

ANTON NIJHOLT
University of Twente
Human Media Interaction
Enschede, The Netherlands
anijholt@cs.utwente.nl

No Grice: Computers that Lie, Deceive and Conceal

1 Introduction

In 1982 Time Magazine made the computer 'Machine of the Year'. Until then, famous people such as Ronald Reagan, Lech Walesa and Ayatullah Khomeini had been chosen as 'Man of the Year'. That tradition was continued after 1982. Interestingly, in 2006 it was again the computer that appeared on the cover of Time magazine in its yearly election of the 'Person of the Year'. However, now it said, 'You.' 'Yes you. You control the Information Age. Welcome to your world.' This change from making the computer the machine of the year to a statement in which it is assumed to be necessary to make explicit that humans are in control illustrates that indeed, there can be doubts who is in control. Interaction with a computer in a human-like way, in particular, using natural language, has been the topic of research since the time of the early computers. Chatbots, question-answering systems and dialogue systems, have been designed and during the 1980s and the 1990s of the previous century such systems have been demonstrated in research environments and have shown their (very) limited use in chatbots that are offered at web pages or spoken language systems that give access, by telephone, to certain information services. There is no way that a user can be considered to be in control. The user has to adapt to the system, he or she is commanded to provide information at a time and in a way the system is assumed to be able to understand. And most of the time the system does not understand. Spontaneous speech cannot be understood and in situations where there are several speakers interacting with each other – a multiparty interaction situation - it is yet impossible for a computer interaction system that plays a role as a computational partner (speaker, listener, both) to determine how, for example, the turn taking process takes place and when and how to place an interruption. However, exactly these topics have received lots of attention in European research projects in the previous years and it has become clear how difficult it is to model natural human-human interaction and to replace one of the humans in human-human interaction by a computational agent .such that we have human-like human-artificial agent interaction.

Despite slow progress in natural and human-like human-computer interaction, it nevertheless remains a main research aim. Progress in speech and natural language processing research has halted. Our ability to model speech processes and natural language understanding has not progressed in the last twenty years, except when you are willing to measure this progress in promilles rather than in percents. Nevertheless there is optimism when looking at modeling human-computer interaction, in particular when looking at modeling nonverbal aspects of such interaction. New sensor technology has made it possible to track nonverbal interaction cues and activities. Current research activity, for example in various large-scale European research projects, is aiming at using sensor technology and sensor data interpretation of nonverbal aspects of human-human interaction and of human behavior activity in general. Again, as has been the leading principle of research in the past, the assumption is that we can model human-human interaction, preferably in a multi-party interaction setting, and that this knowledge can be used to design more 'natural' interfaces between humans and computer-supported environments in 'daily-life' situations.

As discussed in [10,15,17] in the future our daily life interactions with other people, with computers, robots and smart environments will be recorded and interpreted by computers or embedded intelligence in environments, furniture, robots, displays, and wearables. These sensors record our activities, our behavior, and our interactions. Fusion of such information and reasoning about such information makes it possible, using computational models of human behavior and activities, to provide context- and person-aware interpretations of human behavior and activities, including determination of attitudes, moods, and emotions. Sensors include cameras, microphones, eye trackers, position and proximity sensors, tactile or smell sensors, et cetera. Sensors can be embedded in an environment, but they can also move around, for example, if they are part of a mobile social robot or if they are part of devices we carry around or are embedded in our clothes or body. Clearly, we may assume that in the future these sensors and their embedded intelligence are connected and this allows reasoning about and interpretation of a person's activities, behavior and interactions in a comprehensive way, leaving no aspect of his or her life hidden. Our daily life behavior and daily life interactions are recorded and interpreted. The general aim of such research is to equip environments and devices with intelligence that supports their inhabitants or their users in their activities. The environment can support the user in a reactive or in a pro-active way. That is, the environment can observe the user and his or her activities, can draw conclusions about a mental or physical state of the user, can draw conclusions about the intentions of the user (knowing about a user's preferences, an interaction history, a current mental state) and can then decide to offer or suggest the user certain alternatives to choose from, or to adapt the interface to the affective state of the user (where adaption includes offering different interaction modalities or ways to use them).

How can we use such environments and how can such environments use us? Do we always want to cooperate with these environments; do these environments always want to cooperate with us? In this paper we argue that there are many reasons that users or rather human partners of these environments do want to keep information about their intentions and their emotions hidden from these smart environments. On the other hand, their artificial

interaction partner may have similar reasons to not give away all information they have or to treat their human partner as an opponent rather than someone that has to be supported by smart technology.

This will be elaborated in the forthcoming sections. In the next sections we will survey examples of human-computer interactions where there is not necessarily a goal to be explicit about intentions and feelings. In subsequent sections we will look at (1) the computer as a conversational partner, (2) the computer as a butler or diary companion, (3) the computer as a teacher or a trainer, acting in a virtual training environment (a serious game), (4) sports applications (that are not necessarily different from serious game or education environments), and games and entertainment applications. In all these examples it sometimes can be useful or even necessary to hide feelings and intentions. In some examples the 'game' includes not showing emotions, not being honest, or not being cooperative. Conversations, training, and games suppose some cooperation or 'common ground'. The computer can be our opponent, our trainer, or someone we don't agree with but who we need to convince in a negotiation process. There may also be situations where we simply don't trust the decisions of the intelligence that is embedded in our environments. That is, assume a situation where we think we know better than the computer and because of that, can we overrule the computer?

2. Cooperation and Non-Cooperation

2.1 Grice and Cooperation

We should start our observations with an acknowledgement to the Gricean principles that have been followed by Q&A and natural language dialogue systems. The main assumption is that users are cooperative. In [6] this is called the Cooperative Principle: Make your contribution such as it is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged. Speakers (generally) observe the cooperative principle, and listeners (generally) assume that speakers are observing it (conversational implicature). Grice also introduced some Conversational Maxims, such as the Maxim of Quantity: Make your contribution to the conversation as informative as necessary; and Do not make your contribution to the conversation more informative than necessary. There are other Maxims. The Maxim of Quality: Do not say what you believe to be false; and Do not say that for which you lack adequate evidence. The Maxim of Relevance: Be relevant (i.e., say things related to the current topic of the conversation). And, the Maxim of Manner: Avoid obscurity of expression; Avoid ambiguity; Be brief; and Be orderly.

Clearly, although Grice was talking about human-human interaction, one can imagine how attractive these rules are from the point of view of human-computer interaction, where there is already more than enough difficulty for the computer to understand short task-oriented commands and where we certainly don't want these commands hidden in obscure, ambiguous, and non-relevant information that maybe has to be understood before it can be discarded. Obviously, this Gricean view has been attacked, extended and refined. There have been discussions about these principles and maxims and whether they are descriptive and prescriptive, refinements have been introduced that include ethical considerations and refinements have been introduced that look at the importance of nonverbal communication and how nonverbal signals should be included in these views (see e.g. [1]). Many researchers looked at ways to model mutual and cooperative understanding of face-to-face human-human understanding. And, they were quite willing to look beyond these Gricean principles. What makes it possible that we are able to deal with indirect speech acts [23] or with politeness expressions [12]? The use of humor during conversations is another topic that has been addressed [14]. Using indirect speech acts, being polite, or using humor are examples where conversational partners do not follow the Gricean rules, while certainly, these are examples of natural interaction behavior. We will return to these issues in the next sections.

But, it can be concluded that since the 1970's the majority of dialogue or interaction system researchers have been working on modeling the information that is shared by conversational partners. Each new contribution to the interaction is interpreted in the context of the shared information and requires an update of this shared information. Principles of 'grounding' [4], shared plans [7] and joint intentions [5] have been introduced. Indeed, the majority of research on modeling human-human interaction and on applying these models for human-computer interaction necessarily assumes cooperation and modeling some mutual understanding about what is being discussed between conversational participants. Some exceptions will be discussed in forthcoming subsections.

As mentioned before, there are many situations where people don't follow the Gricean rules. That may be because of, for example, self interest, indifference, competition, politeness or respect. People may aim, consciously or unconsciously, at keeping the interaction going, rather than on aiming at satisfying the Gricean maxims. And, there are many real-life interaction situations that we want to model in human-computer or human-robot interaction where the computer, the social robot, or the virtual human may disagree with us, consider us an opponent rather than acting as a virtual agent that always agrees with us and performs our wishes, or simply tries to persuade us to do or forget about doing certain things. We can find such situations in electronic commerce applications, 'negotiation' situations, daily life conversations, training and sports situations, and game and entertainment situations. These situations will be discussed next.

2.2 Conversations and Dialogues

When we look at existing conversational systems it is clear that users have to adapt to the system and no natural conversation can take place. Users of these systems are assumed to be cooperative and in practice also start with showing empathy to the system. Unfortunately, an artificial system is not able to handle this helpfulness, it gets confused, and it bothers the user with its stupidity. In current Question-Answering (Q&A) systems the computer has the initiative and the user has to adapt to a system with poor language understanding, poor speech understanding, and no interaction understanding. Let alone any understanding by the system of context, affective

content and user preferences. Nevertheless, task-oriented Q&A systems can be presented in such a way that a user adapts in a rather natural way to the limitations of a system. Current limitations are not knowing about the user, the task, the context, and the preferences of a user. In fact, not knowing about everything a user finds important, except for that one particular question the user is asking and that probably gets misunderstood by the system since the user does not know about the way the system expects him or her to formulate a question. As an extreme example consider the ALICE conversational system. ALICE is a general conversational system that is well-known and has been rewarded for its performance, while hardly any scientific natural language dialogue research has been included in the system. On the other hand, because of their task and domain-dependent nature, no existing dialogue system can perform as well as ALICE in a conversational setting. Let's look at an example of an Alice dialogue that came up when we researched the possibility of a virtual diary companion for soldiers who performed a mission in Afghanistan. The virtual diary companion was designed to provide empathic support to soldiers, stimulating them to express their feelings while making entries in the diary [16]. While looking at its possible performance we also looked at the possibility of using ALICE for this task, but stopped considering ALICE after seeing the following conversation:

Human: I've had a terrible day.
ALICE: What happened to it?
Human: First we were under attack.
ALICE: And then what happened? And second?
Human: Then our vehicle struck a landmine.
ALICE: Fascinating. It would be nice if our vehicle struck a landmine.
Human: Then my best friend got killed.
ALICE: Good reasoning. How well do you know this person?

Despite the low-quality speech and language processing technology that we have today, it is assumed, and at least topic of research, that in the future there is the possibility to have dialogue and conversational systems that will be used and will be appreciated by their users. This may be expected because the system learns to know about the user from his or her interaction history including preferences, interests, and attitudes. Moreover, apart from speech and language, there are other modalities that provide information from user to system. Gestures, facial expressions, body postures, gaze behavior, and other information that can be measured (e.g., (neuro)physiological information) can be used by a conversational system to keep the conversation going or to provide a dialogue system with the information to provide relevant feedback. It should be mentioned that in several human-human interaction studies it has been shown that there are situations where nonverbal cues are more important in conveying a speaker's message than the meaning of the words that are used. This is in particular true when attitudes or feelings are communicated. Cf. Mehrabian [13] who introduced the following rule: "Total Liking = 7% Verbal Liking + 38% Vocal Liking + 55% Facial Liking." In case of incongruities appearing in the modalities with which a message is communicated, the nonverbal modalities seem to be the most important.

Consider the progress that is being made to make a computer a conversational partner. The computer can be represented as a virtual friend that knows about us and to whom we can talk in a confidential way. It may be a virtual butler that also knows about us and maybe even knows more about us than a real friend, but to whom we can talk to in a less confidential way. In real life, whoever we talk to, we don't display all our feelings or are explicit about all our goals. We don't provide all information we have. Sometimes that is to protect ourselves from unwanted intimacy; sometimes it is to protect our conversational partner from information that may be harmful for him or her. We keep back information, we lie, and we manipulate. There are studies that tell us how many 'lies' we are using every day. In many cases these 'lies' are functional. They are not that important and they keep the conversation going. When we have useful applications for virtual conversational partners, do we always want them to be completely honest and really mean what they say during a conversation? Wouldn't that lead to very short conversations, while having a long and entertaining conversation should be the goal of the artificial conversational system? We think that our artificial conversational partner should be able to get the conversation going even when it is able to recognize consciously produced incomplete, incorrect or misleading information. An important question is whether it should show that it knows that the user is providing incomplete or misleading information, whether it should take advantage of it, or whether it should stay silent and act as a virtual butler that uses such information to better serve his or her human master.

Poggi et al [19] studied deception. They introduce a typology of ways to deceive and they discuss the detection of deception in a multimodal, nonverbal communication setting. This framework is used to design virtual humans or embodied conversational agents that are able to deceive about their emotional state by using deceptive facial expressions. That is, they can fake, mask or inhibit expressions. A virtual human that is able to deceive about her emotional state by using deceptive facial and bodily expressions can play a useful role in conversations, negotiations, or other ways of (nonverbal) interaction (see also the next sections).

Hence, we can conclude that in conversational settings, that is, in a setting where a virtual human or a social robot is used as a conversational partner, there are good reasons to have this artificial partner knowingly accepting that its human partner is not necessarily following the Gricean principles and adapts to a verbal and nonverbal exchange where it does not follow these rules itself (e.g. by displaying deceptive nonverbal behavior, rather than what could be considered spontaneous behavior). Detecting that a human conversational partner is not following these rules (consciously or unconsciously) is becoming possible by technology that senses all kinds of non-verbal communication information (from speech, gaze, head and body movements, and physiological information, including brain and muscle activity measurements).

2.3 Commerce, Negotiation, and Persuasion

During daily conversations we are not always honest. These conversations do not necessarily have a particular aim. As discussed, one aim is to keep the conversation going. The situation is different when we consider verbal and nonverbal interaction between a human and an (embodied) agent in an electronic commerce setting. The aim of the agent, reflecting the aim of its designers and owners, can simply be to sell as many products and services as possible. Such an agent will not follow the Gricean cooperative principle. Neither will an agent that participates in an online auction or an agent that is meant to persuade a citizen to behave in a certain way. These agents take a certain perspective in their interaction and do not necessarily provide fully complete or fully correct information. They are not necessarily sincere. In their interaction with other agents or humans they have to decide when and how to be honest and when and how to deceive and when and how to hide information.

In [3] it is argued that “Agents are and will be designed, selected or trained to deceive, and people will be deceived by and will deceive their own agents.” For example, in multi-agent systems agents can be self-interested and if necessary will compete with other agents in order to achieve a task assigned to it by its user. In an auction an agent will not be honest about the price it is willing to pay if there is a chance to get a product cheaper. Even a personal assistant agent can decide to deceive its ‘owner’ or conceal certain information because it knows more about, among other things, legal consequences of actions, consequences for long-term goals and preferences that a user has, or consequences for health. Castelfranchi mentions various reasons why agents will deceive us and each other: for protecting secrets and privacy, for courtesy, for tutorial persuasion, and for protecting collective interests. Some of these reasons are ‘pro-social’ and can be seen as altruistic. Lying can be ethically justified and is not necessarily immoral.

Clearly, similar to what we mentioned at the end of section 2.2, the agents need to be aware of their behavior and need to be able to make decisions about appropriate behavior. In the case of embodied agents or virtual humans their nonverbal behavior should be chosen in accordance to what they want to achieve.

2.4 Teaching, Training, and Serious Games

Computing intelligence and computing power can be embedded in a virtual teacher or a teaching environment. Teachers do not always act in an explicit cooperative way. It can be useful to provoke, challenge, or tease a student. It can be useful to use humor, to play the role of a non-understanding conversational partner and to display faked emotions. At the same time, a student interacting with a teacher or a virtual teacher has a strategy that aims at getting a good assessment of his or her knowledge and motivation. The computer-controlled virtual teacher needs to be aware of this. The student is not necessarily aware of the strategies of a human teacher or the strategies that have been included in a virtual educational environment and an embodied virtual teacher. Neither the student, nor the teacher is playing according to the Gricean rules. There is nothing wrong with that, but in order to act natural and to be effective, a virtual teacher or teaching environment should be able to detect, analyze and synthesize such behavior in order to generate understanding and empathetic behavior in a virtual agent in order to take care of natural face-to-face interaction. Choosing between strategies has been the topic of research of our ‘Ines’ project [8]. In this project we have a student performing a nursing task on a virtual patient while being monitored by a virtual teacher, commenting on the performance of the student and trying to motivate the student by providing appropriate comments on the progress that is being made. Clearly, such comments are meant to play a role in the teaching process and do not necessarily reflect the opinions or the emotions of the teacher.

Nowadays virtual reality environments are being used for teaching and training situations. In these environments events are simulated and trainees can ‘enter’ these environments in order to learn to collaborate with human or virtual team mates, enter into situations where they have to negotiate with human or virtual partners, or enter situations where they have to fight human or virtual opponents. An example of such a ‘serious’ or role-playing game is the virtual human doctor project [24,25]. The setting is a clinic somewhere in Iraq. The trainee is an army captain who has to persuade a doctor to move his clinic because of a planned military operation. This has to be done without revealing details of the military operation. Obviously, and being part of the training situation, the doctor is not necessarily cooperative. This requires the modeling of non-cooperative behavior. In [24] various factors have been distinguished and some of them have been implemented in such a virtual doctor. They include unilateral topic shifts or topic maintenance, avoidance, competition, unhelpful criticism, withholding of information, lying and deception, antagonism, and rejection of empathy. More globally, mentioned in [26,25], three orientations toward a negotiation can be distinguished: avoidance, distributive, and integrative. Here, ‘avoidance’ means trying to avoid the negotiation, denying the need for it. The ‘distributive’ orientation assumes that there is a winner and a loser, while in the ‘integrative’ orientation there can be a win-win situation. Each orientation needs different strategies. In the case of the virtual human doctor one can have the situation that there is ‘avoidance’ first, but it can change to the ‘distributive’ and after that to an ‘integrative’ orientation. In this research strategies are chosen based on an appraisal model of emotion that looks at these orientation strategies as types of coping strategies.

At the end of section 2.2 and 2.3 we mentioned the need for virtual conversational agents that can process non-Gricean behavior and that can behave themselves in a non-Gricean way. The underlying assumption was that we need such agents in order to model natural interactions between humans and virtual agents (virtual humans or social robots). Now, also at the end of this section we can mention that there are application areas for virtual humans (and social robots) where they have to play roles in simulated training or teaching environments, where they have to negotiate, and where they have to persuade, in both verbal and nonverbal ways. Being non-cooperative and displaying deceptive expressions – as was studied in Poggi et al. [19] – are part of a virtual human’s behavior in these situations.

2.5 Sports, Games and Entertainment

Presently we see research and the development of technology that aim at providing exercise and training environments. Sensors are able to detect human activity in rather detailed ways. Microphones detect speech and sound, cameras detect movements of the body, the limbs and changes of facial expressions, there are sensors that detect positions and proximities, and physiological sensors provide information about body and mental state of a user of these exercise or training environments. These environments aim at improving the health of their users, for example by displaying a motivating virtual environment, a virtual coach, and a fitness exercise program. Interpretation of the information obtained from the sensors allows the environment, probably represented by a virtual human, to match actual behavior with desired behavior, and to adapt its appearance and its feedback strategies to the performance of the user. In our research we looked at the modeling of a fitness trainer [21]. One of the things that we noticed is that a trainer needs to be aware that a user is not necessarily honest in his or her verbal or nonverbal attitude towards a trainer. He or she can hide fatigue or exaggerate fatigue. The virtual trainer has the possibility to know about this and has to decide how to deal with this. This includes deciding whether the trainer's knowledge about a user's deceptive behavior should be communicated to the user. It is not always in the interest of the user or the trainer to speak the truth.

Fitness training using virtual environments and interactive virtual trainers receive much attention. A virtual trainer can monitor actual exercise performance and can match it with desired exercise performance. Other virtual environments are introduced to simulate situations that allow sportsmen to train and improve their performance during a real sports game. Virtual environments have been designed to simulate ice hockey, handball, boxing, or rugby situations. Exertion interfaces have been introduced that offer new sports or games that require physical efforts and 'whole body interaction' by the user of such interfaces. In these environments we have trainers, team mates, and opponents.

In sports and games deceptive actions are part of the game. They are meant to divert attention from one's real purpose. Hence, in virtual training and recreational environments a trainer or in particular a game opponent is not only allowed, but also expected to have nonverbal behavior that is aimed at deception. Just to mention an example, suppose we have a virtual fencing trainer. Its main job will be to exercise recognizing and generating deception behavior. Similarly, we can look at virtual or mediated boxers [9,18], baseball players [11] or rugby players [2]. Training environments have been built where a human handball keeper or a human rugby player has to deal with feinted attacks.

Again here, similar to the conclusions of the previous subsections, in these situations we need agents that have underlying models of not being cooperative and models that aim at nonverbal misleading an interaction partner.

2.6 Not Trusting the Computer

Clearly, in many of the situations described above, the computer, or rather how it appears to us in our interaction activities, can interact with us in playful, exercise, entertainment, sports, and serious gaming environments. In these environments we can expect that situations we can expect that non-cooperative and deceiving behavior is there. It is part of a game, it is part of a training, and it is part of an exercise programme. It may be the case, and it was an essential theme in Stanley Kubrick's movie 2001, that we simply do not trust an advice or a decision made by an extremely intelligent computer and that we verbally and nonverbally try to deceive this computer, assuming that we know better. In a well-known fragment of this movie one of the astronauts (Dave Bowman) takes the decision to hide his suspicion that HAL, the intelligent computer, is not able to handle a particular dangerous situation. Or, at least, not willing to handle this situation in the (life-saving) interest of the astronauts. Dave decides to discuss this situation with his co-pilot, but is not aware that HAL has eyes everywhere and is aware of this discussion. Later, trying to convince HAL to adapt the mission's aims, HAL is able to confront Dave with this overheard discussion and refuses to make any changes to the mission. Nevertheless, Dave's empathy, trying to understand HAL's way of feeling and reasoning, turns out to be stronger than HAL's understanding of Dave's intentions. The '2001' movie is science-fiction, but nevertheless. The discussion between Dave and HAL is about trust, mistrust and assuming that your conversational partner's aim has interests others than your and tries to deceive you.

HAL: This mission is too important for me to allow you to jeopardize it.

DAVE BOWMAN: I don't know what you're talking about, HAL?

HAL: I know you and Frank were planning to disconnect me, and I'm afraid that's something I cannot allow to happen.

DAVE BOWMAN: Where the hell did you get that idea, HAL?

HAL got that idea by observing Dave and Frank discussing how to deal with him while assuming their conversation was hidden from artificial eyes and ears sensors in the environment. Wrong idea, HAL knew.

Clearly, here, we cannot say what has to be done. Do we want to negotiate with the computer, do we want to compromise, or do we want to overrule the computer whatever his arguments are? Or is it up to the computer to choose among these alternatives? Whatever we choose, interaction models aware of different perspectives, different aims and different truths need to be designed.

2.7 On Being Believable

There is another issue that needs to be addressed. Whenever we introduce a humanoid (virtual human or a social robot) to engage in an interaction with a human interaction partner, in today's research our aim is to keep the interaction going, making the interaction believable and making the virtual character or the social robot believable

in its interaction with the human conversational partner. This aim can be traced back to the Turing Test. This test was introduced to measure progress in artificial intelligence. The test looked at intelligent expressed in a verbal way. When we have a human conversational partner and a computer communicating with a human 'referee', can this referee decide whether he has been communicating with the computer or with the human conversational partner? This test does allow all kinds of cheating from the part of the computer agent. Everything is allowed, as long as the human conversational partner accepts the verbal and nonverbal contributions of the computer agent as believable (for some time). In section 2.2 we displayed a conversation between the ALICE-bot and a human partner, and, clearly, the interaction is far from believable. All kinds of research attempts aim at making embodied conversational agents (virtual humans/social robots) more believably. This certainly includes modeling agents in such a way that they attempt to make their human conversational partners believe that they are more intelligent and more social than they really are. There are several strategies that can be used: displaying empathic and intelligent listening behavior, avoiding certain discussion topics, or changing the topic of discussion. In [20] we discuss behaviors and appearances of virtual humans in settings where there is not necessarily cooperation, trust, or empathic behavior.

3 Conclusions and Future Work

There are many reasons why we need to deal with deceptive verbal and nonverbal interaction. In this paper we looked at natural conversations and why such conversations profit from not always showing true feelings, we looked at commerce and negotiation situation where users are not assumed to show their feelings, we looked at game, training and simulation environments where users have to compete, obey and adjust their behavior to demands and preferences of their coaches, their team mates, and their virtual opponents. In all these situations some modeling of non-cooperative behavior, some modeling of empathic behavior, some modeling of persuasive behavior, and some modeling of coaching or teaching behavior is required. In games we see research attempts to make the 'non-playing' characters more autonomous by providing them with intelligence and social behavior. Clearly, when this is done, these characters need to know about competition, disagreement, aggressiveness and violence. There is also discussion about bringing games more into the real world, that is, it is expected that in the future more competitive situations will be designed in the real world that allow playful deception [22].

Acknowledgements This research has been supported by the GATE project, funded by the Netherlands Organization for Scientific Research (NWO) and the Netherlands ICT Research and Innovation Authority (ICT Regie).

References

- [1] J. Allwood. Linguistic Communication as Action and Cooperation. Ph.D. thesis, Göteborg University, Department of Linguistics, 1976.
- [2] S. Brault, B. Bideau, R. Kulpa, and C. Craig. How the global body displacement of a rugby player can be used to detect deceptive movement in 1 vs. 1. Proceedings of the 11th Virtual Reality International Conference, April 22-24, 2009, Laval-France, 161-166.
- [3] C. Castelfranchi. Artificial liars: Why computers will (necessarily) deceive us and each other? Ethics and Information Technology 2, Kluwer Academic Publishers, 2000, 113–119.
- [4] H.H. Clark and D. Wilkes-Gibb. Referring as a collaborative process. *Cognition*, 22 (1986) 1-39.
- [5] P.R. Cohen and H.J. Levesque. Confirmations and joint action. In: Proceedings IJCAI-91, 1991, 951-957.
- [6] H.P. Grice: Logic and Conversation. In: Syntax and Semantics, Vol. 3, Speech Acts, ed. by P. Cole and J.L. Morgan. New York: Academic Press 1975, 41–58.
- [7] B. J. Grosz and C. L. Sidner Plans for discourse In P. R. Cohen, J. L. Morgan, and M. E. Pollack (Eds.), Intentions in Communication. MIT Press, Cambridge, MA, 1990, 417-444.
- [8] D. Heylen, A. Nijholt & R. op den Akker. Affect in Tutoring Dialogues. *Journal of Applied Artificial Intelligence* (special issue on Educational Agents - Beyond Virtual Tutors), Vol. 19, No. 3-4, Taylor & Francis, 2005, 287-311.
- [9] J. Höysniemi, A. Aula, P. Auvinen, J. Hännikäinen, and P. Hämäläinen. Shadow boxer: a physically interactive fitness game. Third Nordic Conference on Human-Computer interaction (NordiCHI 2004), Vol. 82, ACM Press, New York, 2004, 389–392.
- [10] IST Advisory Group (ISTAG). Ambient Intelligence: from Vision to Reality. In Riva, G., Vatalaro, F., Davide, F., Alcañiz, M. (Eds.) Chapter 4 in Ambient Intelligence. IOS Press, 2005, 45-68.
- [11] T. Komura, A. Kuroda, and Y. Shinagawa. NiceMeetVR: facing professional baseball pitchers in the virtual batting cage. ACM Symposium on Applied Computing, 2002, 1060–1065.
- [12] G. Leech. Principles of Pragmatics. Longman, London, 1983.
- [13] A. Mehrabian. *Silent Messages* (1st ed.). Belmont, CA: Wadsworth, 1971.
- [14] A. Nijholt. Conversational Agents and the Construction of Humorous Acts. In: Conversational Informatics: An Engineering Approach. T. Nishida (Ed.), John Wiley & Sons, Chicester, England, 2007, 21-47.
- [15] A. Nijholt. Google Home: Experience, Support and Re-Experience of Social Home Activities. *Information Sciences. Special Issue on Ambient Intelligence. Vol 178/3, Elsevier, Amsterdam, 2008, 612-630.*
- [16] A. Nijholt, F. Meijerink, and P.-P. Maanen. A virtual diary companion. In: Fourth International Workshop on Human-Computer Conversation, 6-7 October 2008, Bellagio, Italy, 1-5.

- [17] A. Nijholt. Don't Give Yourself Away: Cooperation Revisited. In: Logic and the Simulation of Interaction and Reasoning Symposium at the AISB 2008 Convention "Communication, Interaction and Social Intelligence", 3-4 April 2008, Aberdeen, UK. The Society for the Study of Artificial Intelligence and Simulation of Behaviour, 2008, 41-46.
- [18] J.Y. Park and J.H. Yi. Gesture Recognition Based Interactive Boxing Game. *International Journal of Information Technology* 12(7), 2006, 36-44.
- [19] I. Poggi, R. Niewiadomski, and C. Pelachaud. Facial Deception in Humans and ECAs. *Modelling Communication. Lecture Notes in Artificial Intelligence* 4930, Springer-Verlag, Berlin, 2008, 198-221.
- [20] D. Reidsma, Z.M. Ruttkay, and A. Nijholt. Challenges for Virtual Humans in Human Computing. In: *AI for Human Computing. Lecture Notes in Artificial Intelligence* 4451. Springer Verlag, Berlin, 2007, 316-338.
- [21] Z. Ruttkay and H. van Welbergen. Elbows Higher! Performing, Observing and Correcting Exercises by a Virtual Trainer. In: *Intelligent Virtual Agents. Lecture Notes in Computer Science* 5208, Springer-Verlag, Berlin, 2008, 409-416.
- [22] J. Schell. When games invade real life. DICE 2010 presentation. <http://g4tv.com/videos/44277/dice-2010-design-outside-the-box-presentation/>
- [23] J.R. Searle. Logic and conversation. In: *Syntax and Semantics, Vol. 3: Speech Acts*. P. Cole and J. L. Morgan (Eds.), Academic Press, 1975, 59-82.
- [24] D. Traum, W. Swartout, S. Marsella, and J. Gratch. Fight, Flight, or Negotiate: Believable Strategies for Conversing Under Crisis. *Intelligent Virtual Agents (IVA 2005), LNCS 3661*, Springer-Verlag, Berlin, 2005, 52-64.
- [25] D. Traum. Computational Models of Non-cooperative dialogue. Abstract of invited talk at LONDIAL 2008 Workshop on Semantics and Pragmatics of Dialogue, May 2008.
- [26] R.E. Walton, and R.B. Mckersie. *A behavioral theory of labor negotiations: An analysis of a social interaction system*. McGraw-Hill, 1965.