Don't Give Yourself Away: Cooperation Revisited

Anton Nijholt[1]

Abstract. In the future, sensor-equipped environments will be able to detect, interpret and anticipate our intentions and feelings. While on one hand this allows more natural interaction between humans and human activity supporting intelligent environments, on the other hand it allows these environments to collect more information about the user than he or she finds desirable for a satisfactory interaction. That is, there are many human-human interaction situations where it is quite acceptable or even necessary that part of the intentions and feelings of one conversational partner remains hidden for the other. We will discuss research on ambient intelligence and human-computer interaction that allows us to introduce and discuss this problem and we illustrate our viewpoints with examples from our own research on virtual humans interacting with human partners that act, behave, and perform in environments equipped with sensors that capture and interpret human behaviour.

1 INTRODUCTION

Most of our research in human-computer interaction assumes that humans and computers cooperate. And although there is research on adaptive interfaces, most of the time the user has to adapt to the interface by using rather unnatural devices, follow interaction protocols, speak clearly, etcetera. Here we explore human-computer interaction where there is not necessarily cooperation and where it may be in the interest of the user to hide his intentions or feelings. When we talk about interaction then we don't limit ourselves to verbal interaction. On the contrary, in what follows we assume situations and environments where all modalities that can be displayed (movements of body parts, posture and gestures, facial gestures, speech, gaze, (neuro-) physiological) can be observed by the interacting partners.

2 NO DESIRE TO BE FORTHCOMING OR COOPERATIVE

People often hide their feelings, they often hide their thoughts, and they often hide information. People often behave differently depending on when they are alone or when others are around. People sometimes want to hide from others; they are not always in need of an audience, bystanders or partners.

People have their interests and preferences. Depending on them, and their personality and their mood, they voluntary or involuntary give away part of themselves during interactions. People do not always want to be forthcoming. Moreover, they play roles. Implicit or explicit decisions are made about the roles they want to play and what they want to disclose, where, when, and how. This is also known as privacy. They also make decisions how much effort they will make in understanding a conversational partner and his or her interactional goals. Also, too much interest from others in our motivations is not appreciated. We don't want other people to read our mind. We are not always interested in reading other people's mind.

Neither is it always in our interest to be cooperative. However, in general, being cooperative, just as being polite, can sometimes help us to get closer to our interactional goal. In a conversation we can flatter our conversational partner, we can purposely misunderstand our partner in order to be able to make a humorous remark, and we can play the devil's advocate, and nevertheless be cooperative. We play along with the rules of a conversation or a negotiation and therefore we are cooperative despite possible elements of competitiveness. In these situations Grice's maxims on cooperation, i.e. assumptions a listener is supposed to have about the interaction behaviour of a speaker, seem to be violated, but the relevance of the behaviour can be explained from a pragmatic, conversational point of view, rather than from a sentence level point of view. Conversational partners can achieve their goals although they can have competitive interests. To achieve these goals it is acceptable that people hide their feelings and intentions. Moreover, it is quite acceptable that they tell lies.

Clearly, lies don't follow Grice's maxims. People don't follow Grice's maxims, since all people sometimes lie in everyday conversations. They say things they don't mean, they tell self-oriented lies, i.e., lies for one's own benefit, or other-oriented lies, i.e., for the benefit of others. Social lies are meant to benefit relationships. Lies act as a social lubricant [1]. We don't always want to speak the truth; we don't always want to hear the truth. And sometimes we don't want to find out the truth. Lies are in the interest of both conversational partners in an interaction. In addition to telling the truth, during a row people will also exaggerate or say things they don't mean. Lies can be nonverbal. We can nonverbally feign surprise or sincerity. We can pretend to be happy by our facial expression; we can pretend to be rich by our clothes.

Interactions can be completely nonverbal or nonverbal supported by speech. Consider for example, two partners coordinating their behavior during dancing, students coordinating their movements with a fitness trainer or a sports instructor, and a conductor conducting an orchestra. Clearly, in sports (e.g., baseball, soccer, tennis, …) and in games misleading your opponent by feigning movements, introducing sudden interruptions of movements or changes in behavior is essential and is part of the entertainment or sports experience. Nonverbal humor is an example of friendly misleading your partner, e.g. during disco dancing by feigning certain movements.

[1] University of Twente, Human Media Interaction, PO Box 217, 7500 AE, Enschede, the Netherlands. anijholt@cs.utwente.nl

In all these observations the nonverbal characteristics of the interaction are extremely important. Even during conversations these characteristics can say more about what people mean or want others to believe than the content of the verbal utterances. People can read these nonverbal characteristics and they can be misled by these nonverbal characteristics in their interactions.

Obviously, there is not necessarily a balance between capabilities of conversational partners. Partners differ in background, knowledge, attitudes and personality. A partner can be more determined to reach a certain goal, a partner can have more social intelligence and be able to read the mind of its human opponent better than he or she is able to do. Not all partners in interactions have equal means.

3 DISAPPEARING COMPUTERS AND INTERFACES

Human-computer interaction is about designing interfaces between humans and computers. Before there were Personal Computers for most of the computer users the interface consisted of a counter where you could deliver your program punched on tape or cards. After that end users were allowed real-time remote access to computers using terminals that allowed the composing and editing of programs. Personal Computers provided end users with facilities to interact with software designed by software companies or by themselves. User interface software mediated between users and application and system software. Graphical user interfaces aimed at efficient and user-friendly interaction. Interface technologies now include speech and language input, haptic input, and vision input. Moreover, in addition to professional applications where efficiency is important and necessary, home and recreational computer use became extremely important and these applications require interfaces where there is a user access layer where user friendliness, ease of learning, adaptiveness, and fun to use are the main design issues, rather than efficiency considerations that appear in levels below.

As mentioned, interface technologies now include speech and language input, haptic input, and vision input. But there is more. Since we can have sensors embedded in the environment, including walls, furniture, devices, robots and pets, now the environment has become intelligent and it can perform not only reactive, but also pro-active behaviour, trying to anticipate what the inhabitant is doing and doing this by perceiving activities and all kinds of verbal and nonverbal behaviour. Embedded sensors include cameras, microphones, location and movement sensors, and sensors that collect and distinguish various types of physiological information and brain activity patterns. Information about the behaviour of the inhabitants and their implicit and explicit addressing of the environment can be fused and interpreted in order to support the inhabitants by providing appropriate feedback.

This research, in particular when research activities into social and intelligent interfaces are included, has become known as ambient intelligence research. Well-known is also Mark Weiser's vision of disappearing computers ("The most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it.") [2] and the associated question how we can design implicit and explicit interaction, with multiple human inhabitants, for sensor-based interfaces [3]. In our view [4] in these environments humanoids and pet-like devices can play a useful role in observing inhabitants and interacting with them. Agent-modelled virtual humans, (mobile) robots, pets (virtual or robotic animals) can have specific human-oriented tasks in smart environments (e.g., be a friend, assist in cooking, take care of house security, retrieve information, assist in health-care and fitness, be an opponent in games or sports), and they can represent human beings (e.g., family members that are away from home but have there whereabouts or activities visualized). They can also communicate with each other, distributing, sharing and integrating their experiences, and learning about the inhabitants and visitors of these environments, making them transparent, independent of the role they play. Being able to capture and interprete what is going on Capturing and interpreting of events and activities, e.g. in home environments, allows future retrieval and replay of events and experiences in, for instance, 3D virtual environments [5].

4 NOT GIVING AWAY YOUR INTENTIONS OR FEELINGS

The main topic of discussion in this paper is not about privacy or privacy protection. But clearly, when we talk about not giving away intentions and feelings, there is the underlying assumption that you are acting and behaving in a world (ambient intelligence environment, smart environment) that is inhabited by agents that perceive your acts and behaviour and that may profit from the knowledge that is obtained in that way. And, they may have goals that do not necessarily match yours, maybe on the contrary. Hence, such agents may collect information that threatens your privacy by using this information without you being aware of it, without your consent, and sometimes against you. We can not expect that we can always choose our own agents. The providers of our smart environments will make choices, they have commercial interests and they have to follow political burps and fluctuating political and societal demands. Where can we hide from those who provide our environments with sensors and communicating intelligent agents that allow the environment to pretend social and intelligent awareness while collecting information during the interactions with us [6]?

Their agents will have their own interests, or rather their owners' interests, and therefore they will persuade us to adapt to their goals. As mentioned in [7], there will even be situations where our virtual personal assistant will take a decision to deceive us (in our 'own interest', of course) and also situations where we will want to deceive our personal assistant.

In our research we don't look into details of these issues. Rather we look at human behaviour during natural conversations and other interactions and the reasons to hide information, i.e., not display ourselves, not to be forthcoming or, even, wanting to mislead our (artificial) interaction partner and provide him or her with information that is not necessarily complete or true.

In future ambient intelligence environments, are we still able to provide our conversational partners with incomplete and sometimes wrong information about ourselves, our intentions and our feelings just as we are able to do and are used to do, and probably for good reasons, in real-life situations nowadays? This question can be asked for real-time face-to (virtual) face conversational interactions, but there are also many other situations where it is rather natural not to show everything we know or everything we feel. In our research on continuous interaction modelling (in contrast to 'turn-taking' interaction [8]) we designed and implemented applications where it turned out that humans in their interactions with embodied partners felt that there sometimes were advantages in not displaying their intentions and feelings, and, also the other way around. This became clear in our research on a so-called Sensitive Artificial Listener (SAL), developed in the framework of an EU FP6 Network of Excellence on the role of emotions in the interface, in which we participated [9], in our research on an interactive virtual dancer [10], an interactive virtual conductor [11], an interactive virtual fitness trainer [12], and in our research on an educational virtual environment for nurse education [13].

In these applications both the nonverbal interaction behaviour and the fact that during interactions all conversational partners continuously display nonverbal interaction behaviour [14], made clear that continuously decisions are being made about what a human partner would like to become displayed to a virtual interactional partner, and the other way around. Examples that emerged in our applications are: using humour to temporarily mislead a conversational partner, not being sincere by feigning interest in a conversation, not yet wanting to show your fatigue to your fitness trainer or colleagues, feigning movements during a dance interaction, and feigning movements in a virtual reality entertainment game.

5 HUMAN COMPUTING TECHNOLOGIES

Many authors have discussed smart environments, ambient intelligence, ubiquitous computing or pervasive computing. Mostly their starting point is the technology that makes it possible to embed environments with sensors, intelligent sensors, and communicating sensors. Our starting point is the support of human beings in their activities in smart environments. That is, rather than in traditional computer science where the assumption is that human beings should add to the tasks that a system has to perform, we investigate how sensor-equipped environments can support human beings in their daily home and professional activities, how they can contribute to their well-being, and how they can contribute to their curiosity to explore new areas of interest. For that reason we need to look at anticipatory user interfaces that should be human-centred and should be built for humans based on human models.

Our natural environments, whether they are home, office, or public space environments, become smart. They are being equipped with sensors, where the sensors themselves can be intelligent, but, more importantly, where the information that is obtained from various kinds of sensors is fused and interpreted, allowing relevant feedback generation from the environment and its devices to the inhabitants of the environment. Tangible objects, robots, virtual pets, furniture, walls, doors, virtual humans displayed on screens, etc., may re-actively and pro-actively provide support to real humans in their daily professional, home, and recreational activities. While doing this, these environments need to collect information about the user: his activities, his intentions, his moods, and his emotions. We discuss three viewpoints related to the ambient intelligence approach to computing and supporting inhabitants of ambient intelligence environments.

The first viewpoint in our research concerns the level of sophistication at which the environment understands what is going on and is able to provide support [15]. Depending on this level of sophistication models are needed that range from simple rules, such as, turn the light on when someone enters the room, to cognitive user models that are used by the environment to understand, e.g., why a person is choosing a particular chair around a table in a meeting room.

The second viewpoint concerns the way the environment and its devices are asked to use the embedded knowledge or decide to use this embedded knowledge in support of the activities of the environment's inhabitants or visitors. For example, in [16] we find an interaction space divided by two axes. One axis ranges from re-active to pro-active behavior, while the other axis covers the range from low to high level attention of these sensor-equipped smart environments. The main message here is that we can design intelligence in the interface and that due to this intelligence, in particular intelligence that is fed from observations from the environment, the interface can evoke behavior that (1) hides processing from the user, (2) can provide feedback to questions or actions requiring feedback, (3) can anticipate user's actions, and (4) can enter and entertain the interaction with other users.

The third viewpoint is related to the research methodology. Our research starts with observing how people behave and interact. For example, in the AMI and AMIDA project [15] we have a large collection of data concerning people interacting with each other. This data is annotated using annotation tools and annotation schemes. Tools and schemes are not perfect. However, machine learning technologies can be used to improve in an iterative way, analysis and modelling of (multi-party) interaction.

6 (NON-) COOPERATIVE BEHAVIOR BY USERS, PARTNERS, AND OPPONENTS

Our research is on natural cooperative and uncooperative behaviour. Clearly, we look at Grice's maxims, but we also distinguish situations where people may prefer not too be fully informative, not to display all relevant information, and maybe even prefer to mislead their interaction partner. It is not always in your interest to open yourself to a conversational partner. Moreover, interactions can become much more interesting and useful when such conversational rules are neglected or violated. This is certainly the case in applications where your interests forces you to disagree with your partner (or rather opponent), for example in discussions, games, or sports. Clearly, we are not talking about verbal interaction only. In our research we include nonverbal interaction and also all kinds of other activities that can be perceived by our interaction partners (or opponents).

It should be mentioned that there is a friction that emerges when on the one hand our smart environments and processing technologies not only allow, but also invite natural interaction behaviour, while on the other hand the processing technologies become able to extract more information about our intentions and feelings from this natural interaction behaviour than we would like to become known in a natural human-human interaction situation. How to deal with partners that have not necessarily been designed to help us, how to deal with partners, e.g. in games and sports that are opponents rather than friends? In the remaindrr of this paper we will in particular look at entertainment and sports applications. Therefore, in the next section, we will look at such applications and introduce the so-called exertion interfaces.

7 DANCES, SPORTS, GAMES, FITNESS

Entertainment, health, sports, and leisure applications using information and communication technology often require and encourage physical body movements and often applications are designed for that reason. In our research we look at bodily and gestural interaction with game and leisure environments that are equipped with sensors (cameras, microphones, touch, and proximity sensors) and application-dependent intelligence (allowing reactive and proactive activity). Interpretation of the bodily interaction, requiring domain-dependent artificial intelligence, needs to be done by the environment and the agents that maintain the interaction with the human partner. In the display of reactive and pro-active activity embodied virtual agents play an important role. Virtual agents can play the role of teacher, coach, partner or buddy. One underlying assumption is that emphasis on activities in which the experience rather than the result will be used to guide the design of social and intelligent systems that will become part of ambient intelligence home environments [5].

Figure 1. Sports over a Distance

Bodily activity, to be captured by cameras, microphones, pressure and location sensors, has been considered for many applications related to sports, games, entertainment, fitness, and education. Hence, there is a virtual therapist that helps patients to recover from injuries [17], an exercise bicycle as user interface to a virtual environment [18], a Tai Chi training master [19] and a shadow boxer [20] that acts as a fitness environment to help prevent neck and shoulder pain, or a Kick Ass Kung-fu system where the children use Kung-fu to fight virtual enemies displayed on a screen [21].

Interfaces to such environments have been called physical interfaces or exertion interfaces. Especially for the latter it is assumed that the interaction requires intense physical effort, for example, repeatedly shooting balls against a wall. This latter example has been realized in the “Sports over a Distance” project at the Media Lab in Dublin [22]. In their exertion interface users are connected through a video conference screen. For each of them the remote player is displayed on the screen. The screen is divided into blocks and the players have to strike the blocks with a regular soccer ball in order to score (see Figure 1). Blocks can ‘crack’, ‘break’, and ‘disappear’. Players see the results of the other player and can try to break a block that has been cracked by the other player. A really hard strike will break a block at once. Hence, in the game there is a combination of tactics and intense physical effort. More recently “Airhockey over a Distance” has been introduced. In this approach not only airhockey tables are networked and augmented with video conferencing, but there is also a physical puck that can be shot back and forth between the two connected locations by measuring the intensity of hitting the ‘wall’ and ‘firing’ a puck at the remote location from the position where it hit the wall.

In our HMI Lab we have designed three applications in which our ideas about nonverbal and bodily interaction have been implemented. The applications are illustrated in Figure 2. The implementations are there, but they are certainly not final. We looked at the design, implementation and evaluation of a virtual dancer that invites a visitor to her environment to dance with her, a conductor that guides musicians in its environment to play according the score designed by a composer, and a virtual trainer (e.g. in the role of fitness trainer or physiotherapist) that knows about exercises that need to be performed by a user or patient. In all these applications there is a continuous interaction between embodied agent and its human partner. Moreover, rather than have the more traditional verbal interaction supported by nonverbal communication, here the main interaction that takes place is nonverbal, and speech and language, when present at all, take the supporting role. External signals like music being played can also have a role in addition to the multimodal communication.

It should be mentioned that industry (Sony, Microsoft, Nintendo) have become aware of applications where users move away from keyboard, mouse and joystick. The dance pad, the DanceDanceRevolution (DDR) games and the DDR tournaments are examples, but so are the Sony EyeToy games that use computer vision input and the Nintendo Wii motion-sensitive

Figure 2. Virtual dancer, virtual conductor and virtual fitness trainer

input device. And, finally, Microsoft has investigated the use of a dance pad to allow a user to issue commands to his or her email program [23].

8 PLAYING WITH BEHAVIORAL INFORMATION

In ambient intelligence environments we have the technology to capture human behavior in everyday life. In our ambient entertainment view the same technology is available and we can either assume that a particular user or visitor of our ambient entertainment environment already carries a user profile that has been generated from the user's behavior in the past, or we can assume that during a possibly playful interaction with the environment a profile can be obtained and can be used by the environment to adapt to the user's characteristics (for example, personality, preferences, mood and capabilities).

What can we learn about the user when we can observe his or her behavior during some period of time? What can we learn from behavioral information captured by cameras, microphones and other types of sensors? In [24] results are reported from short observations of expressive behavior. Observations include the assessment of relationships, distinguishing anxious and depressed people from normal people, predicting a judges' expectations for a trial outcome, determining political views of television newscasters, et cetera. Personality judgments from 'thin slices of behavior' and their accuracy are also discussed in [25].

An example where real-time behavioral analysis is done by a computer can be found in [26]. In their approach a participant is invited in front of a video camera for about 30 seconds. At the end of this period a personality profile for the earlier mentioned Big Five personality traits will be generated.

Taking into Account Involuntary User Responses

In the examples mentioned in section 7 we have bodily interaction with the computer system. Input to an entertainment environment can be based on conscious decisions made by the human. This is usually the case when keyboard, mouse or joystick is used. In the examples we have body and hand gestures, changes of position, etc., to 'control' the system and to perform a certain entertaining task.

Information collected in a user profile, possibly obtained with the methods discussed in section 3 and 4, can be used to adapt an entertaining activity in advance to a particular user and during the activity it helps to anticipate and interpret the user's actions in the environment.

Behavioral signals and patterns during activities provide (additional) information about the tasks that a user wants to perform, the way they should be performed and the user's appreciation of task, performance, and context. Sensing and understanding these signals is an important issue in 'human computing' [3] and it makes human computing an important area of research for entertainment computing. This kind of input is not always consciously provided by a user and is sometimes beyond the control of the user. Behavioral signals also provide information about the affective state of the user and this information is needed to adapt the environment (more or less control by the user, other challenges, etc.) to the user.

More information about the affective state of the user of an entertainment environment can be obtained by collecting and interpreting information obtained from measuring physiological processes and brain activity. Physiological cues are obtained from, for example, respiration, heart rate, pulse, skin temperature and conductance, perspiration, muscle action potentials and blood pressure [27]. Unfortunately, this information can mostly not be obtained unobtrusively.

Finally, we should mention measured brain activity. Again, measuring brain activity, e.g. by using an EEG cap, can provide information about the affective state of the user (frustration, engagement, etc.) and this can be used to dynamically adapt the interface to the user and provide tailored feedback.

User Control of 'Involuntary' Responses

HAL: I'm afraid. I'm afraid, Dave. Dave, my mind is going. I can feel it. I can feel it. My mind is going. There is no question about it. I can feel it. I can feel it. I can feel it. I'm a... fraid.

From: A Space Odyssey, Stanley Kubrick, 1968

Playing against a computer is not fair. The computer knows about our affective state and can decide to communicate this information to our (virtual) opponents or collaborators in the environment who can use it to their advantage. On the other hand, apart from adapting the environment, the computer can also make the human player aware of his or her affective state so that he or she can make an attempt to control it since it can decrease own performance and give away unwanted information to other players in the game.

In games and sports opponents can be misled. We can as well try to mislead or tease our virtual and human partners who perform in a computer-controlled environment. One step further is that we have entertainment games where misleading the computer is essential part of the game. A simple example is playing soccer against a humanoid robot and the robot's aim is to win rather than to offer its human partner an enjoyable experience. In such a situation misleading means for example making feints. But also, for example, trying to look more tired than we really are and all other kinds of misleading behavior. In our virtual dancer installation (section) it happens that human dancers try to tease the virtual dancer by doing something much unexpected and see how she reacts. In other environments we may want to hide our intentions from the computer by controlling our facial expressions (e.g., in a poker game with a computer that can observe us). That is, once we know that our non-human opponent is receptive for our behavioral, physiological or even brain processes, we can try to cheat in order to obtain more satsfaction from the entertainment game. Although game research in this direction is rare, it is well-known that people can learn to control, up to a certain level, these processes. Research and development in brain-computer interfacing and its medical applications makes clear that interesting new types of entertainment will become available in the future [28, 29].

9 CONCLUSIONS

In ambient intelligence environments the environment learns about the user. The environments will be inhabited by, among others, virtual humans and human-lime robots that are there to support us, but also to support other humans making use of these environments. Human computing technology allows the environment to extract information about the user from his behavior and his activities, including his interaction behavior. Our assumption is that for maintaining natural human behavior in such environments we should be able to hide information from the artificial humans in these environments, we should be allowed to lie, and we should have the opportunity to mislead our artificial partners. This is particular true in situations where we play games or do sports with our artificial partners as we illustrated with some examples from our own research. In that case we look at the computer as our opponent, in addition to being a provider of enjoyable experiences. Rather than providing the computer with information about ourselves we then prefer to mislead the computer and hide information about our affective state or even control and manipulate our behavioral, physiological and brain processes so that we consciously can provide the computer with misinformation in order to become the 'winner' in smart entertainment environments.

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

REFERENCES

[1] A. Vrij. Detecting Lies and Deceit: Pitfalls and Opportunities. 2nd Edition, John Wiley & Sons, Ltd, England, 2008.
[2] M. Weiser. The Computer for the Twenty-First Century. Scientific American, Vol. 265, No. 3, 94-104 (1991).
[3] M. Pantic, A. Pentland, A. Nijholt, and T.S. Huang. Human Computing and Machine Understanding of Human Behaviour: A Survey. In: Human Computing. Lecture Notes in Artificial Intelligence 4451, Springer-Verlag Berlin, pp. 47–71 (2007).
[4] A. Nijholt. Where computers disappear, virtual humans appear. Computers & Graphics 28: 467–476 (2004).
[5] A. Nijholt. Google Home: Experience, Support and Re-experience of Social Home Activities. *Information Sciences*. Vol. 178, No. 3, Elsevier, Amsterdam, pp. 612-630 (2008).
[6] A. Nijholt, T. Rist & K. Tuinenbreijer. Lost in ambient intelligence. In: Proceedings ACM Conference on Computer Human Interaction (CHI 2004), ACM, New York, pp. 1725-1726 (2004).
[7] C. Castelfranchi. Artificial Liars: Why Computers Will (Necessarily) Deceive Us and Each Other. Ethics and Inf. Tech. Vol. 2, No. 2, pp. 113-119 (2000).
[8] H. Sacks, E.A. Schegloff, and G. Jefferson. A simplest systematics for the organization of turn-taking for conversation. Language, 50(4), 696-735 (1974).
[9] D. Heylen, A. Nijholt, and M. Poel. Generating Nonverbal Signals for a Sensitive Artificial Listener. In: Verbal and Nonverbal Communication Behaviours. Lecture Notes in Computer Science 4775, Springer-Verlag Berlin Heidelberg, pp. 264-274 (2007).
[10] D. Reidsma, H. van Welbergen, R. Poppe, P. Bos, & A. Nijholt. Towards Bi-directional Dancing Interaction. Proceedings 5th International Conference on Entertainment Computing (ICEC 2006), Cambridge, UK, R. Harper, M. Rauterberg, M. Combetto (Eds.), Lecture Notes in Computer Science 4161, Springer-Verlag, Berlin, Heidelberg, pp. 1-12 (2006).
[11] P. Bos, D. Reidsma, Z. Ruttkay & A. Nijholt. Interacting with a Virtual Conductor. Proceedings 5th International Conference on Entertainment Computing (ICEC 2006), Cambridge, UK, R. Harper, M. Rauterberg, M. Combetto (Eds.), Lecture Notes in Computer Science 4161, Springer-Verlag, Berlin, Heidelberg, pp. 25-30 (2006).
[12] Z.M. Ruttkay & H. van Welbergen. On the timing of gestures of a Virtual Physiotherapist. 3rd Central European Multimedia and Virtual Reality Conference, C.S. Lanyi (eds), Pannonian University Press, Hungary, pp. 219-224 (2006).
[13] 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, pp. 287-311 (2005).
[14] A. Nijholt et al. Mutually Coordinated Anticipatory Multimodal Interaction. In: Nonverbal Features of Human-Human and Human-Machine Interaction. Lecture Notes in Computer Science, Springer-Verlag Berlin Heidelberg, to appear (2008).
[15] R. Rienks, A. Nijholt, D. Reidsma. Meetings and Meeting Support in Ambient Intelligence. Chapter 18 in: Ambient Intelligence, Wireless Networking, Ubiquitous Computing. Th. A. Vasilakos & W. Pedrycz (eds.), Artech House, Norwood, MA, USA, pp. 359-378 (2006).
[16] W. Ju & L. Leifer. The Design of Implicit Interactions: Making Interactive Objects Less Obnoxious. Design Issues: Special Issue on Design Research in Interaction Design. In Press.
[17] S. Babu, C. Zanbaka, J. Jackson, T-O Chung, B. Lok, M.C. Shin, & L.F. Hodges. Virtual Human Physiotherapist Framework for Personalized Training and Rehabilitation, *Graphics Interface 2005*, Victoria, British Columbia, Canada (2005).
[18] S. Mokka, A. Väätänen, J. Heinilä, & P. Välkkynen. Fitness computer game with a bodily user interface. Proceedings of the Second international Conference on Entertainment Computing, Pittsburgh, ACM International Conference Proceeding Series, vol. 38. Carnegie Mellon University, Pittsburgh, PA, pp. 1-3 (2003).
[19] P.T. Chua et al. Training for Physical Tasks in Virtual Environments: Tai Chi. Proceedings of the IEEE Virtual Reality 2003, IEEE Computer Society, Washington, DC, pp. 87-94 (2003).
[20] J. Höysniemi, A. Aula, P. Auvinen, J. Hännikäinen, & P. Hämäläinen. Shadow boxer: a physically interactive fitness game. Proceedings of the Third Nordic Conference on Human-Computer interaction. NordiCHI '04, vol. 82. ACM Press, New York, NY, pp. 389-392 (2004).
[21] P. Hämäläinen et al. Martial Arts in Artificial Reality, Proceedings ACM Conference on Human Factors in Computing Systems (CHI'2005), ACM Press (2005).
[22] F. Müller & S. Agamanolis. Sports over a distance. Comput. Entertain. Vol. 3, No. 3, pp. 4-4, 2005.
[23] B. Meyers et al. Dance your work away: exploring step user interfaces. In CHI '06 Extended Abstracts on Human Factors in Computing Systems, Montréal, CHI '06. ACM Press, New York, NY, pp. 387-392 (2006).
[24] N. Ambady and R. Rosenthal. Thin slices of expressive behavior as predictors of interpersonal consequences: A meta-analysis. Psychological Bulletin, 111, 2, pp. 256-274 (1992).
[25] P. Borkenau, N. Mauer, R. Riemann, F.M. Spinath, and A. Angleitner. Thin Slices of Behavior as Cues of Personality and Intelligence. Journal of Personality and Social Psychology, Vol. 86, No. 4, pp. 599–614 (2004).
[26] Bechinie Bechinie, M., & Grammer, K. Charisma. Cam: A prototype of an intelligent digital sensory organ for virtual humans. Proceedings *IVA 2003*, LNAI 2792, Springer, pp. 212-216 (2003).
[27] R.W. Picard, E. Vyzas, & J. Healey. Toward machine emotional intelligence: Analysis of affective physiological state. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 23 (2001), No. 10, pp. 1175-1191 (2001).
[28] K. Gilleade, A. Dix, J. Allanson. Affective Videogames and Modes of Affective Gaming: Assist Me, Challenge Me, Emote Me. Proceedings of DIGRA'2005, Vancouver, Canada (2005).
[29] D.S. Tan. Brain-Computer Interfaces: applying our minds to human-computer interaction. Informal proceedings "What is the Next Generation of Human-Computer Interaction?" Workshop at CHI 2006, Montreal (2006).