Permutations of control: Cognitive considerations for agent-based learning environments

Baylor, A. L. (2001). Permutations of control: Cognitive considerations for agent-based learning environments. *Journal of Interactive Learning Research. 12*(4), 403-425.

Amy L. Baylor, Ph.D.
Instructional Systems Program
Department of Educational Research
Florida State University
307 Stone Building.
Tallahassee, FL 32306
(850) 644-5203
(850) 644-8776 fax
baylor@coe.fsu.edu

Abstract

While there has been a significant amount of research on technical issues regarding the development of agent-based learning environments (e.g., see the special issue of Journal of Interactive Learning Research, (1999, v10(3/4)), there is less information regarding cognitive foundations for these environments. The management of control is a prime issue with agent-based computer environments given the relative independence and autonomy of the agent from other system components. This paper presents four dimensions of control that should be considered in designing agent-based learning environments, with the MIMIC (Multiple Intelligent Mentors Instructing Collaboratively) system as an example. The first dimension of control involves instantiating the instructional purpose of the environment on a constructivist (high learner control) to instructivist (high program/agent control) continuum. The second dimension entails managing feedback, and several issues need to be considered: type, timing, amount, explicitness, and learner control of agent feedback. Third, agent vs learner control is further defined through the desired relationship of the learner to agent(s) (e.g., agent as learning companion, agent as mentor, multiple pedagogical agents, agent as personal assistant, or agent as resource). Fourth, to be instructionally effective, the agent(s) must assert enough control so that the learner develops confidence in the agent(s) in terms of believability, competence, and trust. Overall, an array of possible permutations of system versus learner control must be carefully considered.

Introduction

The distinguishing characteristics of intelligent agents may be in their capacity to independently manage cooperation among distributed programs and/or other agents, to provide intelligent assistance to learners when traditional interfaces are insufficient, and to enable more humanlike interaction (Bradshaw, 1997); however, from an educational vantage point, a better description might be that intelligent agents are computer programs that simulate a human relationship by doing something that another person could otherwise do for you (Seiker, 1994). As described in a recent special issue of the Journal of Interactive Learning Research, (1999, v10(3/4)), intelligent agent technology provides a range of possibilities for computer-based learning environments. However, as the special issue illustrates, research in intelligent agent-based learning environments tends to focus on system development principles rather than examining the cognitive foundations of such environments. Specifically, articles were included in three areas as described by the editors (Arroyo & Kommers, 1999, p. 235): 1) historical development of intelligent tutoring and support systems; 2) agent paradigms and agent-based user support systems; and, 3) tendencies in agent development and application, including agents as guides, information assistants, architectural solutions, help systems, and as simulation agents in virtual and interactive learning environments. To add to the mix of agent research, there are a variety of programs that are called " agents," as listed by Bradshaw (1997):

1) those that can be scheduled in advance to perform tasks on a remote machine
2) accomplishing low-level computing tasks while being instructed in a higher-level of programming language or script
3) abstracting out or encapsulating the details of differences between information sources or computing services
4) implementing a primitive or aggregate " cognitive function"
5) manifesting characteristics of distributed intelligence
6) serving as a mediating role among people and programs
7) performing the role of an " intelligent assistant"
8) migrating in a self-directed way from computer to computer
9) presenting themselves to users as believable characters
10) speaking an agent communication language
11) they are viewed by users as manifesting intentionality and other aspects of " mental state"

For the purpose of this review, the focus is on agent features #4, 6, 7, 9, 11 above, as they all represent agent functions that have specific relevancy to education (Baylor, 1999c). As Bradshaw (1997) notes, one of the most striking things about recent research and development in software agents is how little commonality there is between different approaches. Further, the investigation of educationally-related aspects of agent-based learning environments is still in early stages of conceptualization. As stated by Espinosa & Ramos (1999), much work remains to be done to integrate cognitive theory with agent-based system development. While a given instructional medium (such as agent-based systems) is not more or less effective for learning (e.g., Clark & Sugrue, 1991), we do not yet know what features of agent-based learning environments best facilitate learning. As McArthur, Lewis and Bishay (1993) stated several years ago, the pedagogical component of intelligent systems receives relatively little mention with current systems demonstrating little pedagogical expertise.

Part of the value of intelligent agents as a computing paradigm is that they act independently from each other and the systems in which they operate; this potential for autonomous agent control and reactivity is also a possible liability if not properly planned. Consequently, a prime cognitive consideration is the management of control within an agent-based learning environment. In other words, who has the figurative " ball," the agent(s) or the learner? How much, to what extent, and in what capacities? These are important issues to consider because they impact learning and instruction.

The agent metaphor as described by Erickson (1997) brings with it a new conceptual model that is quite different from that which underlies today's graphic user interfaces (GUI). The object-action conceptual model, which is the underlying conceptual model of GUIs, has to do with objects and actions. Specifically, GUI elements are portrayed as objects on which particular actions may be done. The learners know that objects are visible, passive, have locations, and may contain things. Thus, GUIs are easy to use because these expectations are usually met by all components of the interface; e.g., learner knows s/he can move it, open it, and close it, click and drag it, double click to open it. The agent conceptual model, on the other hand, presents animated agents that can respond to events in contrast to the passive objects of the traditional GUIs. Erickson refers to this as the responsive agent conceptual model which carries the following assumptions from the learner: agents can notice things, agents can carry out actions, agents can know things, and, agents can go places. As he points out, the more intelligence or knowledge of the agent is not necessarily better; what is important is the match between the agent's abilities and the learner's expectations (e.g., dog agent to fetch electronic newspaper). Unlike objects, agents can go places or be "off stage" and be summoned by the learner when needed. While objects are safe and predictable and figuratively just sit there and hold things, agents become the repositories for adaptive functionality and they can notice things, use rules to interpret them, and take actions based on their interpretations. Erickson predicts that eventually learners will have awareness of agents like they do of objects and there will be consistent ways of finding out what an agent will notice, what actions it will carry out, what it knows, and where it is. However, until they do have this awareness, he implies that they will lose a sense of control.

Along this line, the balancing of agent versus learner control is an important issue, and different permutations of control are possible depending on the instructional purpose of the system. Four key dimensions of control will be discussed here: 1) the instructional purpose of the system; 2) agent feedback; 3) the relationship of agent(s) and learner; and 4) learner confidence in the agent(s). These cognitive considerations will be discussed in light of the prototype MIMIC (Multiple Intelligent Mentors Instructing Collaboratively) agent-based learning environment (Baylor, 1999a; 1999d; 2000b; 2001; in press). While issues regarding agent architecture are important given that they set the constraints for system development, the scope of this article is limited to the cognitive requirements.

Defining the instructional purpose: The constructivist -- instructivist continuum

The locus of control between a computer-based instructional system (e.g., agent-based learning environment) and the learner is a critical issue. According to Reiber (1994), there are two extremes in intelligent learning systems: 1) the instructionist approach (here referred to as instructivist)— in CAI (Computer Aided Instruction) and curriculum-based intelligent tutoring systems (ITS), and 2) the constructivist approach with microworlds, open learning environments and educational hypermedia systems. Or as Lajoie & Derry (1993) proposed, there are generally two camps in considering the use of intelligent systems for education: 1) those who promote using the system as a cognitive tool to stimulate the student to monitor and diagnose performance (e.g., constructivist approach); and, 2) those who promote using the model building approach to use the system as an intelligent tutor (e.g., instructivist approach). Even though Lajoie (2000) more recently denounces the literal separation of these two "camps," it serves as a useful way to conceptualize the difference in approach.

In traditional ITS, the locus of control tends to focus on the computer program since it selects problems for the student and intervenes with instructional interactions when certain events occur. While it can involve coaches and system-supported negotiation of teaching goals, and collaboration among groups of students, this is not traditionally the case. Rather, an ITS generally provides an instructivist approach toward learning where it systematically directs the learning process.

In contrast, one type of constructivist implementation of computer-based systems for learning is as a cognitive tool. Cognitive tools (e.g. Lajoie & Derry, 1993) are mental and computational devices that support, guide, and extend the thinking processes of their students. In terms of educational psychological theory, these cognitive tools serve as a form of distributed cognitions (Salomon, 1993b) to extend a person's intellectual capacity. By extending the cognitive capabilities of the person, cognitive tools could serve to thus augment a person's zone of proximal development, which is the limit of his/her ability to imitate processes demonstrated by others (Vygotsky, 1962). Developing intelligent agents to serve as cognitive tools proves to be a beneficial learning approach from a constructivist perspective (Baylor, 1999b).

A critical issue from a constructivist approach to agent-based learning environments is in moderating between the agent taking over thinking for the student with the agent training the student to think more effectively (Baylor, 2000a). Salomon (1993a) refers to this as the difference between the effects "of" and "with" technology, with effects "of" technology being more desirable. In this sense, the computer technology in and of itself is of little interest whereas what activities it affords is of interest. Further, in terms of the amount of artificial intelligence that should be used by an agent, it is important to consider that more intelligence is not necessarily better from a pedagogical perspective (e.g., Salomon, Perkins & Globerson, 1991). Intelligent agents for learning can come in both varieties: constructivist and instructivist. Overall, in the instructivist approach, the role of the agent is to teach the student knowledge similar to the role taken traditionally by human teachers. In the constructivist approach the agent would be a medium that does not teach the student directly.

For either instructional approach (constructivist or instructivist), a key issue is regarding the nature and need for instructional planning. Instructional planning is the process of mapping out a global sequence of instructional goals and actions that provides consistency, coherence and continuity in the instructional process. It can be applied at two levels: 1) planning of content - the process of selecting the content for an instructional goal that places the student on an appropriate learning path; and 2) planning of delivery - the process of optimal selecting and sequencing of the tutorial interactions focused on a given content (Wasson, 1996).

Wasson (1996) argues that instructional planning is applicable to more than just strict tutoring style ITS. Planning approaches have been used to implement various instructional and teaching strategies such as tutoring, coaching, cognitive apprenticeship, Socratic dialogue, and combinations of these. In fact, any environment to support learning can be enhanced through the use of planning techniques since they can be used as a basis for implementing certain teaching strategies. Further, she argues that planning is only a technique and is neutral to the particular educational philosophy underlying the system. However, it can be used to implement a given educational philosophy. Further, as Pontecorvo (1993) explains, there are some constructivists (e.g., Bruner (1966), Simons (1993)) who would accept that instruction is a systematic and planned activity that is aimed at learning and includes teaching, and, importantly, that there is no opposition between the view of constructivist learner and the view of a planning teacher.

As Chan (1995) notes, one of the key differences between the two approaches is in terms of whether the computer system is assumed to know all the answers and be able to lead the human student to the right path in problem-solving as opposed to not always knowing the correct answers. In other words, the agent acts as a teacher or as a student. This leads to interesting alternatives in the learning environment, where the learner can work with both mentoring agents and learning companion agents.

It is possible to find a compromise between the instructivist and constructivist approaches where the agents strategically choose pedagogical interventions of an instructivist nature while also allowing the learner to assert his/her control to construct knowledge in the environment. Specifically, the agents can implement instructional planning while still promoting a constructivist approach; however, the system must relinquish some control over the learning process to the learner. Along this line, Espinosa & Ramos (1999) provide evidence to support constructivist pedagogical strategies when applied with systematic instructional planning.

Sharing of control over learning process by the learner and agent(s): Implementation

The pedagogical approach of the system — on the continuum of constructivist to instructivist — has a direct effect upon the locus of control. The learner and the agent(s) must coordinate the sharing of control over the learning process. Specifically, a radical constructivist approach would imply that the learner has absolute control over his/her learning whereas a strict instructivist approach would imply that the agents have absolute control over the learning process. To provide an example of what these permutations may look like, the next subsection will illustrate examples of high program control to examples of increasingly greater learner control, thereby beginning with more instructivist approaches and moving into more constructivist approaches.

Minimal control for the learner: Machine learning

An example of a system with minimal control for the learner is Maes' (1997) approach using machine learning techniques, inspired by the metaphor of a personal assistant. In this approach, the agent can "program itself" to assist the learner given a minimum of background knowledge. The agent can learn appropriate behavior from the learner and other agents. Maes suggests that two conditions must be fulfilled for the machine learning approach to be appropriate. First, the use of the application has to involve a substantial amount of repetitive behavior (within the actions of one learner or among learners) as there must be regularities for the agent to learn. Second, this repetitive behavior is potentially different for different learners. If learners' behavior is the same, a knowledge-based approach may yield faster results (for personal assistant-type agents). Overall, it seems that the machine learning approach may be advantageous for a more simple agent-based learning environment where it is appropriate for the agent(s) to monitor the learner's behavior with the goal of asserting ultimate control and guidance.

More balanced control for learner and agent(s): Semiformal systems

With the semiformal systems approach presented by Malone, Lai & Grant (1997), they propose sharing the control between agent(s) and learner. As they state:

> Don't build computational agents that try to solve complex problems all by themselves. Instead, build systems where the boundary between what the agents do and what the humans do is a flexible one. We call this the principle of semiformal systems because it involves blurring the boundary between formally represented information acted upon by agents and informally represented information acted upon by humans (p. 110).

There are three properties that characterize this approach. First, it represents and automatically processes certain information in formally specified ways. Second, it represents and makes it easy for humans to process the same or other information in ways that are not formally specified. Third, it allows the boundary between formal processing by computers and informal processing by people to be easily changed.

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, however, 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 (p. 118).

A serious risk for learning agents according to Malone and colleagues is that agents will infer incorrect rules (or fail to infer correct ones) when learners could have easily described the rules they actually wanted to use. A semi-formal approach to designing such systems would suggest that any attempts to have agents automatically "learn" from observing learners' behavior should occur only after the system already provides a way for learners to directly specify what they want (p. 119). While this does not imply that learners be exposed to low-level programming language details, it does suggest that through the user interface the learner is aware of what is going on in a system in a form that is meaningful and understandable. One approach to part of this problem is suggested by the "explanation" facilities in traditional knowledge-based systems.

However, Maes (1997) notes that a problem with this approach is in terms of the competence dimension as it requires too much insight, understanding, and effort from the learner. Consider that the learner must do the following: 1) recognize the opportunity for employing an agent; 2) take the initiative to create an agent; 3) endow the agent with explicit knowledge (specifying this knowledge in an abstract language); 3) item maintain the agent's rules over time (as work habits or interests change, etc.).

Overall Malone, et. al. (1997)'s semiformal systems approach seems akin to the use of agents as cognitive tools to extend the learner's capabilities. For example, Kurhila and Sutinen's (1999) system is ill-structured and high in user control with the responsibility of pedagogy belonging to human experts. In their system for special education, the learning goal is not set at the beginning but is adjusted during the process, based on learner's choices so that the user can construct his/her own model of the topic. Another open learning environment with agents is SMILE (Stoyanov & Kommers, 1999). In SMILE (Solution, Mapping, Intelligent Learning Environment), the agent supports the user as a cognitive tool to use concept mapping to solve ill-structured problems. By making the agent(s) reasoning visible via figurative "glass boxes," there is great potential for the agent(s) to model thought processes to the learner and facilitate learner reflection and self-evaluation.

Ultimate control for the learner

In Stagecast Creator (formerly called KidSim) (Smith, Cypher & Spohrer, 1997), the learner has absolute control over the learning process since part of his/her task is to program agents in the context of simulated microworlds. Stagecast Creator is a tool kit for children and non-programming adults to construct and modify simulations by programming simple agent behavior. Of special interest is the application of a GUI to programming the agents. Smith and colleagues apply good user interface principles to the process of programming by including the combination of two techniques: 1) graphical rewrite rules; 2) programming by demonstration. The agents in the Stagecast Creator environment are full objects in object-oriented-programming sense as they have state (properties), behavior (rules), and an appearance. They are similar to those in Logo Microworlds, but with a key difference in how they are programmed: In Logo, kids program objects with LOGO whereas in Stagecast Creator kids construct "graphical rewrite rules." Consequently, through the act of programming the agents, the children learn indirectly about the nature of the simulations. Such an environment encourages "learning by designing" as an effective constructivist instructional strategy where students design instruction to better learn information. Another example of a system with very high learner control is where students teach the agent to teach their peers in the content domains of math and science (Brophy, Biswas, Katzlberger, Bransford, & Schwartz, 1999).

In summary, the choice of the levels of program and learner control is dependent on the instructional pedagogy. A machine learning approach may be beneficial for the agents to monitor the learner and then proceed to instruct him/her regarding his/her performance. The semi-formal systems approach (e.g., Malone et. al., 1997) is successful as a balanced division of control by making the agents' reasoning visible to the learner. Finally, a system high in learner control, such as Stagecast Creator, is beneficial for creating a more exploratory constructivist learning environment.

Managing agent-learner feedback

In terms of feedback, for each instantiation of the system -- from high program control to high learner control — there are different possibilities for implementing feedback. Arroyo & Kommers (1999) categorize feedback in terms of three micro-functions (presentation, testing, and corrective), and the focus here is on corrective feedback. In the case of machine learning the agent is essentially learning patterns of behavior of the learner and taking a lead role to facilitate his/her educational goals; thus, feedback would be controlled completely by the agent. In the case of semi-formal systems, the feedback is automated to some extent by the system, but with opportunities for the learner to specify his/her desires. In the case of an open-ended environment such as Stagecast Creator, feedback from agents is completely controlled by the learner at the risk of creating an environment that is too unstructured.

Cognitive windows: glass boxes, not black boxes

An important consideration in terms of feedback is that the pedagogical agent should not provide too many insights and thereby annoy the student. 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. In this way the learner-agent relationship becomes mutually collaborative as each provide feedback for each other.

A related issue is in terms of how active the agents should be in providing explanations of their pedagogical behavior. Assuming that the agents do have some planning role in the instructional environment, does the learner needs understanding of what happened pedagogically and why? One advantage of having the mentoring agents explicitly reveal teaching strategy differences (as opposed to them being built-in to the system and invisible to the learner) is that it can facilitate reflective and metacognitive thinking for the learner (e.g., Baylor & Kozbe, 1998). Information processing models of cognition (e.g., Pressley & McCormick, 1995) suggest the primary importance of metacognitive skills, particularly as metacognitive ability is a feature of expert problem solvers (Glaser & Chi, 1988). Where novices tend to focus on surface features of a problem, experts tend to better organize and represent the information and use more metacognitive skills. Although a metacognitive state may hinder intuitive ability (Baylor, 1997), it is beneficial for other reasoning processes.

There are potential problems regardless of the level of feedback explicitness. As Erickson (1997) proposes:

> Consider an intelligent tutoring system that is teaching introductory physics to a teenager. Suppose the system notices that the student learns best when information is presented as diagrams and adapts its presentation appropriately. But even as the system is watching for events, interpreting them, and adjusting its actions, so is the student watching the system, and trying to interpret what the system is doing. Suppose that after a while the student notices that the presentation consists of diagrams rather than equations: it is likely that the student will wonder why: ?Does the system think I'm stupid? If I start to do better, will it present me with equations again?' There is no guarantee that the students' interpretations will correspond with the system's. How can such potentially negative misunderstandings on the user's part be minimized? (p. 83)

The trick seems to be in resolving the learner's need for explanations from the system with the need for her to formulate her own explanations. And the solution to this dilemma can be simple, depending on the context. As noted previously, an act such as winking can connote a lot of information without providing explicit details. Further, in terms of timing of feedback, Espinosa & Ramos (1999) point out the problem of "temporal holes" in traditional, instructivist-designed systems. As they describe, important learning events may be taking place but the system misses them due to its predetermined time intervals for response.

To facilitate reflection in MIMIC, the concept of cognitive windows is taken one step further. Here, teaching strategy differences among the mentoring agents are explicit to the learner (as opposed to being built-in to the system and invisible to the learner) in order to facilitate reflective thinking. For example, at the completion of a case study, the learner is provided with a demonstration of how each agent "solved" the case and then encouraged to reflect as to how it compared to his/her approach. Along this line, a key goal of MIMIC is to promote metacognition in learning instructional design. Feedback in MIMIC is both system-controlled and learner-initiated. While there is some feedback that is unsolicited, the learner can always request additional suggestions from the agent(s).

Overall, taking the middle road, it seems that glass boxes are preferable to black boxes and there is educational value for agents to model thought processes. Regarding the amount and timing of feedback, an overall guideline is that agents should not provide feedback that is too explicit. Ideally, the system should provide some means for the learner to control the amount of feedback.

Especially in an environment with high learner control, the learner needs to be familiar with the agent(s) so that s/he can better understand the agents' actions. This required relationship of the learner with the agent(s) differentiates the agent-based approach from the traditional ITS approach. The next section will describe possible combinations of the agent-learner relationship.

Defining the relationship of learner to agent(s)

There are several possibilities in defining the learner-agent relationship. The agent can serve as a personal assistant, a pedagogical expert/mentor, or as a learning companion. Additionally, there can be multiple agents serving roles in these capacities. The inherently social relationship of learner and agent(s) makes sharing the power between/among learner and agent(s) important to define.

Agent as assistant

The role of agent as assistant is reminiscent of Negroponte's (1970) original characterization of intelligent agent as a butler. In this agent-learner relationship, the agent facilitates and automates certain learning-related tasks. For example, In El-khouly, Far, Koono (1999)'s system for teaching computer languages, there is a personal assistant for teachers (PAA-T) that represents the server and a personal assistant for students (PAA-S) that represents the client. They describe a coached problem solving approach where the initiative in student-tutor interaction changes according to student progress, to try to encourage the student along a problem-solving path. In a similar conceptualization, the Digital Teaching Assistant in Angel (A New Global Environment for Learning) (see Jafari, 1999) facilitates course-related logistical tasks that a traditional human teaching assistant would support. Using the agent as a different kind of assistant, Kurhila & Sutinen (1999) help "hyperspace disabled" learners with the "Ahmed" agent, who helps to limit options and makes the learning space more meaningful.

Agent as mentor

Another possible instantiation of the learner-agent relationship is that where the agent serves as a pedagogical expert or mentor (see Baylor, 2000a). Here, the agent can monitor and evaluate the timing and implementation of teaching interventions (e.g., help, feedback). For example, Stone & Lester (1998) describe Herman-the-bug, an animated pedagogical agent that facilitates learning about plant structure. COACH is an intelligent agent system by Seiker (1994) that records learner experience to create personalized learner help for LISP with an adaptive interactive help system. As Seiker describes, (p. 92) "Just as a football coach will stand on the sidelines and encourage, cajole or reprimand, so COACH is an advisory system that does not interfere with the learner's actions but comments opportunistically to help the learner along." Like COACH, an intelligent agent could serve as an advisory-style agent that builds a learner relationship with the explicit goal of educating the individual. Along this line, the intelligent agent will draw from an adaptive user model that selects appropriate advice.

With agent as mentor, the relationship becomes similar to that of a cognitive apprenticeship (e.g., Collins & Brown, 1987), where the student improves his/her performance while working with the more expert performer: the intelligent agent. In such a relationship and learning environment, the agent may facilitate learning by modeling and coaching formative skills. Further, such a relationship would suggest that as the student gains expertise, the agent would fade and allow for more student initiative. Additionally, such a relationship can support an individual's metacognitive processes with the agents serving as a technological "reciprocal teacher" (e.g., Palinscar & Brown, 1984), prompting the individual to engage in analysis of his/her own cognitive processes (see Baylor & Kozbe, 1998). For example, the agents can encourage the individual to assess what cognitive strategies are being used, similar to Salomon's pedagogic computer program, the Writing Partner (Salomon, 1993a), which asks the learner intelligent questions through the writing process.

Having several experts describing the instructional content matter from different points of view can be very rewarding for the learner (Laurel, 1997). An example of this implementation is the MIMIC system, where multiple agents as mentors each represent a theoretical perspective of instructional design (e.g., direct instruction, or constructivism) (Baylor, 2001; in press). Having several experts describing the instructional content matter from different points of view can be very rewarding for the learner (Laurel, Oren & Don , 1990). Since the learner may learn more from a teacher using a different instructional approach or perspective, multiple teachers can help the learner to establish the best personalized approach to understanding the content. Hietala & Niemirepo (1998) refer to this aspect as the need for pedagogical multiplicity of teachers. They suggest that the many levels and complexities of the learning process might be alleviated by providing more alternatives to the learner via an "extended family of intelligent agents." With the possibilities of multiple intelligent agents and/or learners, the problem is deciding which form of collaboration is appropriate for a given learning situation (Kearsley, 1993). Note that while the issue of collaboration versus competition in managing multi-agent systems is important, it is not covered in the scope of this article.

Another multi-agent system called ETOILE, by Dillenbourg, Mendelsohm, and Schneider (1994), is designed to teach educational psychology principles. ETOILE includes five teaching agents, labeled after the teaching styles they implement. They are called by the names of Skinner, Bloom, Vygotsky, Piaget, and Papert. Each tutor is implemented as an independent rule base. The five teaching agents implement decreasing level of directiveness: Skinner works step by stop, Bloom makes larger steps but with close control of mastery, Vygotsky is based on participation, Piaget intervenes only to point out problems and Papert does not interrupt the learner. This ETOILE system also includes a "coach" agent that is in charge of which tutor is used; however, the learner may also select or remove a tutor.

However, it must be noted that there are significant difficulties in developing true pedagogical expertise in an intelligent agent as there are with other artificially-intelligent systems. From a theoretical perspective, providing beneficial features of intelligent agents essentially requires the agents to have metaknowledge, or knowledge about the knowledge they are working with, which is a very difficult requirement. As McArthur, Lewis and Bishay (1993) suggest, most intelligent tutoring systems are constrained to a single method of teaching and learning, while truly expert human tutors can adopt different methods. Yet, the implementation of multiple teaching agents in the learning environment serves to challenge these difficulties.

Agent as learning companions

Another relationship within which the agent can serve to further learning outcomes is to serve as a learning companion agent that learns together with the learner. A learning companion is particularly effective if it replicates the behavior of a novice with whom the learner can "collaborate." A less capable computer-based learning companion may be preferable for the student because it encourages him/her to teach the companion (e.g., Uresti, 1998). Regardless, one or more learning companions provides the learner with the capability of collaborating with others and thus theoretically supported by the idea of distributed cognition. As Dillenbourg (1996) explains:

> Distributed cognition does not simply mean division of labor where you have a task, you split the task, you give different sub-tasks to different people, they go in different rooms, and when they have finished they assemble the results. That's cooperation (1+1=2). The idea of distributed cognition is something more-- it's the idea that the whole is more than the sum of the parts? with collaboration the learner has to do the sub-task which has been allocated to him, but at the same time he has to interact about it (1+1>2). (p.165)

Thus, the addition of a learning companion in the form of an intelligent agent can positively contribute to the learning process as contributing to the personal cognitive formulations of the task at hand. When considering agents as learning companions, Aimeaur, Dufort, Leibu, and Frasson (1997) promote that the learning companion should possess pedagogical knowledge to sometimes purposely disturb the human learner — this "learning by disturbing" strategy aims at making the student to confront his/her weaknesses. Hietala and Niemirepo (1998) propose an environment where there are multiple learning companions of different levels of expertise assisting young children in learning mathematics. They propose that providing a group of heterogeneous companion agents (both strong and weak ones) increases the learner's motivation to collaborate with the agents. Another instantiation of learning companion agents is to have them replicate the behavior of novices, with whom the learner can "collaborate." In Sheremetov and Nunez's (1999) "EVA" system, there are four virtual learning spaces: a) knowledge; b) collaboration; c) consultation; and d) experimentation. In this computer- supported cooperative learning environment the learner receives support from learning community agents and serves as an active collaborator, learning through interaction.

Overall, the addition of agents as assistants, mentors and/or learning companions provides several possibilities. Specifically, pedagogical agents can guide the learner through the learning process, and learning companion agents can work with the learner in collaboration toward learning. For both types of agents, multiple agents provide even more flexibility in providing different modes of teaching assistance and/or levels of learning companion expertise.

Perhaps the most critical consideration is the learner-agent relationship is the need for the learner to be confident in the agent. For learner confidence, the agent must be believable, the learner must be able to trust the agent, the agent must demonstrate that it is competent, and a motivational attachment must be developed between learner and agent. This requires that the agent demonstrate sufficient control over the learning process so as to earn the learner's trust and confidence.

Agent control for learner confidence

In terms of promoting the learner's confidence in the agent(s), the agent(s) must be believable, motivational, competent, and trustworthy, all of which require the agent to assert control in the learning process.

Learner must perceive agent(s) as believable

The Guides (Oren, Salomon, Kreitman & Don, 1990; Salomon, Oren, and Kreitman, 1989) project (as discussed by Erickson, 1997) is an anecdotal study that investigated the issue of believability for agent-like computer programs. The project involved the design of an interface to a CD ROM encyclopedia (focusing on early American history) with a set of travel guides, each of which was biased towards a particular type of information (settler woman, Indian, inventor). They found that students tended to assume that the guides, which were presented as stock characters, embodied particular characters. For example, since many of the articles in the encyclopedia were biographies, learners would assume that the first biography suggested by a guide was its own! Students also wondered if they were seeing the article from the guide's point of view (they weren't). Further, they sometimes assumed that guides had specific reasons for suggesting each story and wanted to know what they were (in line with learners' general wish to understand what adaptive functionality is actually doing). Some of the students got emotionally engaged with the guides; one student getting angry that the guide had betrayed her; in another case the guide inadvertently disappeared and the student interpreted this as " ...the guide got mad, he disappeared." As Erickson (1997) comments, while no controlled experiment was involved in these findings, rather these findings are anecdotal, it is hard to believe that the learner would have made such an inference if the suggested articles had been presented in a floating window that had vanished.

In another project, Nass & Steuer (1993) (as discussed by Erickson, 1997) researched the role of computers as social actors, investigating the tendency of people to use their knowledge of people and social rules to make judgments about computers. They found that very small cues can trigger people's readiness to apply social rules to computers. For example, simply having the computer use a human voice is enough to cause people to apply social rules to the computer. This suggests that the agent metaphor may be invoked very easily, which is surprising given that the cues are so minimal (e.g., voice). In an experimental manipulation, they found that ratings were much more favorable when computer B praised A's tutorial than when computer A had praised itself (which is a social rule with humans).

Both projects suggest that believability of the agent(s) contributes to the confidence that the learner has with the agent(s). Specifically, the learner may tend to personify instructions coming from an agent and this learner-agent relationship can be initiated even with minimal cues such as voice.

Learner must be motivated by agent(s)

A critical factor regarding confidence in 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 feedback. An important part of this feature is that the intelligent agent will resemble a human mentor in terms of motivational qualities (e.g, Baylor, 2000a). 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. Additionally, it is beneficial if the agents are conversational and invoke a human-like persona in their interactions with the learner.

Small cues can lead to motivational attachment. Walker, Sproull, & Subramani (1994) (as discussed by Erickson, 1997) performed a controlled study of human response to a synthesized talking face. Specifically, compared to the people who simply filled out a questionnaire, those who answered the questions delivered by the two synthesized faces spent more time, wrote more comments, and made fewer errors. People who interacted with the faces seemed more engaged by the experience. In terms of the faces, one face was stern; people who answered questions delivered by the stern face spent more time, wrote more comments, and made fewer errors. Yet they also liked the experience and the face less.

When the social relationship of the learner with the agent is motivation, this serves to further engage the learner into the learning experience and thus increase his/her confidence. Also, such a motivational attachment may provide the learner with greater willingness to work together with the agent which would further enhance confidence.

Learner must perceive agent(s) as competent

A key issue regarding the learner's level of confidence with the agent is regarding his/her perceived competence of the agent(s) (e.g., Maes, 1997). Further, as Norman (1997) points out, people have overblown expectations about what an agent can or should do. Overall, the learner must believe that the agent(s) is sufficiently intelligent and efficient in its suggestions.

Learner must trust the agent(s)

The learner must also perceive that the agent(s) are trustworthy. One way that trust is enhanced is when the agent can provide reassurance to the learner. Specifically, Norman (1997) brings up an important issue: assuming that the technology is under control from a technical standpoint, the agent needs to reassure the learner that all is working according to plan. Trust also requires that the learner have confidence in the privacy of his/her interactions with the agent(s) and the confidentiality of the actions.

In summary, the learner's development of confidence in the agent(s) is dependent on his/her perceptions in the following areas: 1) agent believability; 2) motivational attachment with the agent(s); 3) agent competence; and, 4) agent trustworthiness. So again, the issue of assertion of control is important, here from the perspective of the agent so as to establish a beneficial working relationship with the learner.

Conclusion

To summarize, the management of control is a prime issue with the design of agent-based computer environments. There are four dimensions of control to consider, which lead to a variety of permutations. First, determining the instructional purpose of the environment on a constructivist (high learner control) to instructivist (high program/agent control) continuum is necessary. Second, several issues need to be considered to manage feedback: type, timing, amount, explicitness, and learner control of agent feedback. Third, agent versus learner control is further defined through the desired relationship of the learner to agent(s) (e.g., agent as learning companion, agent as mentor, multiple pedagogical agents, agent as personal assistant, or agent as resource). Finally, the agent(s) must assert enough control so that the learner develops confidence in the agent(s) in terms of believability, competence, and trust. Overall, the design of agent-based learning environments requires careful consideration of these dimensions in order to develop systems with appropriate permutations of system and learner control.

References

Aimeaur, E., Dufort, H., Leibu, D., and Frasson, C. (1997). Some justifications for the learning by disturbing paradigm. In B. DuBoulay & R. Mizoguchi (Eds.), Artificial Intelligence in Education: Knowledge and Media in Learning Systems. (pp. 119-126). Proceedings of AI-ED. Kobe, Japan.

Arroyo L. & Kommers, P. Preface? Intelligent agents for educational computer-aided systems. Journal of Interactive Learning Research, 10 (3/4).

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 (1997). A three-component conception of intuition: Immediacy, sensing relationships, and reason. New Ideas in Psychology, 15 (2), 185-194.

Baylor (1999a). Foundations for designing MIMIC, an intelligent agent-based learning environment. Proceedings of International Conference of Technology in Education (ICTE), Tampa, FL.

Baylor (1999b). Intelligent agents as cognitive tools. Educational Technology, 39(2), 36-40.

Baylor (1999c). Intelligent agents for education. Paper presented at Annual Convention of American Educational Research Association, Montreal, Canada.

Baylor (1999d). Multiple Intelligent Mentors Instructing Collaboratively (MIMIC): Developing a Theoretical Framework. Proceedings of International Cognitive Technology Conference, San Francisco, CA.

Baylor (2000a). Beyond butlers: Intelligent agents as mentors. Journal of Educational Computing Research, 22(4), 373-382.

Baylor, A. L. (2000b). A framework for MIMIC, an intelligent agent-based learning environment. Paper presented at the Association for Educational Communication and Technology, Long Beach.

Baylor, A. L. (2001). Investigating Multiple Pedagogical Perspectives through MIMIC (Multiple Intelligent Mentors Instructing Collaboratively). Proceedings of Artificial Intelligence in Education (AI-ED) International Conference, San Antonio, Texas.

Baylor, A. L. (in press). Agent-based learning environments for investigating teaching and learning. Journal of Educational Computing Research.

Bradshaw, J. M. (1997). An introduction to software agents. In J. M. Bradshaw (Ed.), Software Agents (pp. 3-46). Menlo Park, CA: MIT Press.

Brophy, S., Biswas, G., Katzlberger, T., Bransford, J., & Schwartz, D. (1999). Teachable Agents: Combining Insights from Learning Theory and Computer Science. Paper presented at the Artificial Intelligence in Education (AI-ED) International Conference, LeMans, France.

Bruner, J. (1966). Toward a Theory of Instruction. Cambridge, Mass: Harvard University Press.

Chan, T-W. (1995). Social learning systems: an overview. In B. Collis & G. Davies (Eds.), Innovative adult learning with innovative technologies. (pp. 101-122). IFIP Transactions A-61, North-Holland.

Clark, R. & Sugrue, B. (1991). Research on Instructional Media, 1978-1988. In G. Anglin (Ed.) Instructional Technology: Past, present, and future. Englewood, Colorado: Libraries Unlimited.

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.

Dillenbourg, P. (1996). Some technical implications of distributed cognition on the design of interactive learning environments. International Journal of Artificial Intelligence in Education, 7 (2), 161-179.

Dillenbourg, P., Mendelsohn, P., & Schneider, D. (1994). The distribution of pedagogical roles in a multiagent learning environment. In R. Lewis & P. Mendelsohn (Eds.), Lessons from learning (pp. 199-216): Elsevier.

El-khouly, M. Far, B. & Koono, Z. (1999). Agent-based computer tutorial system: An experiment for teaching computer languages. Journal of Interactive Learning Research, 10 (3/4), 275-286.

Erickson, T. (1997). Designing agents as if people mattered. In J. M. Bradshaw (Ed.), Software Agents (pp. 79-96). Menlo Park, CA: MIT Press.

Espinosa, E. & Ramos, F. (1999). Intelligent agency and tutoring: The importance of being timely. Journal of Interactive Learning Research, 10 (3/4), 335-360.

Glaser, R. & Chi, M. (1988). Introduction: What is it to be an expert? In M. Chi, R. Glaser & M. Farr (Eds.), The nature of expertise (pp. xv-xxiix). Hillsdale, NJ: Erlbaum & Associates.

Hietala, P., & Niemirepo, T. (1998). The competence of learning companion agents. International Journal of Artificial Intelligence in Education, 9, 178-192.

Jafari, A. (1999). Intelligent Agents for Distance Learning. CyberLab, IUPUI School of Engineering and Technology, Indiana University Purdue University Indianapolis.

Kearsley, G. (1993). Intelligent agents and instructional systems: Implications of a new paradigm. Journal of Artificial Intelligence and Education, 4 (4), 295-304.

Kurhila, J. & Sutinen, E. (1999). Sharing an open learning space by individualizing agents. Journal of Interactive Learning Research, 10 (3/4) 287-300.

Lajoie, S. P. (2000). Computers as cognitive tools: No more walls. Mahwah: Lawrence Earlbaum Associates.

Lajoie, S. P. & Derry S. (1993). Computers as cognitive tools. Hillsdale: Lawrence Earlbaum Associates.

Laurel, B. (1997). Interface agents: Metaphors with character. In J. M. Bradshaw (Ed.), Software Agents (pp. 67-78). Menlo Park, CA: MIT Press.

Laurel, B., Oren, T. & Don, A. (1990). Issues in multimedia interface design: Media integration and interface agents. Proceedings of CHI'90 Human Factors in Computing Systems, ACM Press, 133-139.

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.

Maes, P. (1997). Agents that reduce work and information overload. In J. M. Bradshaw (Ed.), Software Agents (pp. 145-164). Menlo Park, CA: MIT Press.

Malone, T.W., Lai, K-Y, & Grant, K.R. (1997). Agents for information sharing and coordination: A history and some reflections. In J.M. Bradshaw (Ed.), Software Agents (pp. 109-144). Menlo Park, CA: MIT Press.

McArthur, D., Lewis, M., & Bishay, M. (1993). The roles of artificial intelligence in education: Current progress and future prospects. RAND, Santa Monica, CA, DRU-472-NFS.

Nass, C., and Steuer, J. (1993). Anthropomorphism, agency, and thopoeia: Computers as social actors. Human Communication Research, 19 (4), 504-527.

Negroponte, N. (1970). The Architecture Machine: Towards a More Human Environment. MIT Press.

Negroponte, N. (1997). Agents: From direct manipulation to delegation. In J. M. Bradshaw (Ed.), Software Agents (pp. 57-66). Menlo Park, CA: MIT Press.

Norman, D. A. (1997). How might people interact with agents. In J. M. Bradshaw (Ed.), Software Agents (pp. 49-56). Menlo Park, CA: MIT Press.

Oren, T., Salomon, G., Kreitman, K., & Don, A. (1990). Guides: Characterizing the Interface. In B. Laurel (Ed.), The art of human-computer interface design (pp. 367-381). Reading, Mass.: Addison-Wesley.

Palinscar, A., & Brown, A. (1984). Reciprocal teaching of comprehension-fostering and comprehension-monitoring activities. Cognition and Instruction, 1(2), 117-175.

Pontecorvo, C. (1993). Developing literacy skills through cooperative computer use: Issues for learning and instruction. In T.M. Duffy, J. Lowyck, D.H. Jonassen (Eds.) Designing Environments for Constructive Learning, (pp. 139-160). NATO ASI Series F. Heidelberg: Springer-Verlag.

Pressley, M., & McCormick, C. (1995). Advanced educational psychology for educators, researchers, and policymakers. New York, NY: HarperCollins College Publishers.

Reiber, L. P. (1994). Computers, Graphics, & Learning. Duvuque, IA: Wm. C. Brown Communications, Inc., Brown & Benchmark.

Salomon, G. (1993a). 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. (1993b). Distributed cognitions: Psychological and educational considerations. Cambridge: Cambridge University Press.

Salomon, G., Oren, T., & Kreitman, L. (1989). Using guides to explore multimedia databases. In Proceedings of the Twenty-Second Annual Hawaii International Conference on System Science, 3-11. Washington, D.C.: IEEE Computer Society.

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.

Sheremetov, L. & Nunez, G. (1999). Multi-agent framework for virtual learning spaces. Journal of Interactive Learning Research, 10 (3/4), 301-320.

Simons, P.R.J. (1993). Constructive learning: The role of the learner. In T.M. Duffy, J. Lowyck, D.H. Jonassen (Eds.) Designing Environments for Constructive Learning, 291-313, NATO ASI Series F. Heidelberg: Springer-Verlag.

Smith, D.C., Cypher, A., & Spohrer, J. (1997). KidSim: Programming agents without a programming language. In J. M. Bradshaw (Ed.), Software Agents (pp. 165-190). Menlo Park, CA: MIT Press.

Stone, B., & Lester, J. (1998). Dynamically sequencing an animated pedagogical agent. In M. N. Huhns & M. P. Singh (Eds.), Readings in Agents (pp. 156-163). San Francisco, CA: Morgan Kaufmann.

Stoyanov & Kommers, (1999). Agent-support of problem solving through concept mapping. Journal of Interactive Learning Research, 10 (3/4), 401-425.

Uresti, J.A.R., (1998). Teaching a learning companion. Proceedings of Workshop "Current Trends and Applications of Artificial Intelligence in Education" at The Fourth World Congress on Expert Systems, Mexico City, Mexico.

Vygotsky, L. (1962). Thought and language. Cambridge: Massachusetts Institute of Technology. (Original work published 1934).

Walker, J., Sproull, L. & Subramani, R. (1994). Computers are social actors. Proceedings of Human Factors in Computing Systems: CHI ?94 Conference, 85-91. New York: ACM Press.

Wasson, B. (1996). Instructional Planning and Contemporary Theories of Learning: Is this a Self-Contradiction? In P. Brna, A. Paiva & J. Self (Eds.) Proceedings of the European Conference on Artificial Intelligence in Education, 23-30, Lisbon: Colibri.