Judging Laura: Perceived Qualities of a Mediated Human Versus an Embodied Agent

Renate ten Ham[1], Mariët Theune[1], Ard Heuvelman[2], and Ria Verleur[2]

[1] Department of Computer Science
{theune, hamrh}@cs.utwente.nl
[2] Department of Behavioral Sciences,
University of Twente, P.O. Box 217, 7500 AE, Enschede, The Netherlands
{a.heuvelman, r.verleur}@utwente.nl

Abstract. Increasingly, embodied agents take over tasks which are traditionally performed by humans. But how do users perceive these embodied agents? In this paper, we describe an experiment in which we compared a real person and a virtual character giving route instructions. The voice, the outfit and the gestures were kept (close to) identical for both cases. The participants judged them, among other things, on trustworthiness, personality and presentation style. In contrast to the outcome of earlier investigations, in most categories the agent scored better or comparable to the human guide. This suggests that embodied agents are suitable to take the place of humans in information-giving applications, provided that natural sounding speech and natural looking nonverbal behaviors can be achieved.

1 Introduction

In order to make human-computer interaction similar to face-to-face communication between humans, an increasing number of interfaces are being equipped with human-looking virtual characters that can use natural language and display nonverbal behaviors. These characters are referred to using different terms, including 'synthetic personae' (McBreen et al., 2000), 'embodied conversational agents' (Cassell et al., 2000), and 'animated interface agents' (Dehn & van Mulken, 2000). For brevity, in this paper we will refer to them as 'embodied agents' or simply as 'agents'.

It is generally assumed that for an agent to be optimally engaging and effective, it has to be as lifelike as possible. Several studies showed that when an embodied agent seems more human in its appearance and behavior, more human qualities are accredited to it. King & Ohya (1996) carried out an experiment with stimuli varying from simple geometric shapes to lifelike human forms, which were rated on agency and intelligence. One of their conclusions was that a human-like appearance and 'subtle behavioral displays' - such as eye blinking - have a great effect on the user's appraisal of these capabilities. Embodied agents can offer intelligence, personality and emotion and therefore communication properties that help to make us feel understood and appreciated (Nijholt, 2004). Users have been shown to like embodied agents and find them engaging (Takeuchi & Naito, 1995; Koda & Maes, 1996).

T. Panayiotopoulos et al. (Eds.): IVA 2005, LNAI 3661, pp. 381 – 393, 2005.

Increasingly, agents are used for tasks that are traditionally performed by humans, such as providing information, explaining or answering questions as an instructor or a teacher. More and more companies use an agent on their website, or use an agent to give information in their office building. Cassell et al. (2002) observed that "users' behaviors appeared natural, as though they were interacting with another person" when using MACK (Media lab Autonomous Conversational Kiosk), an embodied agent answering questions about and giving directions to the MIT Media Lab's research groups, projects and people. With respect to educational applications, Lester et al. (1999) state that "… because of their strong visual presence and clarity of communication, explanatory lifelike avatars offer significant potential for playing a central role in next-generation learning environments." But how do users feel, when they get information from an agent instead of a real person? Reeves and Nass (1996) have shown that people respond to computers and other media like they respond to people, treating them as social actors and attributing them with personality. But how will people judge the personality of an embodied agent, compared with the personality of a real person? Will they have the same emotional response to agent and person, and will they trust information given by the agent as much as information given by the person?

This study provides an exploratory investigation into these questions. An experiment was performed in order to compare an embodied agent and a video recording of a real person on, among other things, trustworthiness, personality, presentation style and user's emotional response. In our experiment we focused on the effect of human versus synthetic appearance rather than on behavior.

2 Related Work

Most embodied agent evaluations have focused on comparing interfaces with or without an embodied agent, and on comparing agents with different visual appearances.

Koda & Maes (1996) compared agents in a poker game that were embodied as a smiley, a dog, a cartoon face, a realistic face (a photo of a real person), or not at all. They found that the embodied agents were considered more likeable and engaging than the disembodied agent. The realistic face was found slightly more likeable and engaging than the other faces, but not significantly so.

Sproull et al. (1996) compared the use of a realistic 3D talking head (stern or neutral looking) with a textual interface in a career counseling application. Their subjects responded in a more social way to the talking heads than to the text-only interface. On the other hand, subjects perceived the personality of both versions of the talking head as less positive than that of the text-only interface. As suggested by Dehn & van Mulken (2000) the latter result may have been influenced by the voice of the talking heads, which lacked inflection and thus did not sound entirely natural.

McBreen et al. (2000) compared the following agent embodiments: a photo of a real person with or without lip movement, a 3D talking head, and a video of a real person. They also compared a disembodied condition, where the agent was represented by a voice only. The same (human) speech soundtrack was used in all cases. Overall, the videos were rated best for likeability (friendliness, competence, naturalness) and several other aspects. The talking heads were rated worst on almost all

fronts. However, this might be explained by the fact that the talking face had minimal facial expressions.

Beun et al. (2003) compared a photorealistic 2D talking head, a cartoon character, and an agent represented by only a text balloon, measuring two variables: anthropomorphism (in terms of the agent being helpful, sensible, etc.) and memory performance (subjects had to remember two stories told by the agents). The realistic face scored highest on anthropomorphism. For both embodied agents, memory performance was better than for the disembodied agent.

3 Experiment

The question we try to answer in our experiment is how users perceive an embodied agent as compared to a real person, in the context of an information presentation task. Here, we focus on the user's subjective experience rather than on objective measures such as memory performance. The presentation task chosen for our experiment is that of route description. This is a task where embodiment is quite appropriate: in real life, verbal route descriptions are most often presented in a face-to-face situation, and the speaker typically displays nonverbal behavior (mainly in the form of gestures) while giving the description. This makes route description a suitable task for our experiment.

Applications that involve an embodied agent giving route descriptions include virtual receptionists (e.g., Cyberella, Rist et al., 2002) and virtual guides in real or virtual environments (Cassell et al., 2002; Kopp et al, 2004; Theune et al., 2005). In general, such applications are not aimed at achieving maximal efficiency but rather at giving the user a lifelike experience.

3.1 Design

There were two conditions in our experiment. The participants were initially presented with a route description that was either given by a human guide, recorded on video (condition 1) or by an embodied agent (condition 2). We adhered to methodological standards by making the human guide and the agent guide as similar to each other as possible, only varying the dimension under investigation: i.e., the synthetic versus human appearance of the guide. How we achieved this is discussed in section 3.4.

For both versions of the guide we used the name Laura: the actual name of the human guide. After the participants had watched the route description by the human or the agent guide, they were asked several series of questions, measuring among other things their emotional response and their perception of the guide's personality. Then they were shown a movie with the same route description, but this time presented by the version of the guide they had not seen yet. After this second movie, when the participants had seen both agent and human guide, they were asked their opinion about the quality of the agent, and they had to indicate which version of the guide they preferred.

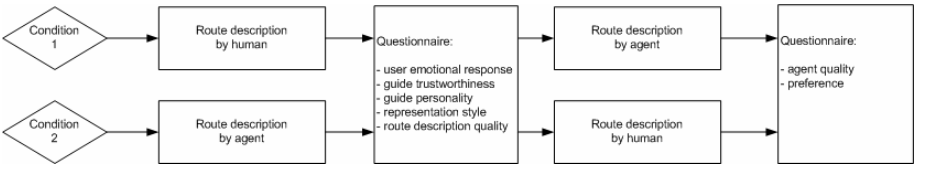

Fig. 1. Graphical representation of the research design

A few limitations of this experimental design are the following. First, arguably the most important property of agents is their ability to interact with users. In our experiment, however, we used an agent for a non-interactive task: presenting route information. We opted for one-way communication so that all participants would get the same information in the same way, thus restricting the variation to the dimension we were interested in. Second, we used a video recording rather than a 'live' person to compare the agent with. However, watching a video is not fully comparable to being face-to-face with another person. For example, Burgoon et al. (2002) found that mediated interaction (video conferencing) in a decision-making task scored much lower than face-to-face interaction on social judgments such as involvement, trust and sociability. On the other hand, this effect of mediation can be expected to be smaller in situations where there is no actual interaction, as in our experiment. People are used to seeing mediated people presenting information, for example newsreaders on television. And an embodied agent is in any case mediated: people need a computer to interact with it. This means that to keep the experimental conditions as similar as possible, the human guide in our experiment had to be mediated too.

3.2 Dependent Variables

After having seen the route description given by either the agent or the human guide, the participants in the experiment answered several questions. In this section we explain how these questions were grouped, and how reliable these groupings are. All questions were measured on a nine-point scale, except the question about preference.

User emotional response was measured using the Self-Assessment Manikin (SAM), a visual scale which represents the user's emotional response to a stimulus, with respect to the dimensions valence (pleasant or unpleasant), arousal, and dominance (Lang, 1985). SAM reflects each dimension with a graphic character arrayed along a continuous nine-point scale. For valence, SAM ranges from a smiling happy figure to an unhappy figure (see Figure 2). For arousal, SAM ranges from an excited figure to a sleepy figure. Finally, the dominance scale goes from a very small figure to a very big figure.

Guide trustworthiness was measured in terms of seven items: expertise, believability, realism, reliability, friendliness, sympathy, and dominance.

Guide personality was measured using Cattell's 16PF, Personality Factors. These 16 factors represent the most important personality factors according to Cattell & Cattell

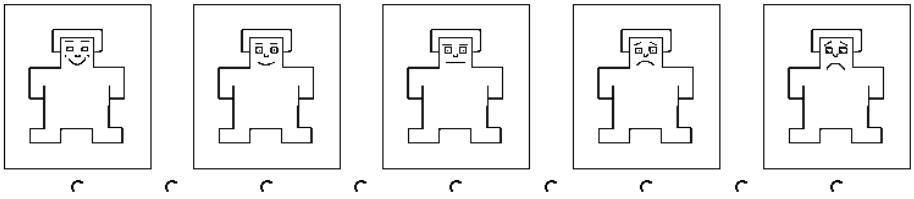

Fig. 2. The SAM scale for valence

(1995). The 16 factors are: warmth, reasoning, emotional stability, dominance, liveliness, rule-consciousness, social boldness, sensitivity, vigilance, abstractness, privateness, apprehension, openness to change, self reliance, perfectionism and tension.

Presentation style regards the way the guide presented the route. This reliable index (α=0,78) was formed by twelve nine point scale items: good-bad, pleasant-unpleasant, polite-impolite, natural-artificial, flowing-clumsy, relaxed-tense, energetic-lethargic, dynamic-static, accurate-inaccurate, calm-excited, exuberant-apathetic, and interested-uninterested.

Route description quality measured the way participants felt about the route description itself. This index was comprised of eight nine point scale items: concise-tedious, simple-complex, easy-difficult, interesting-boring, structured-unstructured, useful-useless, clear-unclear, and comprehensible-incomprehensible. This index is reliable: $\alpha = 0.80$.

Agent quality consists of six items measuring how participants felt about the quality of the embodied agent: good-bad, modern-old fashioned, realistic-unrealistic, advanced-outdated, usable-unusable, innovative-traditional.

Preference was determined using one simple question: "Which of the two do you prefer: virtual person (agent) or real person (video)?"

3.3 Participants

Participants in the experiment were 78 undergraduate students from different faculties in our university. They were all following a course in Media Psychology and were rewarded with bonus points to participate. Participants were randomly assigned to one of the conditions, with age and gender approximately balanced across conditions. The average age of the participants was 21; 60 % of the participants were female.

3.4 Material

For the agent we used the Living Actor™ technology from Cantoche.[1] We wanted to make the agent as human-like as possible, so we selected an agent that looked realistic rather than cartoon-like and had a large repertoire of gestures. The agent that best met these requirements happened to be female, the Cantoche character 'Julie'. We wanted to reduce the differences between the agent and the human guide as much as possible, so that synthetic versus human appearance (the dimension under investigation) was

[1] www.cantoche.com

the only dimension on which the two guides were different. Therefore we asked someone who looked like the agent to play the role of the human guide, and dressed her in exactly the same clothes as the agent.

The movies of the route presentations were created as follows. First, we made a video recording of the human guide as she spontaneously described the route. Then we scripted the agent to simulate the gestures that had been made by the human guide as closely as possible, e.g., pointing left and right. Because of limitations in the gesture repertoire of the agent, this simulation deviated in a few respects from the original recording. Therefore we made a final recording of the human guide as she was describing the route, this time mimicking the agent. The human actor was not asked to imitate the agent in every behavioral detail, only at the more global level of gestures. The use of different gestures would have made the presentations of the guides too dissimilar to allow for a reliable comparison, but we considered the smaller unconscious movements such as blinking and head movements as part of what made the human guide appear human and the agent guide synthetic.

Finally, we added the speech of the human guide to the agent, synchronized the agent's gestures and lip movements with the speech, and created a white background for both movies. This resulted in two route description movies by guides that used exactly the same speech, had roughly the same appearance, and used the same gestures. Also, both guides had a neutral facial expression. Only the more subtle nonverbal behaviors such as blinking, head movements and small posture shifts were different between the guides. Overall, they acted and looked similar, the main difference being that one guide was human and the other an embodied agent (see Figure 3).

Fig. 3. The human guide (left) and the agent (right)

3.5 Procedure

The experiment was performed in a Web environment. After a short instruction, the participants started the questionnaire on their computer. The movies with the route presentations were integrated in the questionnaire. The participants could not see the movies twice, nor could they go back to see or change their previous answers.

Depending on the group they were assigned to, participants would start with watching a movie with either the agent or the human guide presenting the route. Both movies started with the guide introducing herself: "Hi, I'm Laura." She would then thank them for their cooperation and explain she was going to give them a route description. This way the participants could get used to the voice and the appearance of the guide.

4 Results

With the exception of user emotional response, which was measured using SAM (see section 3.2), and preference, where the participants had to indicate whether they preferred the human or the agent guide, a nine-point scale was used for all questions. The ends of the scale correspond to contrastive attributes such as good-bad, pleasant-unpleasant etc. In the results given below, the high end of each scale corresponds to the positive attribute, and the low end to the negative attribute in the pair. For most pairs, e.g., good-bad, it is clear which attribute is positive and which is negative. However, for some pairs, e.g., calm-excited, we had to judge which attribute would be considered most positive given the task performed by the guide. In the tables below, the attribute we judged to be more positive is always listed first.

We used the SPSS program (T tests) to compare the mean of the scores on all dimensions as described in paragraph 3.2. This test compares the mean of each item or index for both conditions. The F-value indicates the difference between the two conditions. Differences where $p < .05$ will be treated as significant.

User emotional response

On this dimension, there were no significant differences between both groups (see Table 1).

Table 1. Separate items for user emotional response (* p<.05)

	Agent	Human	Significance (2-tailed)
Valence	4.37	4.83	.09
Arousal	6.68	6.13	.19
Dominance	4.97	5.45	.21

Guide trustworthiness

The participants felt that the agent was more competent than the human guide ($F = 0.98$, $p<0.05$). The scores on the other items did not differ significantly between human guide and agent. Reliability of the guide was rated exactly the same for both guides.

Guide personality

The agent was seen as more relaxed than the human guide ($F = 1.07$, $p<0.01$), more self assured ($F = 0.73$, $p<0.05$) and less traditional ($F = 1.40$, $p<0.01$). Participants

who saw the agent remarked that they found it likable, but businesslike. On the other hand, several participants who saw the person said that they couldn't really judge personality based on the short presentation. This comment did not occur for the agent.

Presentation style

Table 2 shows all the separate items from this index. The presentation style of the agent was seen as significantly more relaxed, dynamic and interested than the presentation style of the human guide. Overall, there was a significant main effect with regard to the presentation style index ($F = 0.39$, $p<0.05$), such that participants found the presentation style of the agent better than the style of the human guide. A few of the remarks are: "very humanlike" and "neutral, but very accurate and polite". The real person was found "too boring" and "pretended".

Table 2. Separate items for presentation style (* p<.05)

	Agent	Human
Good-bad	4.79	4.70
Pleasant-unpleasant	4.92	4.70
Polite-impolite	6.39	6.23
Natural – artificial	5.47	4.88
Flowing – clumsy	5.82	5.28
Relaxed – tense	6.05	5.35*
Energetic – lethargic	5.29	4.75
Dynamic – static	4.47	3.36*
Accurate – inaccurate	6.42	3.38
Exuberant – apathetic	4.26	4.03
Calm – excited	3.16	3.13
Interested - uninterested	5.53	4.83*

Route description quality

There was a significant main effect with regard to this index ($F = 0.50$, $p<0.05$), such that participants found the route description better when it was presented by the agent. As Table 3 shows, the agent scored higher on every single item, although only one item is significant: the route description given by the agent was considered significantly less boring than the description given by the human guide.

Table 3. Separate items for route description quality (* p<.01)

	Agent	Human
Concise - tedious	4.05	3.30
Simple - complex	3.82	3.33
Easy - difficult	3.97	3.58
Interesting - boring	3.95	3.00*
Structured - unstructured	5.92	5.55
Useful - useless	4.45	4.08
Clear - unclear	5.34	5.00
Comprehensible - incomprehensible	5.63	5.28

Agent quality

After the participants had answered the previous questions (which they answered for either the human or the agent guide, i.e., the only guide they had seen so far), they were shown the version of the guide they had not yet seen. Then, when it was certain that all participants had seen the agent guide, they were asked to rate its quality.

Table 4 shows the results, split between participants who had first seen the agent (and had rated its personality, etc.) and participants who had first seen the real person (and had rated its personality, etc.). We can see that participants who first saw the real person, and then the agent, regarded the agent as less realistic ($F = 0,92$, $p<0,05$) and less advanced ($F = 0,89$, $p<0,05$) than the participants who saw the agent first.

Table 4. Separate items for agent quality index, split between participant groups (* p<.05)

	Agent first	Human first
Good - bad	6.05	6.08
Modern - old-fashioned	5.95	6.00
Realistic - unrealistic	6.18	5.20*
Advanced - outdated	6.03	5.03*
Usable - unusable	6.24	5.53
Innovative - traditional	5.16	5.35

Table 5 shows the overall scores, averaging over both groups. In general, we can see that the agent is regarded as fairly modern, realistic and usable.

Table 5. Separate items for agent quality index

	Overall score
Good - bad	6.06
Modern - old-fashioned	5.97
Realistic - unrealistic	5.68
Advanced - outdated	5.51
Usable - unusable	5.87
Innovative - traditional	5.26

Preference

About half (52%) of the participants preferred the real person; the other half preferred the agent. The most mentioned reason for choosing the agent is that people felt less distracted. Several participants indicated that curiosity about age, profession, or what she is wearing will distract when working with a human guide. When working with an agent they could concentrate more on the message instead of the guide. Participants who chose the real person commented mostly that the real guide was “more personal” or “more intimate”. Or just said that they liked a real person better. It seems as though subjects preferring the agent had a clearer motivation to do so than subjects preferring the human guide.

Previous experience
Participants who didn't have any previous experience with embodied agents found the agent to be more believable ($F = 1.38$, $p<0.05$), and its presentation style more relaxed ($F = 1.05$, $p<0.05$), flowing ($F = 1.19$, $p<0.05$) and exuberant ($F = 1.12$, $p<0.05$) than participants who had no prior experience with embodied agents.

5 Discussion

In previous experiments, videos or photos of real persons were generally preferred over 2D or 3D agent embodiments (e.g., Koda & Maes, 1996; McBreen et al., 2000). Remarkably, in our experiment the findings are reversed: overall, the embodied agent received more positive ratings than the video recording of the human guide. There are several factors that may help explain these results. First of all, the agent was of good quality (as confirmed by the participants' ratings), having a realistic appearance, a natural human voice,[2] and quite natural movements that included not only gestures but also more subtle behaviors such as blinking, head movements and posture shifts. All in all, despite being an animation, the agent appeared fairly realistic and this may have led to more positive judgments than were found in previous experiments.

Another factor that may have caused a preference for the agent is that the participants in our experiment were young people, who are generally open to new technology and may appreciate a novelty, like a virtual character, more than a well-known phenomenon like a real person. With an older age group, the outcomes might well have been different. A comparison between an older group of participants and the original group of students might make clear if age is of influence on the results.

At the same time, there are also some factors that may have negatively influenced the scores of the human guide. One of these is the fact that in the final version of our recording, she was acting instead of behaving spontaneously. She had to recreate her earlier spontaneous description, this time keeping in mind which gestures the agent could and could not make. For this reason she may have come across as less self-assured and less relaxed, and thus also as less competent. The participants may also have had higher expectations of the human guide than of the virtual guide: when people see a real person explaining a route, they may expect more spontaneous gestures than were actually performed by the actress. This could have caused the participants to judge the route description by the human as relatively static and boring. On the other hand, one of the participants remarked: "I can imagine an agent explaining something in a very boring way". An agent will give a steady, always similar performance, and people expect this to happen. This may explain as well why the real person was found to be more static and boring. In addition, the combination of a hu-

[2] We used the same voice for both human and agent in order to reduce the differences between the two conditions, and to avoid possible negative effects of a synthetic voice (cf. Sproull et al. 1996). However, some subjects (in particular those having previous experience with embodied agents), commented that they found the combination of a human voice with a synthetic agent somewhat unnatural. So, the use of a natural voice might as well have been an advantage as a disadvantage.

man voice and appearance with artificial behavior (as in some sense the human guide was mimicking the artificial agent) may have been perceived as inconsistent. As shown by Isbister & Nass (2000), people tend to dislike inconsistency within agents. However, our agent was also inconsistent in the sense that it coupled an artificial appearance and behavior with a human voice (i.e., the voice of the human guide). Some participants remarked that they found this unnatural, although this did not lead to a more negative judgement.

Most participants who reported previous experience with embodied agents, referred to computer games. The characters in computer games have much more advanced graphics and animation than the agent used in our experiment, and this probably explains why this group of participants was significantly less positive about agent quality than the group who had no experience with embodied agents at all. Also remarkable is that participants who saw the agent first rated the quality of the agent significantly higher than participants who saw the human guide first. The explanation for this may be that people who watched the agent first focused more on the information it presented, whereas people who watched the agent second were already familiar with the information, and therefore had more attention to spare for the inherent properties qualities of the agent.

An interesting factor is also the time people spend with the embodied agent. One might expect that if people get the time to get used to how an agent presents the information, they might be able to focus even more on the message instead of the presenter. In our experiment people were confronted with only very short movies; but if there is extended usage, the differences in perception will probably be more pronounced. Of course with respect to the use of ECA's in applications it is also very relevant to find out how 'syntheticness' influences task performance; in fact we have investigated this too, but we will report on that in another paper.

6 Conclusions and Further Research

The question we tried to answer in our experiment is how users perceived an embodied agent as compared to a real person. We carried out an experiment with 78 participants who either received a route description from a human on video, or from an embodied agent. The equal scores on emotional response to, and trustworthiness of, the agent and the real person indicate that agents have strong potential as a guide, tutor or advisor. A striking result was that the comparison in presentation style turned out in favor of the agent rather than the human guide. The quality of the route description given by the agent was also perceived more positive on every dimension. Even though these results may have been partially influenced by the set-up of our experiment (with the human guide acting not entirely spontaneously), this is encouraging news for developers of interface agents. The fact that agent and human scored about the same on personality is encouraging as well. An important caveat is that to be comparable with a human guide, the agent has to sound natural and display human-like nonverbal behaviors. Especially in fully interactive situations, which go beyond pure information presentation as in our experiment, achieving this still presents an important challenge.

We see several options for further research. For instance, repeating the experiment with a male agent or with an older age group might very well produce different outcomes. Also, as mentioned above, speech and interaction are a very important part of the communication between humans and agents. Further research on the influence of these factors will help to determine how people perceive agents.

Acknowledgements

We thank Cantoche for allowing us to use their agents, and Jan Oosterhuis for technical support in setting up the experiment. We also thank Zsófia Ruttkay and three anonymous referees for their useful comments on the first version of this paper. This work was carried out in the context of the ANGELICA project, sponsored by NWO (Netherlands Organization for Scientific Research).

References

Beun, R.J., de Vos, E., & Witteman, C. Embodied conversational agents: effects on memory performance and anthropomorphisation. Proc. IVA (2003), 315-319.

Burgoon, J.K., Bonito, J.A., Ramirez, A., Dunbar, N.E., Kam, K., & Fischer, J. Testing the interactivity principle: Effects of mediation, propinquity, and verbal and nonverbal modalities in interpersonal interaction. J. of Communication special issue: Research on the Relationship between Verbal and Nonverbal Communication: Emerging Integrations (2002), 657-677.

Cassell, J., Sullivan, J., Prevost, S., & Churchill, E. Embodied Conversational Agents. MIT Press (2000).

Cattell, R.B. & Cattell, H.E.P. Personality structure and the new fifth edition of the 16PF. Educational and Psychological Measurement, (1995), 6, 926-937.

Dehn, J., & van Mulken, S. The impact of animated interface agents: A review of empirical research. Int. J. Human-Computer Studies (2000), 52, 1-2.

Isbister, K. & Nass, C. Consistency of personality in interactive characters: Verbal cues, nonverbal cues and user characteristics. Int. J. of Human Computer Studies 53 (2000), 251-267.

King, W.J. & Ohya, J. The representation of agents: Anthropomorphism, agency and intelligence. CHI ' 96 Conference Companion, Vancouver, B.C., (1996), 289-290.

Koda, T. & Maes, P. Agents with faces: The effect of personification. Proc. IEEE Robot-Human Communication (1996).

Kopp, S., Tepper, P. & Cassell, J. Towards integrated microplanning of language and iconic gesture for multimodal output. Proc. of the International Conference on Multimodal Interfaces (ICMI), 2004.

Lang, P.J. The Cognitive Psychophysiology of Emotion: Anxiety and the Anxiety Disorders. Hillsdale, N.J.: Lawrence Erlbaum (1985).

Lester, J. Zettlemoyer, L, Gregoire, J., & Bares, W. Explanatory lifelike avatars: Performing user centered tasks in 3D learning environments. Proc. Autonomous Agents'99, ACM Press (1999).

McBreen, H., Shade, P., Jack, M., & Wyard, P. Experimental assessment of the effectiveness of synthetic personae for multi-modal E-retail applications. Proc. 4th Int. Conf. on Autonomous Agents (2000), 39-45.

Nijholt, A. Where computers disappear, virtual humans appear. Computers and Graphics 28 (2004).

Reeves, B. & Nass, C. The Media Equation: How People Treat Computers, Television, and New Media like Real People and Places. Cambridge University Press, New York (1996).

Rist, T., Baldes, S., Gebhard, P., Kipp, M., Klesen, M., Rist, P., & Schmitt, M. CrossTalk: An interactive installation with animated presentation agents. Proc. of COSIGN'02 (2002).

Sproull, L., Subramani, R., Kiesler, S., Walker, J., & Waters, K. When the interface is a face. Human-Computer Interaction, 11 (1996), 97-124. (Reprinted in B. Friedman, (1997) Human values and the design of technology, CLSI, Publications.)

Takeuchi, A., & Naito, T. Situated facial displays: Towards social interaction. Proc. SIGCHI conference on Human factors in computing systems (1995), 450-455.

Theune, M., Heylen, D. & Nijholt, A. Generating embodied information presentations. In O. Stock and M. Zancanaro (eds.), Multimodal Intelligent Information Presentation, Kluwer Academic Publishers (2005), 47-70.