

Running head: BEYOND BUTLERS: INTELLIGENT AGENTS AS MENTORS

Beyond butlers: Intelligent agents as mentors

Journal of Educational Computing Research, Vol 22(4), 2000, pp. 373-382.

Amy Baylor, Ph.D
Assistant Professor of Instructional Systems
Department of Educational Research
Florida State University
305B Stone Building
Tallahassee, FL 32306
(850) 644-5203
Email: baylor@coe.fsu.edu

Abstract

This paper discusses pedagogical issues for intelligent agents to successfully serve as mentors for educational purposes. Broader issues about the nature or persona necessary for an intelligent agent as mentor are discussed, incorporating usability and human-computer interaction issues such as the anthropomorphic qualities of the agent and the social relationship between learner and agent. Overall, to be effective for learning, it is argued that there are three main requirements for agents as mentors: 1) regulated intelligence; 2) the existence of a persona; and, 3) pedagogical control.

Introduction

Imagine having a super-smart friend with whom you could study at any time of the day, who was always there to encourage and guide you. Or imagine having a job coach who could lead you through a difficult work situation. Imagine now that the friend is not a human friend, but rather a computer program called an intelligent agent. People have always been fascinated with the idea of non-human computer assistants: androids, humanoids, robots, and science fiction, the mechanical maid in the cartoon “The Jetsons.” But what about the possibility of a personalized computer-based instructor, mentor, personal coach? While we are not there quite yet in terms of technical capabilities, we are continually improving artificial intelligence techniques. This paper will discuss important human-like pedagogical characteristics for intelligent agents to successfully serve educational purposes, covering broader issues about the persona necessary for an intelligent agent mentor (for a more comprehensive review of intelligent agents for education see Baylor & Jafari, 1999).

What is an intelligent agent?

Intelligent agents are independent computer programs operating within software environments such as operating systems, databases, or computer networks (Roesler and Hawkins, 1994). One typical non-educational use of intelligent agents is to help users find or organize information according to personal specifications. Negroponte (1970) originally described intelligent agents as electronic “butlers,” performing such tasks as filtering email, scheduling appointments, informing regarding investments, and making travel arrangements. For example, an intelligent agent could search on the Internet to find the best price for an old Beatles record, thereby obtaining additional information so that the user does not have to spend the time and effort. Another simple intelligent agent is the personified paper clip “Clippit” in Microsoft Office that follows the user's actions and makes suggestions.

While intelligent agent technology involves artificial intelligence (reasoning, planning, natural language processing, etc.) and system development techniques (object-oriented programming, scripting languages, human-machine interface, distributed processing, etc.), it also requires cognitive understanding of the task at hand. Especially for educational applications, expertise is required in terms of specifying how and what the learner should learn. Only then can intelligent agents provide valuable assistance educationally, serving as cognitive tools (see Baylor, 1999b), while enabling a more humanlike interaction than other computer-based implementations. Along this line, Seiker (1994) describes intelligent agents as computer programs that simulate a human

relationship by doing something that another person could otherwise do for you. Importantly, intelligent agents have been predicted to be one of the most important computing paradigm in the next few years, with some claiming that soon every significant application will have some form of agent functionality (e.g., Janca, 1995). Consequently, the use of intelligent agents for educational purposes is of interest. The development of intelligent agents for learning is also very relevant to current research on electronic performance support systems (EPSS) and online interactive help.

Here it is important to differentiate between two different meanings of an intelligent agent that are often confused: adaptive functionality and the agent metaphor (Erickson, 1997). According to adaptive functionality, agents can serve as autonomous computer programs with intelligence, adaptivity or responsiveness. Their role involves noticing, interpreting, and responding, like Negroponte's original description of agents as "butlers". The agent metaphor, on the other hand, involves what is portrayed to the learner, where the agent is a program that appears to have the characteristics of an animate being. This metaphor suggests a different characterization of the agent, its relation to the learner, and its capabilities and functions. Thus, it is not necessarily preferable to have both the agent metaphor **with** the adaptive functionality; for example, using a meeting scheduling program that "shows" available times with a certain color after discerning what times are available -- no "agent" is needed and would be distracting in this case (Erickson, 1997). Yet for agents as mentors, it is critical for them to have a character that interacts with the learner. Consequently, moving beyond adaptive functionality and agents as "butlers" to the agent metaphor provides opportunities for intelligent agents in education.

This paper will specifically consider *intelligent agents as mentors* by first presenting overall features of an intelligent agent mentor (in terms of its persona, appearance, and social relationship with the learner) and secondly by describing necessary pedagogical requirements.

The persona of the intelligent agent

In terms of the persona of intelligent agents, Laurel (1997, 1990) argued that intelligent agents should be distinctive characters similar to actors in a play. For example, as part of a social studies program for studying the time period of the Civil War there could be a historian agent, an explorer agent, and a politician agent to represent different points of view. According to Laurel, the more explicit their definition, the easier it will be for learners to predict their behavior. If their definition is too transparent, we will likely assume more intelligence than they possess and expect more than they can deliver.

She describes three benefits of this anthropomorphic approach where agents are characters. First, this representation optimally uses our ability to make accurate inferences about how a character is likely to think, decide, and act on the basis of its external traits. Second, the agent as a character invites conversational interaction. Third, the metaphor of character successfully draws our attention to just those qualities that form the essential nature of an agent: responsiveness, competence, accessibility, and the capacity to perform actions on our behalf (Laurel, 1990). She stresses that we want to be able to predict their actions with greater certainty than those of "real" people. For example, increasingly in the world of adventure and role-playing computer games, designers are implementing characters with traits that are dynamic (modified by learning and experience) and relational (modified in relation to objects and situations).

S[Picard, 1997 #795]hneiderman (1997) suggests that maybe it is not good to compare agents to humans since these terms shape the thoughts of everyone regarding agent expectations; thus, limiting them to human terms of ‘intelligent’ and ‘smart.’ On the other hand, Norman (1997) points out that people tend to have overblown expectations about what an agent can or should do. Specifically, some believe that it is wrong -- immoral even -- to offer artificial systems in the guise of human appearance, for to do so makes false promises (Norman, 1997). However, research by Reeves and Nass (1996) suggests that people treat computers as human, even when the computer interface is not explicitly anthropomorphic.

As summarized in Kearsley (1993), Friedman & Kahn (1992) discuss the moral implications of delegating decisions to computers and allowing computers to exhibit human characteristics. They argue that a clear distinction between human and computer actions should be preserved. Such concerns may be reasonable fears, according to Bradshaw (1997): “The more intelligent the robot, the more capable of pursuing its own self-interest rather than its master’s. The more humanlike the robot, the more likely to exhibit human frailties and eccentricities.”(p. 3) Along this line, it is probably not a good idea to provide a user interface agent that is very sophisticated, qualified and autonomous from the start (Baylor, 1999a; Maes, 1994). Such an agent would leave the user with a feeling of loss of control and understanding (Shneiderman, 1992).

But does the learner want agents to be believable, to demonstrate emotion? As stated by Bates (1994) emotion is one of the primary means to achieve this believability, this illusion of life, because it helps us know that characters really care about what happens in the world, that they truly have desires. Lester, Converse, Kahler, Barlow, Stone, & Bhoga (1997) tested the affective impact of animated pedagogical agents by implementing an agent that offered various types of feedback during a learning task. They found that learners preferred the agent in the fully expressive condition (where all types of feedback were present) to all other conditions (which offered limited types of feedback). Further, performance on the learning task was best in the fully expressive condition. However, in a preliminary study, Dietz & Lang (1999) found that while users preferred agents showing more emotion and performed better on a memorization task with the emotion-showing agents, the results were not statistically significant.

How should it look?

Overall, the agent is preferably portrayed with graphic or iconic representations rather than realistic animations or video, although more realistic animations have been explored such as Microsoft’s conversational parrot “Peedy” (Ball, Ling, Kurlander, Miller, Pugh, Skelly, Stankosky, Thiel, Van Dantzich & Wax, 1997), or Lester & Stone's (1997) "Herman the Bug" in the Design-a-Plant learning environment. But Shneiderman (1992) suggests that user interfaces should not attempt to mimic human interaction but maintain a “neutral” status. A good compromise may be the use of different faces, which Maes (1997) describes in an email agent application. Specifically, the agent has different faces to indicate what the agent is doing: thinking, working, suggesting, unsure, pleased, confused. In terms of the value of agents exhibiting life-like expressions, Lester and Stone (1997) and Koda & Maes (1996) suggest that such agents have greater motivational impact.

Social relationship of learner and agent

A critical factor regarding the person-agent relationship includes developing a social relationship of the agent with the learner. This relationship requires the intelligent agent to be perceived by the learner as trustworthy, honest, and cooperative while providing appropriate feedback (e.g., Baylor, 1999a). It is also important that the intelligent agent resemble a human tutor in terms of motivational qualities. As Lepper & Chabay (1987) propose, motivational components are as important as cognitive components for an intelligent tutor. Taking it one step further, they propose that bringing empathy to computer tutors is conducive to learning. However, there is limited empirical evidence to support such claims.

Laurel (1997) points out that human-like qualities are often attributed to the agent, e.g., one interface designer has described error messages as “wrist-slapping grannies.” Specifically Laurel suggests the two distinctly human-like qualities of an agent are responsiveness and the capacity to perform actions. Along this line, Norman (1997) points out the need for the agent-human interaction to provide *reassurance* for the user. To do so, the technology needs to be working seamlessly with the agent reassuring the user that all is working according to plan.

Overall, the learner must feel confident that the agent will perform the desired task and that the agent interpreted the desired task correctly. Or, as stated by Maes (1997), a leader in intelligent agent research at MIT, two key issues are of competence and trust. In her approach the agent gradually develops its abilities so that the user is also given time to gradually build up a model of how the agent makes decisions, thereby improving trust. The agent acquires competence from four sources: monitoring the user, noticing his/her behavior; providing direct and indirect user feedback; training from examples given explicitly by the user; and, asking for advice from agents that assist others with same task.

Serving as a pedagogical mentor

The social relationship between the learner and agent is further defined to be a pedagogical relationship, where the intelligent agent (in theory) can monitor and evaluate the timing and implementation of teaching interventions (e.g. provide help, feedback). The relationship of the intelligent agent to the learner can be described as a cognitive apprenticeship, where the student improves his/her performance while working with the more expert performer: the intelligent agent. In support of this feature, Collins & Brown (1987) suggest that students may learn best in environments including modeling and coaching of formative skills. A favorable situation with intelligent agents might be that as the student gains expertise, the agent would fade and allow for more student initiative. An important consideration in terms of feedback is that the intelligent agent should not provide too many insights and thereby annoy the student. Furthermore, as Negroponte (1997) suggests, the human act of winking can connote a lot of information to others simply in the *lack* of information. This sort of familiarity is needed for the pedagogical agent to avoid relentless explicitness. To address this issue, part of the pedagogical task should include the monitoring of the timing and implementation of the advisements. With the principle of minimal help as the default, there could also be the possibility for the student to select a feedback option depending on the amount of structure, interaction, and feedback s/he desires when problem-solving.

There are significant difficulties in developing pedagogical expertise in an intelligent agent. As McArthur, Lewis and Bishay (1993) state, the pedagogical component of intelligent systems receives relatively little mention with current systems demonstrating little pedagogical expertise. As they suggest, most intelligent tutoring systems are constrained to a single method of teaching and learning, while truly expert human tutors can adopt different methods. Dillenbourg, Mendelsohn & Schneider's (1994) agent-based ETOILE system, however, consists of multiple pedagogical agents taking different approaches according to educational psychology principles. The five teaching agents are named Skinner, Bloom, Vygotsky, Piaget, and Papert where Skinner works step by step, Bloom takes larger steps but with control of mastery, Vygotsky is based on participation, Piaget intervenes only to point out problems, and Papert does not interrupt the learner. Further, their system includes a "coach" agent that manages the roles of the pedagogical agents and controls the interactions among all of them and the learner. Another system under development is Baylor's (1999c) MIMIC (Multiple Intelligent Mentors Instructing Collaboratively) system, which consists of three pedagogical agents (from three different theoretical perspectives) to teach instructional design to pre-service teachers. Specifically, the three pedagogical agents take three different perspectives of instructional design: the *instructional systems design* pedagogical agent; the *alternative views of instruction* pedagogical agent; and, the *teaching as an art* pedagogical agent. In this system, each agent interacts with both the learner and the other pedagogical agents to provide a unique and holistic learning environment for learning instructional design.

The agent needs to explain its reasoning

A critical issue in terms of educational value is in moderating between the agent taking over thinking for the student with the agent training the student to think more effectively. Malone, Lai & Grant (1997) propose sharing the control between agent(s) and learner through a semiformal systems approach. Usually computer systems are at one extreme or the other: either highly structured such as databases with strict requirements and structured procedures, or non-structured such as word processing where the computer's role is to record, store and transmit information without having to "understand" or process the information it stores. In the semiformal systems approach, the information is semi-structured with the reasoning visible to the learner. In other words, rather than creating intelligent agents whose operations are "black boxes" designers should try to create "glass boxes" where the essential elements of the agents' reasoning can be seen and modified by learners (Malone, Lai & Grant, 1997).

Together with the importance of sharing the agent's reasoning with the learner, the issue of feedback is critical (see Baylor, 1999a). As Erickson (1997) points out, the user needs understanding of what happened and why. Keeping this issue at the forefront, Baylor & Kozbe (1998) developed initial specifications for an intelligent agent, the Personal Intelligent Mentor (PIM), that has special potential for tapping learners' metacognitive processing. As a pedagogical expert the intelligent agent could serve as a technological "reciprocal teacher" (e.g. Palinscar & Brown, 1984), prompting the individual to engage in analysis of his/her own cognitive processes. This use of an intelligent agent would serve to encourage the individual to assess what cognitive strategies are being used, similar to Salomon's pedagogic computer program, the Writing Partner (Salomon, 1993), which asks the learner intelligent questions through the writing process.

Conclusion: Beyond Butlers

Agents as mentors for learning are much more than the “butlers” that Negroponte (1970) first described. To be effective for learning, there are three main requirements for agents as mentors: 1) regulated intelligence; 2) the existence of a persona; and, 3) pedagogical control.

First, while agents as mentors should have intelligence (in the artificial intelligence sense), as do other non-educational intelligent agents, it is critical for this intelligence to be moderated differently. While a mentoring agent must demonstrate competence to the learner, in order to be effective the mentoring agent should not be *too* intelligent (e.g., see Salomon, Perkins & Globerson, 1991) because this may lead the user to have unrealistic expectations, a loss of control and limited understanding as to the agent’s reasoning.

Second, agents as mentors must have an educationally appropriate persona. From a motivational standpoint, the agent should provide appropriate feedback and reassurance to the learner. In a relational sense, the agent should be dynamic, trustworthy, honest and cooperative. In terms of how this is operationalized, it seems that the more emotion and expressiveness, the better (even though there is some contradictory empirical evidence). Yet as Kearsley (1993) claims, to be truly effective the agent must also have understanding of the entire personality of the learner (not only just what the learner has learned).

Third, agents as mentors should control the pedagogical interventions. In terms of monitoring and evaluating the timing and implementation of teaching interventions (e.g. providing help, feedback), the agent should not provide too many insights and thereby annoy the student; consequently, setting minimal help as the default is advisable. Additionally, the learner should be given the option to control the level and type of feedback s/he receives from the agent (Baylor, 1999a). Given these considerations, implementing a mentoring agent as a reciprocal teacher, with the learner as a cognitive apprentice is a favorable solution. In this implementation, as the student gains expertise, the agent would fade and allow for more student initiative thereby addressing the need for more flexible control. A promising possibility in terms of regulating the pedagogical interventions is the instantiation of *multiple* pedagogical agents in a learning environment. In this sense, building beyond Laurel’s (1990; 1997) suggestion to have agents represent roles or characters in a play, agents for learning could represent different instructional roles such as in MIMIC (Baylor, 1999c) and ETOILE (Dillenbourg et. al, 1994).

References

- Ball, G., Ling, D., Kurlander, D., Miller, J., Pugh, D., Skelly, T., Stankosky, A., Thiel, D., Van Dantzich, M., & Wax, T. (1997). Lifelike computer characters: The persona project at Microsoft. In *Software Agents*, edited by J.M. Bradshaw (pp. 191-222). Menlo Park, CA: MIT Press.
- Bates, J. (1994). The role of emotion in believable agents. *Communications of the ACM*, 37(7), 122-125.
- Baylor, A. L. & Kozbe, B. (1998). A personal intelligent mentor for promoting metacognition in solving logic word puzzles. Proceedings of Workshop "Current Trends and Applications of Artificial Intelligence in Education" at The Fourth World Congress on Expert Systems, Mexico City, Mexico.
- Baylor, A. L. (1999a). Cognitive guidelines for agent-based learning environments: The three cognitive considerations of confidence, collaboration, and control. Manuscript in review for *Journal of Interactive Learning Research*.
- Baylor, A. L. (1999b). Intelligent agents as cognitive tools. *Educational Technology*, 39(2), 36-40.
- Baylor, A. L. (1999c). Multiple Intelligent Mentors Instructing Collaboratively (MI MIC): An Intelligent Agent-Based Learning Environment for Learning Instructional Design. Invited full proposal to National Science Foundation in Knowledge Distributed Intelligence: Learning Intelligent Systems division, May, 1999.
- Baylor, A. L., & Jafari, A. (1999). What are the possibilities of intelligent agents for education? Proceedings of International Conference of Technology in Education (ICTE), Tampa, FL.
- Bradshaw, J.M. (1997). An introduction to software agents. In *Software Agents*, edited by J.M. Bradshaw (pp. 3-46). Menlo Park, CA: MIT Press.
- Collins, A. & Brown, J.S. (1987). The computer as a tool for learning through reflection. In H. Mandl and A. Lesgold (Eds.), *Learning issues for intelligent tutoring systems*. New York: Springer-Verlag.
- Dietz, R. B. & Lang, A. (1999). *Affective agents: Effects of agent affect on arousal, attention, liking & learning*. Proceedings of International Cognitive Technology Conference '99, San Francisco, CA.
- Dillenbourg, P., Mendelsohn, P. & Schneider, D. (1994). The distribution of pedagogical roles in a multiagent learning environment. In Lewis, R., & Mendelsohn, P. (Eds.), *Lessons from learning*. Elsevier, 199-216.
- Erickson, T. (1997). Designing agents as if people mattered. In *Software Agents*, edited by J.M. Bradshaw (pp. 79-96). Menlo Park, CA: MIT Press.
- Friedman, B. & Kahn, P. (1992). Human agency and responsible computing: Implications for computer system design. *Journal of Systems Software*, 17, 7-14.
- Janca, P. (1995). Pragmatic application of information agents: theory and practice. *BIS Strategic Decisions*.

- Kearsley, G. (1993). Intelligent agents and instructional systems: Implications of a new paradigm. *Journal of Artificial Intelligence and Education*, 4 (4), 295-304.
- Koda, T., Maes, P. (1996). *Agents with face: the effects of personification of agents*. Proceedings of HCI '96, London, UK.
- Laurel, B (1990). Interface agents: Metaphors with character. In B. Laurel (Ed.), *The art of human-computer interface design*. Reading, MA: Addison-Wesley.
- Laurel, B. (1997). Interface agents: Metaphors with character. In *Software Agents*, edited by J.M. Bradshaw (pp. 67-78). Menlo Park, CA: MIT Press.
- Lepper, M.R. & Chabay, R.W. (1987). Socializing the intelligent tutor: bringing empathy to computer tutors. In H. Mandl and A. Lesgold (Eds.), *Learning issues for intelligent tutoring systems*. (pp. 242-257) New York: Springer-Verlag.
- Lester, J. C., Converse, S. A., Kahler, S. E., Barlow, S. T., Stone, B. A. & Bhoga, R. S. (1997). *The persona effect: affective impact of animated pedagogical agents*. Proceedings of CHI'97 Human Factors in Computing Systems.
- Lester, J. C., & Stone, B. A. (1997) Increasing believability in animated pedagogical agents. In *Proceedings of the First International Conference on Autonomous Agents*.
- Maes, P. (1994), Agents that reduce work and information overload. *Communications of the ACM*, 37(7), 31-40.
- Maes, P. (1997). Agents that reduce work and information overload. In *Software Agents*, edited by J.M. Bradshaw (pp. 145-164). Menlo Park, CA: MIT Press.
- McArthur, D., Lewis, M., & Bishay, M. (1993, November). The roles of artificial intelligence in education: Current progress and future prospects. RAND, Santa Monica, CA, DRU-472-NFS.
- Malone, T.W., Lai, K-Y, & Grant, K.R. (1997). Agents for information sharing and coordination: A history and some reflections. In *Software Agents*, edited by J.M. Bradshaw (pp. 109-144). Menlo Park, CA: MIT Press.
- Negroponte, N. (1970). *The Architecture Machine: Towards a More Human Environment*. MIT Press.
- Negroponte, N. (1997). Agents: From direct manipulation to delegation. In *Software Agents*, edited by J.M. Bradshaw (pp. 57-66). Menlo Park, CA: MIT Press.
- Norman, D.A. (1997). How might people interact with agents. In *Software Agents*, edited by J.M. Bradshaw (pp. 49-56). Menlo Park, CA: MIT Press.
- Palinscar, A., & Brown, A. (1984). Reciprocal teaching of comprehension-fostering and comprehension-monitoring activities. *Cognition and Instruction*, 1(2), 117-175.
- Reeves, B., & Nass, C. (1996). *The Media Equation*. Stanford, CA: CSLI Publications.
- Roesler, M. & Hawkins. D. (1994). Intelligent Agents: Software servants for an electronic information world (and more!). *Online*, 19-32.
- Salomon, G. (1993). On the nature of pedagogic computer tools: The case of the writing partner. *Computers as cognitive tools*. (pp. 179-196). Hillsdale: Lawrence Earlbaum Associates.

Salomon, G., Perkins, D.N., & Globerson, T. (1991). Partners in cognition: extending human intelligence with intelligent technologies. *Educational Research*, 20 (3), 2-9.

Seiker, T. (1994). Coach: A teaching agent that learns. *Communications of the ACM*, 37(7), 92-99.

Shneiderman, B. (1992). *Designing the User Interface, 2nd Edition*. Reading, MA: Addison-Wesley.

Shneiderman, B. (1997). Direct manipulation versus agents: Paths to predictable, controllable, and comprehensible interfaces. In *Software Agents*, edited by J.M. Bradshaw (pp. 97-108). Menlo Park, CA: MIT Press.