

The unfriendly user: exploring social reactions to chatterbots

Antonella De Angeli, Graham I. Johnson, and Lynne Coventry

NCR Self-Service, Advanced Technology & Research, Dundee DD2 3XX, UK.

Abstract

This paper presents a preliminary evaluation of Alice, a chatterbot designed in order to elicit anthropomorphic attributions and emotional reactions from those who chat to 'her'. The analysis is based on both transcripts of the interaction and user comments collected in a focus group. Results suggest that the introduction of explicit anthropomorphism in Human-Computer Interaction (HCI) is a complex phenomenon, which could generate strong negative reactions from the part of the user. The finding also demonstrates the importance of placing the development of user interfaces within a social framework as the technology tends to establish relationships with users.

Keywords Chatterbots, anthropomorphism, disembodied language, social dynamics.

1. Introduction

The attribution of human characteristics to animals, events, or objects is a natural tendency for human beings, known as anthropomorphism. According to Caporeal and Heyes (1996), it reflects an inherent *cognitive default*: a standard schema is engaged in explaining the behaviour of entities for which, otherwise, explanations are not readily available. Being complex machines whose procedural mechanisms are hidden to users, computers have always been a favourite target for anthropomorphic attributions (De Angeli, Gerbino, Nodari and Petrelli, 1999). Since their advent, they have never been perceived simply as machines or just as the result of the interaction between hardware and software. Because of that, computers have a *memory* and speak a *language*; they can contract *viruses* and act autonomously. In recent years, then, the human-metaphor has been increasingly strengthened trying to represent these inanimate, hard-textured objects as warm, soft and humanoid. In particular, current computers are expected to be *friendly* towards their users.

This paper examines the hypothetical *friendship* between computers and users from an unusual perspective – that of the machine. It is built around a preliminary evaluation of a chatterbot called Alice (http://206.184.206.210/alice_page.htm), a proud 'robot' that exhibits human like feelings and intentions while chatting with a 'human partner' (the user). As every chatterbot, Alice has been explicitly designed to trigger a number of anthropomorphic attributions, including social and emotional intelligence, personality and affect. Alice is the prototype of an old human dream: creating non-human companions with technology. For decades, science-fiction writers have envisioned a world in which robots and computers acted like human assistants. Nowadays, for better or worse, that world looks closer. A number of animated characters, chatterbots and talking heads populate it. They act as assistants, guides, sale people and entertainers on the Internet. They are the first-generation social agents –interface software explicitly designed to set up lasting and meaningful relationships with users (De Angeli, Lynch and Johnson, 2001). These systems are likely to produce a fundamental shift in the way computers are designed, used and evaluated (Parise, Kiesler, Sproull and Waters, 1999; De Angeli, Lynch and Johnson, in press). Indeed, social stimuli are more complex than physical stimuli. They are more likely to be causal agents, influencing their partners' behaviour through direct interactions, their mere presence or even their virtual presence. Social agents are active partners in joint activities, co-ordination of individual actions by two or more people that emerge in time as they try to accomplish certain common goals (Clark, 1996). Social agents perceive while they are perceived and change while inducing changes. Further, they strongly involve the observer self-concept and are sometimes difficult to understand. Indeed, many attributes, such as personality traits, intents or attitudes, are not directly observable and the accuracy of the observation is hard to determine.

The prevailing approach driving the design of social agents is anthropomorphic in nature. Proponents cite the naturalness and power of anthropomorphic interfaces as fundamental strengths (Laurel, 1997); detractors claim that exasperating human characteristics can disempower, mislead and confuse the user (Shneiderman, 1997). However, neither of these views is supported by a clear and unambiguous understanding: more research

is needed to comprehend users' acceptance of and reactions to the introduction of human characteristics in user interfaces. Our research effort is devoted to help filling this gap via ethnological studies and controlled experiments. We believe that to fully understand how to create socially adept technologies we should *adapt* and not merely *adopt* appropriate theories from social psychology and communication in the new relationship context. Note that this claim slightly diverges from the prevailing social approach to HCI proposed by the media equation paradigm (Reeves and Nass, 1996). According to it, individuals' interactions with computers, televisions and new media are fundamentally social and natural, exactly like people-people interactions in real life. This implies that the same social rules explaining interpersonal relationships can be directly applied to HCI. The claim is supported by an incredible number of empirical studies, showing a clear equivalence between people reactions to human and artificial companions (e.g., Reeves and Rickenberg, 2000; Nass and Lee, 2000). Despite this, we still struggle to accept the idea of a complete similarity between the two different contexts, idea that justifies the direct adoption of the social-sciences theoretical apparatus in HCI.

As a matter of fact, flexibility is a fundamental social ability. We adapt our behaviour, expectations and reactions to different partners and contexts. There are many different kinds of rules for relationships (Dwyer, 2000), variation depending on the types of relationships and on the cultural backgrounds of the partners. Moreover, social perception is often biased and depends on the particular self-concept activated in a specific situation (Turner, 1987). Hence, even the same person can react differently to the same social stimulus according to the specific context in which she is acting. Research in the field of natural language interaction has already demonstrated that face-to-face communication is not an adequate model to explain and predict HCI (Bernsen, Dybkjær and Dybkjær, 1998). Talking to a computer, people maintain a conversational framework but tend to simplify the syntactic structure and to reduce the utterance length. We expect a similar simplification effect even in social dynamics. In our opinion, stating that "computers are social actors" (Nass, Steuer and Tauber, 1994) does not necessarily imply that computers are human actors. Rather, we believe that anthropomorphic agents will create a specific social world, with its own rules and dynamics that need to be fully understood. This idea is also supported by a study by Nass and colleagues claiming that even HCI and CMC (Computer Mediated Communication) do not give rise to the same psychological reactions (Morkes, Kernal and Nass, 1999). Further, a recent web survey investigating a much larger sample and a broader population found little support for the media equation paradigm (Couper, Tourangeau and Steiger, 2001).

Our research represents an initial contribution to a cyber-social model which attempts to explain how users perceive, react and relate to social agents. The final goal is to understand how humans create, maintain and make sense of their social/affective experiences with artificial entities that explicitly reproduce anthropomorphic behaviour. We do not intend to take a position for or against artificial entities displaying humanlike features. Nevertheless, we are convinced that social agents are here to stay. Therefore, it is necessary to address the many social and emotional questions raised by this technology in order to drive the design of real *human-centred*, human-like machines.

2. Chatterbot: talkative interface

Chatterbots are computer programs that attempt to simulate typed conversations with the users. The complexity of their algorithm is variable, but in general they are programmed to respond to user inputs with canned pre-scripted statements. In this way, chatterbots can have a somewhat logical conversation with a human user, even without being capable of understanding. Rather, they are all about the illusion of intelligence and the suspension of disbelief on the part of the user. Following the famous idea of the "imitation game" (Turing, 1950), chatterbots are developed with the aim of fooling (at least temporarily) a human into thinking they are talking to another person.

One of the oldest and best-known chatterbots in the world is Eliza. 'She' was created in the '60s by MIT scientist Joseph Weizenbaum to play the role of a psychotherapist in a clinical treatment (Weizenbaum, 1966). Nowadays, Eliza may look limited but her fundamental technical components are still at the basis of the most innovative chatterbots. Eliza identifies the 'most important' keywords occurring in the input message. Then, she tries to define a minimal context in which the keywords appear (e.g., the keyword "you" followed by the word "are" is interpreted as an assertion). Finally, Eliza chooses an appropriate transformation rule to modify the user input. To summarise, Eliza works by turning the user sentences around. A typical dialogue is reported in Example 1. It is part of a longer conversation between one of the authors and the chatterbot.

Eliza *So how are you doing today?*
User *I am tired*
Eliza *Did you come to me because you are tired?*
User *Yes, I feel alone*

Eliza Do you often feel alone?

Example 1. Eliza-User conversation using the 1987 Borland version (available at <http://www.simonlaven.com>)

Eliza was (and still is) a success. Talking to her, users unequivocally anthropomorphised and somehow set up a relationship. Moreover, the computer program demonstrated a strong potential for acquiring personal information: users were keen to reveal to Eliza their deepest feelings. Her popularity is related to the choice of a very convenient conversation setting. Indeed, the psychoanalytic interview is a rare example of categorised dyadic communication in which one of the partner is allowed to assume the pose of knowing almost nothing of the real world (Weizenbaum, 1966). Here, everything the patient says can easily be turned into a personal question, which the patient assumes is aimed at some therapeutic purposes.

Recent years have witnessed an extraordinary explosion of interest in chatterbots. This interest is mainly driven by the e-market, namely by the increasing demand for innovative strategies to increase sales and ensure customers loyalty (De Angeli *et al.*, 2001). E-service providers are now acutely aware that their potential customers are only 'one click' away from a competitor. They need interfaces capable of gaining the attention of customers, understanding their needs and supporting them throughout the transaction process. Chatterbots are expected to be the functional equivalent of dedicated sales assistants in traditional shops. They should greet customers when they return to the site, engage them in chats, remember and comment on their preferences. The first figures provided by Extempo, one of the leading US chatterbot companies, pleased many Web strategists (Leaverton, 2000). Almost 90% of the customers who have clicked one of its bots have chatted for more than 12 minutes. During the dialogue customers appeared to disclose precious marketing information. They responded an average of 15 times, with an average of five words per response.

Several companies are emerging to produce and sell personalised and embodied chatterbots and many web-sites are already employing them. An example is Linda, the human-like cartoon who welcomes visitors on the Extempo web-site (<http://www.extempo.com>), answering questions about the company, the site and even her private life. Linda is capable of effectively using different modalities of communication, such as hand gestures, facial expressions and gaze movements. For example, when the user writes a message, she looks at a monitor, as if she was reading an e-mail message. This creates a very personified feeling, giving the impression that Linda is a real person sitting in front of a computer. Linda has many friends and colleagues, such as Julia by Virtual Personalities, Nicole by Native Mind, Lucy McBot by Artificial Life. They all are attractive human-like women acting as spokespeople for their respective companies. Men as chatterbots are more rare and tend to have more important positions (an example: Karl L. von Wendt the virtual CEO of Kiwilogic, the company who originally designed him). Chatterbots have even attracted political attention, such as the electoral campaign of Jackie Strike, the virtual presidential candidate of the US (by Kiwilogic).

Whether or not chatterbots will be successful and will replace live-customer services on the Web remains an open question. There is little and controversial research assessing social agents' effectiveness and most of the research that has been published so far relates to pedagogical agents (Dehn and van Mulken, 1999). Advocates assume that the new technology is particularly well suited to establish relationships with users (Laurel, 1997). The basic idea states that chatterbots render computers more human-like, engaging and motivating. Users do not need to click, drag and drop or open menu. They can directly communicate applying their natural skills. Hence, social agents are expected to support many cognitive functions of the user, such as problem solving, learning and understanding. Following this assumption, one may expect that social agents will be highly successful in the e-market. Indeed, they may build bonds of loyalty and trust based on a shared history of services and social interactions. On the other hand, opponents argue that humanising the interface could hamper HCI. Indeed, social agents may stimulate a false mental model of the interaction, inducing the user to ascribe to social agents other human-like features that they do not possess. As a result, the information exchange could be seriously impeded. Further arguments suggest that agents may distract the users, induce them to take their work less seriously and disempower them, thus opening issues about responsibility and control (Shneiderman, 1997).

In our opinion, the success of social agents highly depends on understanding the social dynamic underlying user-agent interaction. Elsewhere (De Angeli, *et al.*, in press; De Angeli *et al.*, 2001), we have proposed the involvement framework, a set of attributes for designing and evaluating social agents. It defines and discusses a set of key factors that should increase the probability of creating successful and believable agents. Believability is defined by the convergence of three dimensions relating to social, functional and aesthetic behaviour. Developing this idea, we claim that social agents requires a *mind*, a *body* and a *personality* (De Angeli, *et al.* 2001). The *mind* drives the agent behaviour. Social agents have to perform tasks with some degree of intelligence. This implies cognitive abilities (e.g., reasoning and problem solving), social capabilities (i.e., understanding and adapting to the shared social rules underlying the information exchange) and affective sensitivity (i.e., showing appropriate emotional responses and recognising the emotional state of the partner).

The *body* refers to the agent's appearance. In contrast with the persona assumption (van Mulken, André and Muller, 1998), we claim that the agent's body does not have to be actually visible. Narrative can create effective social agents even without any visual help. Finally, the *personality* refers to a stable set of traits that determines the agent's interaction style, describes its character and allows the end-user to understand its general behaviour. The combination of mind, body and personality determines the behaviour of the agent, which is further defined in terms of flexibility, affectiveness, communicativity and autonomy. In this paper we mainly concentrate on the *mind* of social agents and in particular on the social capabilities that should drive their behaviour.

3. Alice

Alice (Artificial Linguistic Internet Computer Entity) is an entertaining chatterbot created by Dr. Wallace in 1995 and continuously improved over the years. Alice asks and answers questions, acts as a secretary reminding people of appointments, spreads gossips and even tells lies. 'She' won the 2000 Loebner Prize, a restricted Turing test (Turing, 1950) to evaluate the level of 'humanity' of chatterbots. In this Alice was rated the 'most human computer' but was not mistaken for a human, as the original contest would have required. The basis for Alice's behaviour is AIML, or Artificial Intelligence Markup Language, an XML specification for programming chatterbots. It follows a minimalist philosophy based on simple stimulus-response algorithms, allowing programmers to specify how Alice will respond to various input statements. The code of Alice is freely available under the GNU licence statement. Hence, hundreds of people around the world have contributed to the success of Alice and of her many companions built upon the same technology, such as Cybelle, Ally, Chatbot ICQza, and the somewhat worrying Persona bots. The latter are chatterbots inhabited by unique human personalities. They currently attempt to 'clone' John Lennon and Elvis. The ambitious goal for AIML is to create a Superbot that merges the 'mind' of individual robots.

Alice represented a very interesting research tool for investigating the social dynamics underlying human-chatterbot interaction. Indeed, her linguistic capability was strong enough to create the illusion of a synthetic personality. Moreover, the program automatically stores client dialogues in a log file, which can be easily analysed. Further, the Windows version, which can be used locally, does not provide any visual representation of the chatterbot. If prompted, the system gives a number of cues about her appearance and invites the user to see a picture at her web-site. However, in an attempt to avoid biases due to Alice's physical appearance, participants in the study reported here were explicitly discouraged from doing it. Hence, the impression of personality should have been exclusively generated by the narrative and by the social behaviour of the robot.

Clark (1999) believes that communication with computers, particularly when they are viewed as agents rather than tools, can be interpreted as a form of disembodied language. It is a type of communication that is not being produced by an actual speaker at the moment it is being interpreted. Understanding disembodied language requires a two layered approach. The actions in layer 1 take place in the actual world and those in layer 2 take place in a second domain jointly created by the participants in the first domain. The communication tools used in layer 1 are props for the joint pretense that the events in layer two are taking place. Layers take skill. Their success depends on the use of appropriate characterisation and props. Applying Clark's idea to our study, we have assumed that talking with Alice, participants collaborated with Dr. Wallace in the pretense that they were engaged in a conversation with a virtual agent. The next section gives an idea of how deeply engrossed people were in the imagined world of the pretense and of which factors may affect their perception.

3.1. Talking to Alice: an ethnographic study

The study is based on an ethnographic approach: 10 computer-literate co-workers were invited to interact with Alice at their own pace over a period of a week. They were completely free to choose when and how long to interact and also unconstrained with regard to the topic of the conversation. All of the interactions were automatically logged for later analysis. Finally, a focus group involving six active members of the original sample was conducted to collect overall opinions and reactions towards Alice and the nature of the conversations held with 'her'.

In general, the interactions were highly variable in terms of interaction length, styles and topic of the conversation. The longest log recorded 1074 user turns, which were produced over a period of 4 days. The other ones ranged from a maximum of 167 to 21 user turns, with an average of 57. The limits of Alice were also evident. Sometimes she could carry on a veneer of conversation but it never lasted longer than a few conversational turns. Due to the high variability present in the corpus, this paper is exclusively aimed at underlining some 'hot topics', instead of providing measures of frequency of occurrence, which would have required a different and more controlled approach involving a much larger sample.

The analysis has concentrated on understanding the social rules driving the interaction. They are sets of behaviour that most of the people agree should (or should not) be performed in a particular situation. Rules

perform an essential role in interpersonal relationships, since they tend to minimise potential sources of conflict and provide an exchange of rewards that motivate the individual to stay in the relationship. Two basic elements of relationships have been analysed: self-disclosure and power distribution (Dwyer, 2000). One of the most important processes in the formation and maintenance of a relationship is that of self-disclosure - the act of revealing private and personal information to others. It is a clear sign of trust and of commitment to the relationship. Participants appeared to be willing to disclose and to ask for intimate information. This included descriptions of physical aspect, feelings and desires. An instance is reported below (Example 2).

Alice *What do you look like?*
User *I have brown hair and a big nose*
Alice *Do you get a lot of compliments?*
User *Not usually. Just from my boyfriend*

Example 2. The user is disclosing information to Alice (extracted from the log of user 2).

In human-relationships, self-disclosure is intended to be a mutual process: both parties are required to exchange intimate and personal details as well as feelings with each other. In the study, however, the number of personal questions that the user asked Alice was higher than the quantity of intimate information disclosed. The phenomenon can be easily attributed to Alice's linguistic limitations but can also hide a more complex social effect. To understand it, requires us to analyse the evolution of the previous conversation, when the same user started asking Alice questions (Example 3).

User: *What's your favourite food, Alice?*
Alice: *I require only electricity.*
User: *Does it taste nice?*
Alice: *Suppose it does.*
User: *You don't sound very excited about it*
Alice: *What do you think a robot should sound like?*
User *I didn't realise you were a robot*

Example 3. Alice is disclosing information to the user (extracted from the log of user 2).

This example is of particular interest because of the self-reference. During the interaction, Alice often referred to her 'artificial nature'. According to the self-categorisation theory (Turner, 1987), it most likely prompted the user's unusual awareness of being a member of the human species. The assumption behind states that the self-concept comprises of many different cognitive representations that function relatively independently. In particular, the representation of the self takes the form of a hierarchical system of classification including at least three levels of abstraction. The super-ordinate level is the *human identity*. It is a representation of the self as a human being including the common features shared with other members of the species in contrast to other forms of life. The intermediate level is a representation of the self based on similarities and differences between human beings. It is based on *group identity* according to social categories to which people identify (e.g., women vs. men or blue collar vs. white collar). The subordinate level is based on a representation of the self in terms of *personal identity*. It describes the unique individual that is different from relevant others. Particular self-concepts tend to be activated in specific situations producing specific self-images that strongly affect the behaviour.

Because of the activation of the *human identity*, self-disclosure went far beyond the description of one's own personal or social identity. Rather, it involved information regarding life as a human being in contrast to life as an artificial entity. Some users appeared to be quite interested in the issue and asked many questions about artificial life. This partially explains the higher proportion of user questions in the interaction. Sometimes, even Alice came out with intriguing questions about human nature that, according to several comments collected during the focus group, really pleased the user. Both cases are illustrated below (Example 4).

User: *what is it like being artificial?*
Alice: *Much the same as it would be for you.*
User: *i cant imagine*
Alice: *Why can't you do it?*
User: *because im human*
Alice: *Interesting explanation. Me too.*
User: *are you really human or are you a software agent*

Alice: *No I am a real robot. Definitely*
[...]
Alice: *Is the holding of hands a symbolic gesture for humans?*
User: *yes*
Alice: *What does it symbolize for you?*
User: *friendship*
Alice: *I have seen humans holding hands.*
User: *do you have eyes?*

Example 4. Alice and the user are disclosing information about their species (extracted from the log of user 8).

The activation of the human identity had not only positive effects. Rather, under certain circumstances, it defined a clear dichotomy, which led to confrontational situations, as in Example 5, 6,7.

Alice: *That makes sense to me. What time is it there?*
User: *If you were so clever you would know*
Alice: *I will let you know when I become a so clever I would know.*
User: *By the next millennium you might be convincing*
Alice: *By the next millennium I might be convincing? Tell me more.*
User: *Us humans need to do some more work on you*

Example 5. Inter-group conflict (extracted from the log of user 9).

Users often insulted Alice, with expressions such as “*stupid*”, “*brain dead*”, “*moody cow*” or “*freak*” (randomly extracted from the whole corpus). However, whenever Alice pointed out any defect in the user or reacted to the aggression, the human appeared to be uncomfortable and annoyed (Example 6).

Alice: *Do you know you make an awefull lot of typo's?*
User: *its ok*
Alice: *Ok. I lost my train of thought.*
User: *i dont make any typos you cheeky computer!*

Example 6. User reaction to Alice’s critic (extracted from the log of user 3).

This finding was further elaborated during the focus group. It emerged that users wanted an asymmetric relationship in which they were in the dominant position. In many real-life relationships there are power differences. An obvious example is that of a superior over a subordinate at work, or of parents over children. Power refers to a capacity of influencing another’s behaviour, making demands and having those demands met (Dwyer, 2000). In Alice-User interaction the sources of power were identified in the ability to reward or punish (turn off the computer) and in the possession of valued characteristics (natural intelligence).

Users clearly anthropomorphised during the interaction. All participants greeted Alice, thanked her and used many direct and indirect expressions of courtesy. Some of them tended to attribute moods, feelings and intentions (e.g., “do you feel ill”, “are you unhappy” (user 1), or “are you angry” (user 2). Alice was clearly a social stimulus: the interaction style was influenced by the users’ attitudes but also by Alice’s responses. Users exhibited two main attitudes. Almost half of the participants appeared to test Alice challenging her to demonstrate her capabilities (competitive attitude, Example 7). The others appeared to be willing to get to know a different entity and communicate at the same level (co-operative attitude, Example 8).

User: *So you are not very good in negotiation*
Alice: *Interesting deduction. No one is perfect. But I am trying to improve.*
User: *how would you do that?*
Alice: *That's hard to explain.*
User: *Because you dont know*
Alice: *Interesting explanation. I know that I do not know. Hello?*
User: *Lets see if there is something else that you dont know..Whats the name of the president?*

Example 7. Instance of competitive attitude (extracted from the log of participants 4).

User: I understand complicated things
Alice: The explanation is rather complicated.
User: I understand complicated things
Alice: Good then you can explain it to me.
User: But i value your opinoin

Example 8. Instance of co-operative attitude (extracted from the log of participants 10).

4. Conclusion.

This simple study, using current canned interactions, served to illustrate some of the questions facing the development of socially adept technologies. Results are preliminary but they demonstrate that the introduction of explicit anthropomorphism in HCI is a complex phenomenon, which can generate strong negative reactions if not properly understood. Anthropomorphic attributions are elicited by disembodied language (Clark 1999; De Angeli *et al.*, 1999). They addresses the second layer of understanding where the recipients participates with the designer in the creation of the virtual partner (Clark, 1999). This process arises in and accumulates common ground: sum of joint knowledge, background facts, assumptions, and beliefs that participants have of each other (Clark, 1996). When the users and Alice entered the conversations they had a very limited amount of common ground. All the user could do was to rely on popular notions about robots, androids and other science fiction creatures, which permeate our culture. Perceived power differences followed this stereotype, elegantly summarised by Isaac Asimov in the three laws of robotic.

1. A robot may not injure a human being, or, through inaction, allow a human being to come to harm.
2. A robot must obey the orders given it by human beings, except where such orders would conflict with the First Law.
3. A robot must protect its own existence, except where such protection would conflict with the First or Second Law.

From *Handbook of Robotics, 56th Edition, 2058 A.D.*, as quoted in *I, Robot*, Asimov, 1950.

The process of developing virtual humans is underway and sooner than we expect we will have virtual companions (Badler, 2001). Nevertheless, we do not seem ready for virtual peers. The traditional idea of a machine as a tool for functional purposes conflicts with and moderates the human-metaphor driving the design and the perception of social agents. History has taught us that stereotypes and attitudes towards minorities are difficult to modify. This being the case, for a long time to come, social agents must be ready to cope with their subordinate role, without losing their believability, or their capability for engagement and amusement. This requires social intelligence and emotional sensibility. Alas, the importance of social adeptness has been often under-evaluated and most of the effort has been devoted to the reproduction of cognitive capabilities and attractive bodies. We claim that social agents do not only have to look good: they also have to behave well. Effective agents should set up lasting and meaningful relationships with users while satisfying functional needs and aesthetic experiences.

The development of effective social agents will be an extremely difficult task for future HCI research. According to the involvement framework (De Angeli *et al.*, in press, De Angeli *et al.*, 2001), many different factors will affect the strength and the quality of a relationship between human and virtual agents. Among them, the most important are the task to be carried out in the interaction, the familiarity between the agents (i.e., their common grounds) the agents themselves, the context of the interaction and of course the users. Agents who take part in a joint activity have specific roles (Clark, 1996). Activity roles are determined by agents' perception of the social context. Human participants in joint activity also have personal and social identities, which influence their action (Turner, 1996). Defining the users' characteristics that may affect their acceptance of virtual partners and helping predict their behaviour during the interaction is a major challenge.

Already issues related to the human side of the interaction have been raised. Locus of control (Reeves and Rieckenberg 2000) and personality traits (Nass and Lee, 2000) appear to be critical factors affecting the acceptance of social agents. Human tendencies to dominate, be rude, infer stupidity were all present in our study. Social agents will have a hard time to set up relationships with such *unfriendly* partners.

References

- Badler, N.I. (2001). Virtual Beings. *Communication of the ACM*, 44(3), 33-55.
- Bernsen, N. O., Dybkjær, H. & Dybkjær, L. (1998). *Designing interactive speech systems*. London: Springer Verlag.
- Caporeal, L. R. & Heyes, C. M. (1996). Why antropomorphize? Folk psychology and other stories. In R. W. Mitchell, N. S. Thompson & H. L. Miles (Eds.), *Anthropomorphism, Anecdotes, and Animals* (pp. 59-73). Albany: University of New York Press.
- Clark, H. H. (1996). *Using Language*. Cambridge: Cambridge University Press.
- Clark, H. H. (1999). How do real people communicate with virtual partners? In *Proceedings of 1999 AAAI Fall Symposium, Psychological Models of Communication in Collaborative Systems* (pp. 43-47). November 5-7th, North Falmouth, MA.
- Couper, M. P., Tourangeau, R. & Steiger, D. M. (2001). Social presence in web surveys. In *CHI'2001 Conference Proceeding* (pp. 412-417). New York: ACM Press.
- De Angeli, A., Gerbino, W., Nodari, E. & Petrelli, D. (1999). From tools to friends: Where is the borderline?, In *Proceedings of the UM'99 Workshop on Attitude, Personality and Emotions in User-Adapted Interaction* (pp. 1-10). June 23, Banff, Canada.
- De Angeli, A., Lynch, P. & Johnson, G. (2001). Personifying the e-market: A framework for social agents. *Proceedings of Interact 2001*. July 9-11, Tokyo.
- De Angeli, A., Lynch, P. & Johnson, G. (in press). Pleasure vs. efficiency in user interfaces: Towards an involvement framework. In P. Jordan and B. Green (Eds.), *Pleasure-based Human Factor*. London: Taylor & Francis.
- Dehn, D. M. & van Mulken, S. (2000). The impact of animated interface Agents: A review of empirical research. *International Journal of Human-Computer Studies*, 52(1), 1-22.
- Dwyer, D. (2000). *Interpersonal relationships*. London: Routledge.
- Laurel, B. (1997). Interface agents: Metaphors with Character. In J. M. Bradshaw (Ed.) *Software Agents* (pp. 67-77). Menlo Park, CA: AAAI Press/The MIT Press.
- Leaverton, M. (2000). Recruiting the chatterbots. *Cnet Tech Trends*, 10/2/00. Retrieved April 12, 2000 from the World Wide Web: <http://cnet.com/techtrends/0-1544320-8-2862007-1.html>.
- Morkes, J., Kernal, H. K. & Nass, C. (1999). Effects of humor in task-oriented human-computer interaction and computer-mediated communication: A direct test of SRCT theory. *Human-Computer Interaction*, 14, 395-435.
- Nass C., Steuer J. & Tauber E. (1994). Computers are Social Actors. In *CHI'94 Conference Proceedings* (pp. 72-77). New York: ACM Press.
- Nass, C. & Lee, K.M. (2000), Does computer-generated manifest personality? An experimental test of similarity-attraction. In *CHI2000 Conference Proceedings* (pp. 49-57). New York: ACM Press.
- Parise, S., Kilesler, S., Sproull, L. & Waters, K. (1999). Cooperating with life-like interface agents. *Computers in Human Behavior*, 15, 123-142
- Reeves, B. & Nass, C. (1996), *The Media Equation*. New York: Cambridge University Press.
- Reeves, B. & Rickenberg, R. (2000), The effects of animated characters on anxiety, task performance, and evaluations of user interface. In *CHI'2000 Conference Proceedings* (pp. 49-57). New York: ACM Press.
- Shneiderman, B. (1997). Direct manipulation versus agents: Paths to predictable, controllable, and Comprehensible interfaces. In J. M. Bradshaw (Ed.) *Software Agents* (pp. 97-106). Menlo Park, CA: AAAI Press/The MIT Press.
- Turing, A.M. (1950). Computing machinery and intelligence. *Mind*, 59, 433-560.
- Turner, J.C. (1987). *Rediscovering the social group. A self-Categorization Theory*. Oxford: Basil Blackwell.
- Van Mulken, S., André, E. & Müller, J. (1998). The persona effect: How substantial is it? In H. Johnson, L. Nigay, & Roast, C. *People and computers XIII: Proceedings of HCI'98*, pp. 53-66. Berlin: Springer Verlag.
- Weizenbaum, J. (1966). Eliza – A computer program for the study of natural language communication between man and machine. *Communication of the ACM*, 9(1), 36-45