

When is a Communicative Agent a Good Idea?

Patrick Doyle

Knowledge Systems Laboratory
Gates Computer Science 2A
Stanford, CA 94305-9020
+1 82 (650) 723-6707

pdoyle@cs.stanford.edu

ABSTRACT

Research suggests humans treat computers in social ways. We understandably have an interest in building interfaces that use embodiment and natural language to conform to these social roles. However, we must be careful how and when we embed these agents so they actually improve on existing interfaces, rather than simply making them more social.

1. INTRODUCTION

The history of human evolution has been one of adaptation both to the environment and to ourselves as a species. As our brains grew more subtle, so did the mechanisms by which we communicate with one another. Our verbal and non verbal behaviors are wonderfully expressive and rich with meaning.

Interactions with machines, in contrast, are crude and rather stilted. Once again we are forced to adapt, this time to the computing environment, traditionally incapable of sending or receiving communications along the channels to which we are naturally accustomed. Present interfaces rely on decades-old designs that are ergonomic and efficient, but only narrowly communicative.

The work of Reeves and Nass [6] has been greatly influential for researchers seeking new interfaces. They argue, plausibly, that because humans have both evolved and learned a certain system of interaction with other intelligences – a system of *social interaction* – that we naturally (if unconsciously) attempt to employ that system when dealing with our computers. That is, we implicitly have expectations about the social and emotional behavior of machines, and even treat them in social ways, though as a rule these expectations are unmet.

An obvious inference to draw from this is that we can improve our interfaces by making them *explicitly* social. In this way they can adjust to human behaviors rather than requiring that humans adapt to them. Most researchers attempt this by creating embodied agents, so that rather than communicating with the “ghost in the machine,” instead we see an animated character that behaves in recognizably human or lifelike ways. These attempts have met with mixed results.

2. CHARACTER IN CONTEXT

Education and entertainment are the two domains in which embodied characters seem most likely to fit and, indeed, it is in these areas where most of the successes have come.

Lester, for example, has built impressive animated characters as tutors in educational software programs that teach domain knowledge in a constructivist environment [4]. Experiments with

these programs have shown that the mere presence of the animated tutor makes children more attentive and may even improve retention [5]. The STEVE project [7] has studied the use of a talking, embodied agent that teaches users how to perform physical tasks by demonstrating them in a virtual environment. Doyle and Isbister [1] and André and Rist [8] are experimenting with Web-based characters as tour guides and lecturers. [2] and [3] give an overview of current work on such characters.

There has been less recent work on entertainment agents. Although the computer game industry is eager to build games in which intelligent characters interact with the players, the difficulties of building agents that behave intelligently under the constant and prolonged attention of the user has hampered research in this area. Generally, pre-scripted characters or agents with very narrow abilities are used instead, though even these crude approximations have been successful for years.

Embodied agents in work environments have fared still less well. The notable example here is the Office Assistant in Microsoft's Office package. The assistant occupies a small window in a corner of the screen from which it makes occasional task-specific suggestions, provides an interface to the help system, and gives status information through its animations and dialog bubbles. Unfortunately, instead of delivering any of the benefits we might hope to get from these communicative agents – benefits such as increased comfort with the system, increased ease of use, more efficiency in accomplishing one's tasks, or simply more pleasure in performing them – the conventional wisdom¹ is that most users find these agents annoying, intrusive and distracting, and disable them far more often than not.

As I see it, the problems with the Office Assistant are neither that embodied agents are fundamentally a bad idea nor that these particular agents are badly designed. The problem is that they were chosen for the wrong domain. Building a spreadsheet, for example, is essentially a mechanical task – entering numbers and equations – and not one about which a user is likely to want a discussion, especially not with the spreadsheet software itself.

3. QUALITIES OF INTERACTION

From these and other examples, we can make some hypotheses about where communicative agents will work. They depend on the nature of the interaction between the user and the computer.

¹ I blush to cite “conventional wisdom” as a source, but despite being the most widely-used embodied agents, no formal analysis of users' reactions is available. My informal conversations with users, with some of Nass and Reeves' students, and with the original character designers support this assertion, however.

Uncertain vs. decisive. Communicative agents are more valuable to users who are attempting to make choices, particularly choices about which they are unsure. Here the agent is a guide or advisor, making suggestions, offering advice, and presenting facts to assist the user. A lifelike agent, with personality and emotions, can sympathize with and even reassure the user, especially if it takes on an authoritative role in the domain. Most educational software falls in this category; the user is learning what to do, and the agent is providing help and feedback during the learning process.

Satisfying vs. correct. For certain tasks it is more important to be pleased with the result than for the result to be right. Communicative agents can help produce a satisfying result, both by changing the nature of the task and making it more pleasant, and by encouraging the user's belief that the final product is acceptable. This is obvious for entertainment; interacting with our agents is part of the story, and the story is enjoyable to the extent that the interaction is². Commercial examples arise in areas where the user needs expert advice about subjects where preference is more important than substance, such as home decorating: chintzy curtains might be more palatable when recommended by an interactive Martha Stewart.

Dialog vs. command-driven. When the user knows both the nature of the task and how to accomplish it, the job of the interface is to make that as fast and painless an operation as possible. In such a circumstance we want to approximate a *direct manipulation* interface, one that is transparent to the user. For most non-entertainment tasks (since in a game interacting with an agent may well be the goal), social interaction is anything but transparent. It involves explanations, clarifications, redundancy in instructions – all things that are valuable if teaching someone else about something, but intensely intrusive if trying to do it oneself.

The communicative agent is far more appropriate for dialog. The value of conversation lies in *synthesis*; through conversation we clarify our thoughts, we receive criticism and suggestion, we learn new information, our perceptions change. Adventure games offer an example where both dialog and command occur simultaneously; in Infocom's *Planetfall*, for example, talking with and listening to your robot sidekick Floyd is a great deal of fun, but giving him instructions or asking for help is rather tedious and frustratingly open to error. A good interface for dialog, but poor for issuing direct commands.

Social vs. natural. We should be careful to make a distinction between a communicative agent and one that happens to use natural language to communicate. The term "communicative agent" implies, for me, a social communication – a conversation. If we can ask our system, "What is the current temperature in Palo Alto?" and it tells us, "The temperature in Palo Alto is now 68 degrees," it is a natural language interaction, but the qualities we associate with conversation are absent. Conversation implies an exchange of information and ideas, is literally "keeping company with" one another. This is highly appropriate for entertainment, where the agents are meant to keep us company, and for

education, where this social quality can make learning more interactive and more engaging. It can be appropriate for other tasks, but they must be such that the users would *want* to talk to a human being in order for a social interface to be worthwhile. When calling a customer support number with a problem, a social interface (as opposed to an endless list of numeric menus) is critically important to the user's satisfaction, but for filling up the cells in a spreadsheet, although natural language might be desirable, conversation is not.

4. CONCLUSIONS

Whether our agents are truly intelligent or only limited approximations to intelligence, we must remember that the point of adding an embodied character is to improve the experience for the user – and more sophistication is not invariably better. Here we've considered some guidelines for using a communicative agent, and I suggest that social tasks, uncertain goals or goals of satisfaction rather than correctness, and dialog rather than command-structured interactions are all good indications that a communicative agent may work well.

5. ACKNOWLEDGMENTS

This work was supported by NSF Grant IRI-9424261, gifts from Microsoft and Intel, and an Intel Foundation Fellowship. I thank Katherine Isbister for her valuable comments.

6. REFERENCES

- [1] Doyle, P., and Isbister, K. Touring machines: Guide agents for sharing stories about digital places. Submitted to *AAAI Fall Symposium on Narrative Intelligence*.
- [2] Elliott, C. and Brzezinski, J. Autonomous agents as synthetic characters. *AI Magazine*, Summer 1998, 13-30.
- [3] Hayes-Roth, B., and Doyle, P. Animate characters. *Autonomous Agents and Multi-Agent Systems*, 2 (1998), 195-230.
- [4] Lester, J., and Stone, B. Increasing believability in animated pedagogical agents. In *Proceedings of 1st Intl. Conf. on Autonomous Agents (Marina del Rey CA, Feb. 1997)*, ACM Press, 16-21.
- [5] Lester, J.; Converse, S.; Kahler, S.; Barlow, T.; Stone, B.; and Bhogal, R. The persona effect: Affective impact of animated pedagogical agents. In *Proceedings of CHI '97 (Atlanta GA, Mar. 1997)*, ACM Press,
- [6] Reeves, B., and Nass, C. *The Media Equation: How People Treat Computers, Televisions, and New Media Like Real People and Places*. Cambridge UP: New York, 1996.
- [7] Rickel, J., and Johnson, W. Integrating pedagogical capabilities in a virtual autonomous agent. In *Proceedings of 1st Intl. Conf. on Autonomous Agents (Marina del Rey CA, Feb. 1997)*, ACM Press, 30-38.
- [8] Rist, T., André, E. and Müller, J. Adding animated presentation agents to the interface. In *Proceedings of IUI '97 (Orlando FL, Jan. 1997)*, ACM Press, 21-28.

² Presto Studios' games *The Journeyman Project 2* and *3* feature an "AI" sidekick named Arthur. This character, although entirely scripted, is the most engaging example I've seen of what a truly believable agent might do in entertainment or educational software.