus of the virtual characters is on the conversation and only secondary on the task of teaching the learner something. In order to facilitate a more appropriate conversation, the behavior selection of the virtual characters will be less restricted than that of tutoring agents where the main focus of their behavior lies on the tutor task. Consequently the virtual characters in an Interactive Pedagogical Drama can select from a much wider range of conversational behaviors and thus will display more diverse emotions and social skills in their dialogues. As our goal is to make an embodied conversational agent that acts like a human, both behavioral and mentally, the insights in the relationship between intentions and behavior and the insights in the emotion and social features gained by research on Interactive Pedagogical Drama are very useful.

4   A Bad News Agent

Having looked at several state-of-the-arts ECA systems, we have gained some insight on which features are important in conversations and how these features influence a (virtual) human's behavior. Whether a conversation with an ECA involves tutoring, counseling, entertaining or is related to another task, it is of great importance that the virtual character's actions are plausible. This can be achieved by trying to aim for realism both externally and internally, by which we mean that the agent performs close to human-like behavior (i.e. physically) and is driven by close to human-like cognitive models. At first glance there are two benefits. First, by modeling human cognitive processes insight is gained in the practical workings of human cognition, albeit at an abstract level and secondly virtual characters will be able to provide a more understandable and realistic explanation of why they have chosen to perform certain behaviors. All in all, this will contribute in making virtual human be more human-like.

In order to realize a realistic virtual human, we are currently designing a virtual character that makes use of cognitive models that are involved in the selection of appropriate conversational behavior i.e. when the virtual human is engaged in an interaction, what kind of behaviors does it select, why does it select them and how do these selection mechanisms influence the manifestation of the behavior? The cognitive models are based on psychological and sociological theories of human cognition. The purpose of the system is to assist physicians (in training) to practice

holding bad news conversations. To this end the function of the virtual character is twofold. Primarily it plays the role of the virtual patient that is receiving the bad news and secondly it performs a tutor role, giving feedback to the learner about why it responded the way it did (see figure 3).
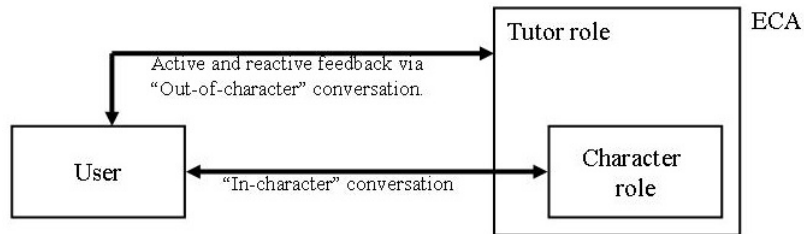

Fig. 3. This is the overall layout of the agent. Notice that the *Character role* is included in the *Tutor role*. This has the advantage that the *Tutor role* can easily refer to the behaviors the *Character role* has performed when the *Tutor role* gives feedback to the *User*.

The conversational behaviors of the virtual human are not restricted to facilitating the optimal learning experience (e.g. steering the conversation so the learner will be challenged more) as is the case in more traditional tutoring systems, but instead the virtual human is only limited to perform behaviors that are appropriate given the situation.

4.1 Bad News Domain

It is important to understand what we mean by "bad news", so as to place this research in the correct context. In general, a bad news conversation is a dialogue in which the "speaker" discloses information that is unfavorable to the "listener". In this research the topic of the unfavorable information will involve the "listener's / receiver's" medical condition, such as with patients that have a terminal disease. Definitions of bad news as they are used in different studies are: "Any information which adversely and seriously affects an individual's view of his or her future." [19] or "news that will change a patient's outlook for the future in a very negative way. Such bad news can be about a severe illness, prospects of death or increasing levels of limitations." [20]. By using bad news conversations as the task of the interaction with the virtual agent and subsequently trying to model the cognitive processes involved, we hope to gain insight in a variety of advanced cognitive behaviors such as affective and social aspects.

A significant amount of research has been done on how someone should conduct bad news conversations, resulting in several detailed protocols and strategies that describe the best way of delivering bad news [19, 21, 22]. Unfortunately, few of these studies have looked at the response behaviors of the receivers of bad news. We are particularly interested in the way receivers respond, both verbally and non-verbally, to receiving the bad news, why they respond in such a manner and most importantly how such behaviors can be modeled into a virtual human. To this end psychological

and sociological literature has been studied on the subject of coping mechanisms and strategies [23, 24]. Of particular relevance is the work of Elizabeth Kübler-Ross, who in her work has gained a great understanding on behavior of terminally ill patients [25]. This has led to the formation of the well-known categorization of coping strategies that are utilized by dying patients: Denial and Isolation, Anger, Bargaining, Depression and Acceptance. By analyzing these different coping mechanisms in terms of affective social aspects we hope to gain an idea on how they influence conversational behavior. So far we have split the influence of coping strategies into two categories: one category that influences the selection of behavior (i.e. which behavior should the virtual agent perform) and another category that influences the manifestation of the selected behavior (i.e. in what manner is the behavior carried out). Contained in the coping strategies are emotional and social features that cause the influence of behavior. A good example is the Anger coping behavior to receiving bad news. When this strategy is adopted the virtual human will select appropriate conversational behavior to cope with the situation. This may be result in selecting the "assigning blame" speech action (attribution emotion) and its accompanying gestures. The selected speech action will be impolite (social relation) and gesturing will have characteristic that are associated with anger (short, strong movements). However in order to use these coping strategies the virtual humans must possess emotions and social skills. To that end we incorporate models derived from cognitive theories.

4.2 Bad News Agent Architecture

The basis of the agent architecture of our Bad News agent is the Beliefs, Desires and Intentions (BDI) cognitive model [26-28], one of the most well-known and most studied models used in creating reasoning intelligent agents [29]. One of the basic components of the architecture is a belief-base that contains all the virtual character's beliefs about the world (including other agents, such as it human interlocutor, and itself). Beliefs describe the agent's subjective interpretations of the situation and not an objective representation. For instance, if a patient is given an estimation of one year of remaining life and the normal prognosis is four to six months, he still might believe that one year is short.

For the aspect of desires we intend to include a goal-base that contains a variety of goals the agent may adopt. The main difference between desires and goals is that desires are roughly unrestricted objectives or situations that the agent *would like* to achieve or bring about. Goals on the other hand are those desires that are actively pursued and have to be consistent with each other. For example, an agent might have the desires to *Get out of the hospital* or to *Get treatment in the hospital*. Obviously these are two desires that are not consistent with each other and as such only one can be adopted as a goal.

As described in the BDI model, intentions represent the deliberative state of an agent, i.e. which goal it is committed to and is it trying to achieve by executing its plan of action. However, in our architecture we make a distinction between the terms intention and communicative intent. While the former may be a proper representation of commitment to a goal resulting in a plan of communicative behaviors, the latter is a description of which behavior or state of mind a communicative act is trying to elicit

from the interlocutor. For example, if the goal the agent is committed to, (i.e. its intention) is to be comforted by getting a positive reassurance from the doctor, the agent might ask "Everything is going to be okay, isn't it?". This question contains two communicative intents. The first intent is to cause the doctor to believe that the agent believes that everything is going to be alright (if he did not think so already). Secondly the question intends to elicit a response in which the doctor confirms the belief of the agent that everything is going to be alright. If the agent forms a communicative intent, this intent will be expressed by its behavior. More specifically, the wished-for effect that a communicative act brings about (i.e. the communicative intent) is contained in the communicative act itself, but it depends on the interpretation by the interlocutor if this effect is achieved. In case of the example, if the doctor ignores or fails to understand the communicative intent he might answer truthfully that everything is NOT going to be alright. A communicative action that does not allow for the possible realization of the communicative intent has no purpose and as such will not be performed. This wished-for effect should not only cause a change in the belief, goal or intention state of the addressee, but should be extended to include the user's affective state. The purpose of the wished-for effect is to get the interlocutor to believe something, know something, do something, feel something etc. that causes the agent to be closer to achieving its own goal.

Although beliefs, desires (or goals) and intentions form a basis for an embodied conversational agent, it is necessary to also incorporate affective and social features in order to create believable interactive virtual humans. To that end we also include a component in the agent architecture that deals with affect. The main content of this component at this point is an emotion appraisal model of our own design, which is based on a conglomeration of features obtained from existing emotion and emotion appraisal models and fitted together. It incorporates features from the OCC model [4], the Affective Reasoner (AR) model [2] and the EMA model [3]. The OCC model is a well-known theoretical model of human emotion. It has been the basis for several state-of-the-art emotion appraisal systems such as EMA and that of FearNot![1]. The OCC model evaluates how the state of the world influences the emotions of the person. However the OCC model does not take into account the mental states of other agents when it is determining which emotion should be elicited based on the situation. In order to make a believable and realistic virtual human this capability needs to be included, so that the virtual agent can generate appropriate responses to the human interlocutor. Therefore we extract features from the Affective Reasoner model that are able to deal with the mental state of others. Also the OCC model does not describe how the formed emotions influence the selection and execution of behavior, which is a problem that needs to be addressed if a virtual human is to be created. Some of the features of the EMA model are included in our affective model as the EMA model utilizes a set of appraisal variables that can be used quite easily in the context of bad news conversations.

In addition to emotion, social skills also play an important role in natural conversation. They influence both the selection of conversational behavior and the realization of that behavior. Each social situation requires a particular type of behaviors. This is represented in the architecture by dividing the virtual human's conversational behavior into categories that correspond with different social situations. Only behaviors from the category that corresponds to the social situation

can be selected. Also each category contains a set of labels that are passed on to the behavior realizer to dictate how the behavior should be executed (e.g. volume of speech, specific facial expressions, and characteristics for gestures).

For the input of the cognitive processes we make the assumption that the virtual human has interpreted input signals from the environment. These signals are all the features that make up the conversational behavior of the learner. Verbally this entails, amongst others, the learner's speech act (both the prosody and the content), when the learner speaks and non-verbally this refers to the learner's facial expression, his gaze behavior, his head movement and his body posture as detected by the virtual human. Consequently the virtual human forms a *causal interpretation* [3]. This is a configuration of the agent's belief-base, goal-base and intention at a certain time-point that represents the agent's subjective interpretation of the relationship between the agent and the environment plus the interpreted but not yet fully processed input signals. The causal interpretation is then handled by cognitive processes that are based on cognitive models. These processes include updating the belief-base (including beliefs about the social relationship between the virtual agent and the human interlocutor), updating the goal-base, emotional appraisal, adjusting plan of behaviors and monitoring the environment. As a result of the cognitive processing, an intention (and communicative intentions) is selected. This selection is influenced by the social relation between the virtual human and the environment, the coping strategy the agent has selected and the output of the emotion appraisal. Based on this intention an appropriate category of behaviors is selected from a behavior library. For example, if the intention is to *give an answer to the interlocutor*, then the "answering"-category of behavior is selected. Subsequently, the behavior that is most appropriate, according to the virtual human's current the social state and the emotion appraisal, is selected. In the situation of a bad news conversation it is likely the social relationship with the interlocutor (a doctor) is quite formal. This will lead to a formal, polite type of answer such as "I understand doctor" instead of "yeah okay doc". Additionally the emotional state of the virtual human (most likely dismayed and sad) will result in a terse answer: "I understand doctor" instead of "It is perfectly clear what you are saying to me doctor." The manner in which the behavior is performed (e.g. variables in prosody and gestures) also depends on the emotion state and the social state of the virtual human.

5 General Discussion

Embodied conversational agents can benefit significantly from the inclusion of social and affective features. Although these features can be included in many ways, a good way of creating virtual humans that represent real humans seems to be to incorporate cognitive models of psychological and sociological theories into the design of such agents. Particularly the inclusion of cognitive models that deal with emotions and social skills are important as they greatly increase the capability of virtual human to behave as a real human. Both emotions and social features influence a virtual human's behavior in two manners: first by influencing which conversational behavior

the agent selects and secondly by influencing the way this behavior is executed. In our development of the Bad News Agent, we are trying to combine both of these.

References

1. Aylett, R., Louchart, S., Dias, J., Paiva, A., Vala, M.: FearNot! - An Experiment in Emergent Narrative. In: IVA 2005. LNCS, vol. 3661, pp. 305—316, Springer, (2005)
2. Elliott, C.D.: The affective reasoner: a process model of emotions in a multi-agent system. Northwestern University, Evanston (1992). (PhD Thesis)
3. Marsella, S., Gratch, J.: EMA: A computational model of appraisal dynamics. In: Gratch, J., Marsella, S., Petta, P. (eds.) Agent Construction and Emotions, pp. 601—606. Vienna (2006)
4. Ortony, A., Clore, G.L., Collins, A.: The Cognitive Structure of Emotions. Cambridge University Press (1988)
5. Scherer, K.R.: What are emotions? And how can they be measured? Social Science Information. 44 (4), pp. 695—726 (2005)
6. Bickmore, T.W.: Relational Agents: Effecting Change through Human-Computer Relationships. MIT, Cambridge (2003) (PhD Thesis)
7. Bickmore, T.W., Schulman, D., Yin, L.: Engagement vs. Deceit: Virtual Humans with Human Autobiographies. In: IVA 2009. LNCS, vol. 5773, pp. 6—19, Springer, (2009)
8. Cassell, J., Bickmore, T.W., Billinghurtst, M., Campbell, L., Chang, K., Vilhjálmsson, H., Yan, H.: In: CHI 1999. pp. 520—527, (1999)
9. Gratch, J., Marsella S.: Tears and Fears: modeling emotions and emotional behaviors in synthetic agents. In: Agents 2001. pp. 278—285, ACM Press, (2001)
10. Johnson, W.L., Rizzo, P., Bosma, W., Kole, S., Ghijsen, M., van Welbergen, H.: Generating socially appropriate tutorial dialog. In: Affective Dialogue Systems 2004. LNCS, vol. 3068, pp. 254—264, (2004)
11. Rickel, J., Johnson, W.L.: STEVE: A Pedagogical Agent for Virtual Reality. In: Agents 1998. pp. 332—333, ACM Press, (1998)
12. Heylen, D., Nijholt, A., op den Akker, R.: Affect In Tutoring Dialogues. In: Applied Artificial Intelligence 19(3-4), pp. 287—311 (2005).
13. Hospers, M.A., Kroezen, E., Nijholt, A., op den Akker, R., Heylen, D.: An agent-based intelligent tutoring system for nurse education. In: Applications of Intelligent Agents in Health Care. pp. 143—159, Birkhauser Publishing Ltd. (2003)
14. Poel, M., op den Akker, R., Heylen, D., Nijholt, A.: Emotion based Agent Architectures for Tutoring Systems: The INES Architecture. In: Cybernetics and Systems 2004. Workshop on Affective Computational Entities (ACE 2004). pp. 663—667, Vienna, (2004)
15. Heylen, D., Theune, M., op den Akker, R., Nijholt, A.: Social Agents: the First generations. In: Proceedings of the International Conference on Affective Computing and Intelligent Interaction. (2009)
16. Marsella, S., Johnson, W.L., LaBore, C.: Interactive Pedagogical Drama. In: Agents 2000, pp. 301—308, ACM Press, (2000)
17. Varni, J.W., Sahler, O.J., Katz, E.R., Mulhern, R.K., Copeland, D.R., Noll, R.B., Phipps, S., Dolgin, M.J., Roghmann, K.: Maternal problem-solving therapy in pediatric cancer. Journal of Psychosocial Oncology, 16, pp. 41—71 (1999).

18.Marsella, S., Johnson, W.L.: An Instructor's Assistant for Team-Training in Dynamic Multi-Agent Virtual Worlds. In: ITS '98 . LNCS, vol. 1452, pp. 464—473, Springer (1998)

19.Baile, W.F., Buckman, R., Lenzi, R., Glober, G., Beale, E.A., Kudelka, A.P.: SPIKES-A six-step protocol for delivering bad news: application to the patient with cancer. The Oncologist, 5 (4), pp. 302—311 (2000)

20.Orlander, J.D., Graeme Fincke, B., Hermanns, D., Johnson, G.A.: Medical residents' first clearly remembered experiences of giving bad news. Journal Of General Internal Medicine, 17(11), pp. 825—840 (2002)

21.Friedrichsen, M.J., Strang, P.M.: Doctors' strategies when breaking bad news to terminally ill patients. J Palliat Med, 6(4), pp. 565—574 (2003)

22.Garg, A., Buckman, R., Kason, Y.: Teaching medical students how to break bad news. CMAJ: Canadian Medical Association Journal, 156, pp. 1159—1164 (1997)

23.Carver, C.S., Scheier, M.F., Weintraub, J.K.: Assessing coping strategies: A theoretically based approach. Journal of Personality and Social Psychology. 56, p. 267-283, (1989)

24.Folkman, S., Lazarus, R.S.: An analysis of coping in a middle aged community sample. Journal of Health and Social Behavior, 21, pp. 219—239 (1980)

25.Kübler-Ross, E.: On Death and Dying, Scribner (1969)

26.Cohen, P.R., Levesque, H.J.: Intention is Choice with Commitment. Artif. Intell., 42(2-3), pp. 213—261 (1990)

27.Rao, A.S., Georgeff, M.P.: Modeling Rational Agents within a BDI-Architecture. In Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR'91). pp. 473—484, Morgan Kaufmann publishers Inc. (1991)

28.Wooldridge, M.: Reasoning about Rational Agents. MIT press, (2000)

29.Georgeff, M., Pell, B., Pollack, M., Tambe, M., Wooldridge, M.: The Belief-Desire-Intention Model of Agency. In: Proceedings of the 5th International Workshop on Intelligent Agents: Agent Theories, Architectures, and Languages (ATAL-98). LNCS, vol. 1555, pp. 1—10, Springer-Verlag, (1998)