

Abuse: the dark side of human-computer interaction

Co-Chairs: Antonella De Angeli¹ Sheryl Brahnam² Peter Wallis³

¹ Manchester University

`Antonella.De-angeli@manchester.ac.uk`

² Southwest Missouri State University

`research@knockloudly.com`

³ Sheffield University, NLP Group

`pwallis@acm.org`

There seems to be something innate in the human/computer relationship that brings out the dark side of human behaviour. Anecdotal evidence suggests that of the utterances made to chat-bots or embodied conversational agents (ECA) in public places, 20-30% are abusive. Why is that? Is it simply that a quarter of the human population are 'sick' and find abusing a machine to be in some way therapeutic? If so, it says something about human nature that is quite disturbing and in need of further study. Perhaps the phenomena is directly caused by a new technology. In the early days of computer-mediated communication there was a tendency for people to abuse each other, but this has become far less common. Will the extent to which people abuse ECAs just naturally become a thing of the past? The Turing Test has also had a considerable influence on appropriate behaviour when talking to a computer. To what extent is abuse simply a way people test the limits of a conversational agent? Perhaps the problem is one of design: the aesthetics of everyday things is key to their success, and perhaps abuse is simply one end of a continuum of the 'aesthetics' of interactive things with human-like behaviour. The extent to which abuse occurs seems to indicate something fundamental about the way humans interact. Is abuse simply the most noticeable phenomena connected to something fundamental in the way we humans communicate?

The purpose of this workshop is to bring together engineers, artists and scientists who have encountered this phenomenon, who might have given some thought to why and how it happens, and have some ideas on how pro-active, agent-based interfaces, typified by chat-bots or ECAs, should respond to abuse.

Agent-based interfaces

Unlike conventional software, agent-based interfaces can be expected to deal with abuse. Whereas most software is designed as a tool to be used by a human, conversational agents convey a sense of autonomous action. With the computer-as-tool metaphor, the human operator is responsible for outcomes. One cannot really blame the hammer when one hits one's thumb, and one cannot, really, blame MS-Word when one hasn't read the manual. With the agent metaphor of software, the agent is in some sense aware of its surroundings (situated) and

responsible (autonomous) for its behaviour. An ECA can be expected to be proactively helpful - they can be expected to take responsibility in some sense for the user's experience. If the user is expecting this, and he or she starts abusing the system, then the system's response will be interpreted as a social act. As Reeves and Nass point out in "The Media Equation" (CSLI Publications, 1996) we often treat machines as if they are in some way human. Whereas we know we are playing make believe when we abuse a hammer for a sore thumb, we are often not aware of anthropomorphic behaviour when dealing with computers. In some contexts, ignoring abuse simply encourages more. Are the ECA that ignore abuse simply "pushing our buttons" and encouraging more?

ECA, by their very nature, interact with us at a social level and must play by our rules. It is hard for us as social actors to separate the machine's actions from those we see other people use everyday. Consequently, until we have a better understanding of the relationship between human and virtual agents, the commercial potential of such agents is questionable. A virtual personal assistant for instance cannot simply provide timely information; if that information is to be believed, the agent must present itself as a reliable source. Although we might know, in a conceptual way, that the interface does not change the quality of the data from a particular source, we cannot help but respond to the conventions used by the human agents we have grown up with. If abuse, or the threat of it, are part of those conventions, then a trusted virtual assistant will need to be able to play that game. Similar concerns arise using ECA on corporate web-sites. To what extent does the ECA represent the organisation? If potential clients are abusing your virtual representative, then to what extent are they abusing your organisation? How can the agent change the views of clients, and to what extent is the process not about providing information, but about behaving well on the social front? Those interested in automated learning and virtual tutors must also consider the social skills relevant to the process of imparting knowledge. If the student's interface is perceived as a social actor, and the student has no respect for that agent, then to what extent might a student be expected to learn? In the other direction, virtual characters in computer based games might benefit from being able to generate abusive behaviour at an appropriate level. To what extent is abuse a means of social positioning and a key process in the establishment of identity and group membership? Games such as "Leisure Suit Larry" are explicitly about social relations but the behaviour of the characters involved is, like that of the monsters in "Quake," obviously not human. A better understanding of how we humans manage our relations could be key to the next generation of products for this multi-billion dollar industry.

Abuse is prevalent and easily detected in human computer interfaces that use conversational agents. Studying the phenomena we hope will lead to a better understanding of human behaviour in a social context, and to human-computer interfaces that interact in a natural manner. What is natural however needs more study, and the way people abuse computers is hopefully a key to next generation of pro-active computer interfaces.