

Mixed Initiative Problem Solving with Animated Pedagogical Agents*

James C. Lester Charles B. Callaway Brian A. Stone Stuart G. Towns

Multimedia Laboratory
Department of Computer Science
North Carolina State University
Raleigh, NC 27695-8206

{lester,cbcallaw,bastone,sgtowns}@eos.ncsu.edu

Abstract

Knowledge-based learning environments provide an ideal testbed for developing and evaluating computational models of mixed initiative interaction. In each problem-solving episode, learners incrementally develop solutions for problems posed by the environment. To maximize learning effectiveness and learning efficiency, we have been developing animated pedagogical agents that dynamically provide explanatory advice in response to changing problem-solving contexts in learning environments. Animated pedagogical agents monitor students' problem-solving activities and intervene with explanations in appropriate contexts. When students reach impasses, agents take control of the interaction and provide appropriate assistance. Dynamically controlling complex problem-solving interactions requires animated agents to make runtime decisions about when to intervene, how to select the content, with what level of directness to present hints, and with which media the advice should be delivered. This paper outlines some fundamental issues in mixed initiative problem solving with animated pedagogical agents and presents implemented solutions to these problems.

Introduction

Mixed initiative problem solving lies at the heart of knowledge-based learning environments. Since the birth of the field more than twenty-five years ago (Carbonell 1970), it has become apparent that developing computational models of mixed initiative is critical to the learning environment enterprise. While learners are actively engaged in problem solving activities, learning environments should monitor their progress and provide them with feedback in a manner that contributes to achieving the twin goals of learning effectiveness and learning efficiency. By carefully monitoring students' progress, learning environments should control the course of the interaction in such a way that

*Support for this work was provided by the IntelliMedia Initiative of North Carolina State University and Novell and equipment donations from Apple and IBM.

they maximize the quality of the learning experience as well as the speed with which learning occurs.

We have recently begun to see the emergence of a new generation of knowledge-based learning environments that are inhabited by *animated pedagogical agents* (Rickel & Johnson ; Stone & Lester 1996). Building on developments in believable agents with life-like qualities (Bates 1994; Tu & Terzopoulos 1994; Granieri *et al.* 1995; Blumberg & Galyean 1995; Kurlander & Ling 1995; Maes *et al.* 1995; André & Rist 1996), these intriguing interactive characters couple key feedback functionalities with a strong visual presence. Introduced immersively into learning environments, they observe students' progress and provide them with visually contextualized problem-solving advice. Learning environments with animated pedagogical agents provide an excellent "laboratory" for developing computational models of mixed initiative interaction. By designing an agent, implementing it in a learning environment, and observing student-agent interactions, we can acquire a body of empirical evidence which can inform decisions about which intervention strategies are effective and which ones are less so.

Mixed initiative control of animated pedagogical agents poses many difficult challenges for agent designers. Agents must be able to track students' activities through a series of problem-solving episodes, intervene when appropriate, and provide advice that promotes learning. Moreover, because agents exhibit multimedia behaviors—these include both verbal utterances and physical actions—their successful deployment depends to a large degree on the extent to which the advice presented by the agent employs media that are appropriate for the context at hand.

To investigate these issues, we are engaged in long-term research program in animated pedagogical agents. The current focus of this work is developing *behavior sequencing engines* that dynamically control the behaviors of animated pedagogical agents in response to the rapidly changing problem-solving contexts of knowledge-based learning environments. To produce effective interactions, behavior sequencing engines must address two fundamental issues in mixed



Figure 1: DESIGN-A-PLANT's animated pedagogical agent, Herman the Bug

initiative: controlling initiative and providing explanations. We have designed and implemented an animated pedagogical agent, Herman the Bug (Figure 1), for DESIGN-A-PLANT (Lester *et al.* 1996; Lester, FitzGerald, & Stone 1997), a learning environment in the domain of botanical anatomy and physiology for middle school students.¹ The agent interactively provides advice to students as they graphically assemble plants from a library of plant structures, i.e., roots, stems and leaves. In DESIGN-A-PLANT, a student's goal in each problem-solving episode is to design a plant that will thrive in a given natural environment with specified conditions such as the amount of available sunlight. As they construct plants to solve the problem, the agent provides them with advice about botanical anatomy and physiology.

This paper overviews our work in mixed initiative interaction for animated pedagogical agents. It is structured as follows. We first examine two fundamental problems in mixed initiative problem solving with animated pedagogical agents in knowledge-based learning environments. We then outline our solutions to these problems and describe the behavior sequencing engine which orchestrates the behaviors of the implemented

¹DESIGN-A-PLANT is a multi-disciplinary project involving computer scientists, multimedia designers, animators, and cognitive scientists. All of the 3D graphics and animations were designed, modeled, and rendered on Macintoshes and SGIs by a twelve-person graphic design team.

agent of DESIGN-A-PLANT. We conclude by describing evaluations of this agent and discussing current activities in the design of animated agents for mixed initiative learning environments.

Mixed Initiative Interaction in Learning Environments

For the past twenty-five years, researchers in knowledge-based learning environments have endeavored to create systems that provide highly customized problem-solving experiences that are tailored to the individual needs of each student (Anderson, Boyle, & Yost 1985; Clancey 1987; Collins & Stevens 1982; Lesgold *et al.* 1992; Woolf *et al.* 1987; Hollan, Hutchins, & Weitzman 1987). While the recent introduction of animated pedagogical agents into learning environments has created the opportunity to significantly improve students' problem solving, controlling mixed initiative student-agent interactions is an imposing task. Animated pedagogical agents' behaviors are guided by behavior sequencing engines that select and assemble agents' behaviors in real time. To cope with rapidly changing problem-solving contexts, behavior sequencing engines must deal effectively with two fundamental problems in mixed initiative problem-solving interaction: controlling initiative and providing explanations.

Controlling Initiative. Student-agent interactions are characterized by problem-solving episodes where control of the initiative frequently changes. At the beginning of each episode, students are unfamiliar with the problem, so the agent should take control and introduce it. For example, in *DESIGN-A-PLANT*, Herman describes the environmental conditions that hold in a particular environment for which a plant will be designed. Once students begin solving problems, the initiative typically changes many times. Students should be able to take control while they are performing problem-solving actions, agents should regain control when it appears that students are experiencing difficulty or when students ask a question, and control should be relinquished to students so they may continue their problem solving. For example, in *DESIGN-A-PLANT*, Herman monitors students as they assemble plants and intervenes to provide explanations about botanical anatomy and physiology when they reach an impasse. Finally, once problem solving is successfully completed by the student, the agent should again regain control to complete the problem solving transaction. This may involve a simple statement that a correct solution has been proposed or perhaps a more elaborate congratulatory utterance accompanied by a visual display. For example, at the end of each successful problem-solving episode in *DESIGN-A-PLANT*, Herman congratulates students.

Providing Explanations. When an animated pedagogical agent takes the initiative, it is typically to provide an explanation. While the surface reason for the intervention is to provide problem-solving advice, the deeper reason is to provide students with a clear explanation of concepts in the domain. It is our belief that animated pedagogical agents should provide explanations that are situated (Suchman 1987; Brown, Collins, & Duguid 1989): they should clearly communicate fundamental conceptual knowledge and advice in concrete problem-solving contexts. For example, suppose students are interacting with a learning environment that is designed to provide instruction about the structure and function of a complex device. Rather than adopting a didactic approach in which the nomenclature, components, and processes are introduced in the abstract, the learning environment should provide students with explanations as they engage in goal-driven problem-solving activities such as designing a component of the device or diagnosing the device's malfunctions. To illustrate, when Herman interacts with students using *DESIGN-A-PLANT*, he should help them learn about leaf morphology in the context of selecting a particular type of leaf as they design a plant that will thrive in particular environmental conditions. In addition to situatedness, the agent should deliver its advice using appropriate combinations of media. For example, introductions of concepts can make use of rich visual animations while reminders of

previously presented material can rely on brief verbalizations uttered by the agent without disrupting problem solving.

Mixed Initiative Problem Solving with Animated Pedagogical Agents

Empirically studying mixed initiative interaction requires three entities:

- *Testbed Learning Environment:* A learning environment whose problem-solving is sufficiently complex to *require* rich student-agent interactions. To this end, we have developed a design-centered learning environment in which students learn by designing artifacts to satisfy given specifications.
- *Problem-Solving Context Model:* A representation of the problem-solving context that encodes the critical features of students' learning episodes. To support design-centered learning, our context model consists of an environmental context, artifactual context, and an advisory context.
- *Animated Pedagogical Agent:* An implemented life-like character whose behavior sequencing engine monitors the model of the problem-solving context to control the initiative and to provide explanations. The implemented agent's behaviors are controlled by a sequencing engine that attends to the state of the context models.

We discuss each of these in turn.

Design-Centered Learning Environments

To create a challenging testbed for studying mixed initiative interaction, we have been exploring mixed initiative in the context of constructivist problem solving. Constructivism holds as its most central tenet that learning *is* building knowledge structures (Piaget 1954), and one of the most promising techniques in the constructivist's arsenal is *design*. Whether it is a child assembling a house from building blocks, an engineering student laying out a circuit, or a graduate student in mathematics constructing an axiomatic theory, the learner is actively engaged in a process that requires them to grapple with fundamental issues in their respective domains. Design-centered learning environments can be developed for domains as diverse as biology (e.g., designing plants), chemistry (e.g., synthesizing compounds), or the social sciences (e.g., building a society, as in the popular Maxis *SIM* series).

DESIGN-A-PLANT is a design-centered learning environment we created for botanical anatomy and physiology. Its problem-solving episodes play out as follows. Students are given an environment that specifies biologically critical factors in terms of qualitative variables. Environmental specifications for these episodes include the average incidence of sunlight, the amount of nutrients in the soil, and the height of the water table. Students consider these conditions as they inspect

components from a library of plant structures that is segmented into roots, stems, and leaves. Each component is defined by its structural characteristics such as length and thickness. Employing these components as their building blocks, students work in a “design studio” to graphically construct a customized plant that will flourish in the environment. Each iteration of the design process consists of inspecting the library, assembling a complete plant, and testing the plant to see how it fares in the given environment. If the plant fails to survive, students modify their plant’s components to improve its suitability, and the process continues until they have developed a robust plant that prospers in the environment.

A Tripartite Context Model

To achieve effective mixed initiative interactions in learning environments, it is critical to dynamically model the problem-solving context. Modeling problem solving is particularly critical in design-centered learning environments because animated pedagogical agents must assist students in making decisions that require them to consider multiple environmental factors, multiple components, and multiple constraints simultaneously. To furnish agents’ behavior sequencing engines with a representation of students’ problem solving, we model problem solving contexts with a tripartite contextual representation of design episodes (Lester *et al.* 1996) that is dynamically maintained by the learning environment. The model consists of an environmental context, an artifactual context, and an advisory context:

- *Environmental Context*: Critical features of the environment which have been presented to the student:
 - Current Environment: Environmental factors (and their values) in the current design episode.
 - Environmental Intent: Associated with each environment presented to the student is the set of object types from the artifact library which that environment is intended to exercise, e.g., some environments are presented to exercise students’ knowledge of leaf morphology.
- *Artifactual Context*: Critical features of the artifact under construction:²
 - Partial Solutions: Selections of components for the current artifact under construction, e.g., the student may have completed the roots and leaves.
 - Focused Component: Artifact component to which the student is currently attending, e.g., the stem.
 - Design Evaluation: When the student completes the design, the artifact is evaluated as successful or not successful in the current environment.

²A detailed discussion of the artifact-based task modeling mechanism may be found in (Lester, FitzGerald, & Stone 1997).

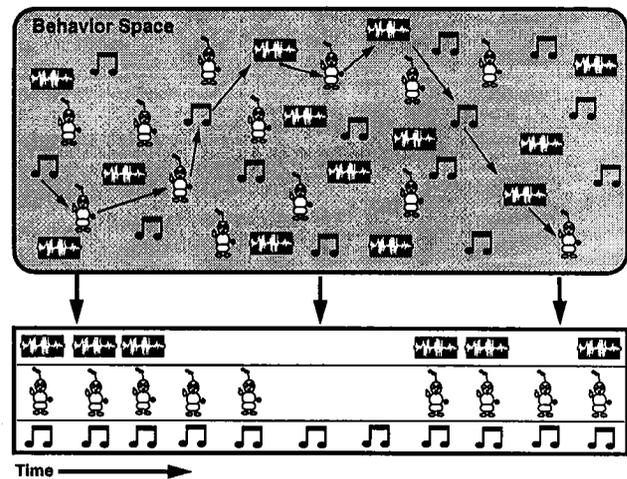


Figure 2: Sequencing coherence-structured behaviors

- *Advisory Context*: Critical features of the advisory dialogue:
 - Topic: Environmental factors, artifact components, constraints, and/or design decisions which have been addressed in advice.
 - Frequency Annotations: Indicate the number of times that the student has been advised about the topic(s).
 - Problem-Solving Idle Time: Time expired since the student’s last action.
 - Media Annotations: Indicate the media, e.g., audio or animation, that were employed to communicate the advice.

Animated Pedagogical Agents

To control the initiative and to provide explanations, we employ the *coherence-structured behavior space* framework (Stone & Lester 1996) for sequencing the behaviors of animated pedagogical agents. Applying the coherence-structured behavior space framework (Figure 2) to create an agent entails: (1) constructing a *behavior space*, which contains very short animated segments of the agent performing a variety of actions, e.g., a gesture or a walk cycle, and (2) constructing a *behavior sequencing engine* that creates global behaviors in response to the changing problem-solving context by navigating coherent paths through the behavior space and assembling them dynamically.

To provide agents with the flexibility required to respond to a broad range of problem-solving contexts, their behavior spaces are populated with a large, diverse set of behaviors. Using the the behavior canon of the animated film (Noake 1988) as our starting point, we build on previous computational models of classical animation principles (Lasseter 1987) to specify behaviors. Behavior spaces contain: *manipulative* behaviors (e.g., the agent picks up and manipulates objects), *vi-*

sual attending behaviors (e.g., the agent watches a student perform an action), *re-orientation* (e.g. the agent stands up), *locomotive* behaviors (e.g., the agent moves from one location to another), *gestural* behaviors (e.g., the agent gestures with its body and arms), and *verbal* behaviors (e.g., the agent utters a phrase). In response to changing problem-solving contexts in DESIGN-A-PLANT, the agent's sequencing engine controls the initiative by monitoring the tripartite context model and then providing explanations by selecting and assembling behaviors from its behavior space.

To illustrate, DESIGN-A-PLANT's agent, Herman, is a talkative, quirky, somewhat churlish insect with a propensity to fly about the screen and dive into the plant's structures as it provides students with problem-solving advice. His behavior space consists of more than 30 animated behaviors, 160 utterances, and a large library of runtime-mixable, soundtrack elements. Throughout the learning session, he remains onscreen, standing on the plant assembly device when he is inactive (Figure 1) and diving into the plant as he delivers advice visually. In the process of explaining concepts, he performs a broad range of activities including walking, flying, shrinking, expanding, swimming, fishing, bungee jumping, teleporting, and acrobatics. All of his behaviors are sequenced in realtime on a Power Macintosh 9500/132.

The behavior sequencing engine monitors the state of the artifactual context to determine when the student requires assistance. If the student makes an incorrect design decision (as indicated by his or her partial solutions), or if the problem-solving idle time exceeds a threshold, then the agent will take control of the interaction. When an intervention is triggered, the sequencing engine must determine the topic of the advice it will provide to the student. If the current partial solution indicates that only a single component is inappropriate for the current environment, the agent will provide advice about that component. If multiple components are inappropriate, the sequencing engine inspects the focused component (the component to which the student is currently attending); if this component is incorrect, the agent will provide advice about it. Otherwise, the sequencing engine inspects the environmental intent of the current environment and determines if one of the inappropriate components is the subject of the environmental intent. If so, the agent will provide advice about that component.

Once the decision to take control has been made and the topic of the intervention has been determined, the content of the explanation must be determined. This decision is governed by presentation strategies that first select the appropriate level of *directness* of the advice and then select a communication medium. The directness spectrum of advisory presentations represents the degree of explicitness with which presentations inform students about design decisions. The least direct advice discusses the information about constraints, e.g.,

the functional relation between environmental factors and artifact components; the most direct advice suggests a specific design decision. While direct advice is easily operationalized, the opportunity for learning is reduced, so indirect advice is generally preferred.

The sequencing engine first selects a point on the directness spectrum, and then uses this decision to determine the presentation media. To select a point on the directness spectrum, four factors are considered: (1) if the topic of the advice is the environmental intent, indirect advice is preferred; (2) if the student has advanced to problems of greater complexity, indirect advice is preferred; (3) if the student is experiencing difficulty (as indicated by his or her partial solutions), more direct advice is preferred; and (4) if advice about a particular topic has already been presented (as indicated by the advisory context), more direct advice is preferred. The selected level of directness is then used to make media selection decisions: the more indirect the advice is, the more likely it will be presented as animations depicting interactions between environmental factors and artifact components; the more direct the advice is, the more likely it will be presented primarily as speech. Finally, the topic, the selected level of directness, and the selected medium are used to select behaviors from the behavior space and to assemble them for exhibition in realtime.

Example Interaction

To illustrate how the behavior sequencing engine controls the initiative and provides advice, suppose that a student has just completed designing a plant for a simple environment. The agent now takes the initiative and describes an environment, in which the roots and leaves are both in focus. The key environmental factors (those that affect objects in the environmental intent) are a low temperature and a high water table, which are displayed onscreen.

Animated Agent: "Hooray, a pretty place. It's absolutely lovely. Of course the ground reminds me of a skating rink. Maybe that's because of the low temperature and high water table. Make sure that the stems are thick and well protected and that the roots and leaves can handle the pretty but harsh conditions."

To focus on roots and leaves, the agent's introduction gives specific advice about the stem only.

Student: Clicks on the upper portion of the plant construction area to begin working on leaves. Spends thirty seconds going back and forth with the mouse from the rollover textual descriptions of the environment to the rollover textual descriptions of the leaf choices but cannot make a decision.

The problem-solving idle time threshold is exceeded. There are three constraints that map

cold temperature to leaf features. Because this is the first time that the student has required advice during leaf selection in this environment and leaves are in focus, three animated explanations are played in sequence. They explain the relationship between cold temperature and leaf size, leaf thickness, and leaf skin thickness. These explanations are followed by the first animated lesson on leaf anatomy.

Student: Makes two more unsuccessful attempts at selecting a correct leaf.

The leaf choice still violates a constraint, but now the student has already seen a detailed animated explanation and has been given a verbal reminder. Therefore, short direct verbal advice is given.

Animated Agent: “Why don’t you try a small thick leaf with nice thick skin?”

Evaluation

To gauge the effectiveness of the animated agent’s behaviors, formative observational studies were conducted with thirteen middle school students using DESIGN-A-PLANT. Each student interacted with the learning environment for forty-five minutes to one hour. As the students designed plants for a variety of environmental conditions, the agent introduced problems, explained concepts in botanical anatomy and physiology, provided problem-solving advice, and interjected congratulatory and off-the-cuff remarks. These studies suggest that a behavior sequencing engine that employs a tripartite problem-solving context model can effectively control the initiative and provide explanations. Herman’s verbal reminders enabled students to continue with their problem solving uninterrupted, and during the study students made frequent (and unprompted) positive comments about Herman’s physical actions and remarks. The variety of his behaviors maintained their interest throughout the sessions, and Herman’s quirky verbal asides were well received.

The study also revealed three problems with the initial algorithm. First, in the original version, the agent provided its advice *before* giving the conceptual explanations. Students tended to forget this advice because, we hypothesize, there were intervening conceptual explanations. The sequencing engine’s assembly mechanism was therefore modified to present problem-solving-specific advice after explanations of fundamental concepts. Second, students were irritated by the repetition of the agent’s explanations. We therefore modified the selection mechanism to ensure that explanations would be repeated only if sufficient time had elapsed. Third, the initial version permitted only isolated explanations to be exhibited, so methods for sequencing multiple explanations were developed. Each of these problems has been addressed in the current implementation.

Current Directions

The encouraging results from the formative evaluation set the stage for the next phase of the work. This phase involves conducting large-scale formal evaluations and developing a new approach to behavior sequencing that promises greater flexibility. With the results from the formative study in hand, we have set out to conduct a series of formal evaluations of the cognitive effects of animated pedagogical agents in mixed initiative problem solving. To this end, we have conducted a formal study of the effectiveness of animated pedagogical agents in different conditions. We recently completed the data collection in a study involving more than 100 middle school students. While analyses of this data is still underway, our first analyses demonstrate *persona effect*, which is that the presence of a lifelike character in an interactive learning environment—even one that is not expressive—can have a strong positive effect on student’s perception of their learning experience (Lester *et al.* in press).

To complement the cognitive evaluation work, we have begun developing second generation agents whose ability to provide effective mixed initiative interactions considerably exceeds their predecessors. In particular, we are developing agents that: (1) make use of a rich dialog history; (2) exhibit believability-enhancing behaviors that makes them more lifelike (Lester & Stone in press); (3) exhibit fine-grained behaviors (as opposed to coarse-grained behaviors, each of which is performed over a longer period of time); and (4) can accompany students in learning environments whose topology is created at runtime and that can be navigated by students in the course of their problem-solving activities. Together, it appears that these capabilities will significantly increase the quality of mixed initiative student-agent interactions.

Conclusion

Knowledge-based learning environments can benefit considerably from animated pedagogical agents which engage in mixed initiative problem-solving interactions with students. This work represents a promising first step toward creating animated pedagogical agents with the ability to control the initiative and to provide appropriate explanations. By designing animated pedagogical agents, introducing them into implemented learning environments, and empirically evaluating the effectiveness of their initiative control and explanation mechanisms, we can gain significant insight into the fundamental nature of mixed initiative interaction. As a result of these efforts, we can accelerate the development of learning environments that effectively promote rich problem solving. More generally, this methodology promises to hasten the deployment of human-agent dialog technologies for a broad range of applications.

Acknowledgements

Thanks to: the animation team directed by Patrick FitzGerald; the students in the Intelligent Multimedia Communication, Multimedia Interface Design, and Knowledge-Based Multimedia Learning Environments seminars; and Chris Tomasson and her seventh grade class at Martin Middle School for participating in the formative evaluations.

References

- Anderson, J.; Boyle, C.; and Yost, G. 1985. The geometry tutor. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, 1-7.
- André, E., and Rist, T. 1996. Coping with temporal constraints in multimedia presentation planning. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, 142-147.
- Bates, J. 1994. The role of emotion in believable agents. *Communications of the ACM* 37(7):122-125.
- Blumberg, B., and Galyean, T. 1995. Multi-level direction of autonomous creatures for real-time virtual environments. In *Computer Graphics Proceedings*, 47-54.
- Brown, J.; Collins, A.; and Duguid, P. 1989. Situated cognition and the culture of learning. *Educational Researcher* 18(1):32-42.
- Carbonell, J. R. 1970. AI in CAI: An artificial-intelligence approach to computer-assisted instruction. *IEEE Transactions on Man-Machine Systems* 4:190-202.
- Clancey, W. J. 1987. *Knowledge-based Tutoring: The GUIDON Program*. MIT Press.
- Collins, A., and Stevens, A. L. 1982. Goals and strategies of inquiry teachers. In Glaser, R., ed., *Advances in Instructional Psychology, Volume 2*. Hillsdale, NJ: Lawrence Erlbaum. 65-119.
- Granieri, J. P.; Becket, W.; Reich, B. D.; Crabtree, J.; and Badler, N. I. 1995. Behavioral control for real-time simulated human agents. In *Proceedings of the 1995 Symposium on Interactive 3D Graphics*, 173-180.
- Hollan, J. D.; Hutchins, E. L.; and Weitzman, L. M. 1987. STEAMER: An interactive, inspectable, simulation-based training system. In Kearsley, G., ed., *Artificial Intelligence and Instruction: Applications and Methods*. Reading, MA: Addison-Wesley. 113-134.
- Kurlander, D., and Ling, D. T. 1995. Planning-based control of interface animation. In *Proceedings of CHI '95*, 472-479.
- Lasseter, J. 1987. Principles of traditional animation applied to 3D computer animation. In *Proceedings of SIGGRAPH '87*, 35-44.
- Lesgold, A.; Lajoie, S.; Bunzo, M.; and Eggan, G. 1992. SHERLOCK: A coached practice environment for an electronics trouble-shooting job. In Larkin, J. H., and Chabay, R. W., eds., *Computer-Assisted Instruction and Intelligent Tutoring Systems: Shared Goals and Complementary Approaches*. Hillsdale, NJ: Lawrence Erlbaum. 201-238.
- Lester, J. C., and Stone, B. A. in press. Increasing believability in animated pedagogical agents. To appear in *Proceedings of the First International Conference on Autonomous Agents*.
- Lester, J. C.; Stone, B. A.; O'Leary, M. A.; and Stevenson, R. B. 1996. Focusing problem solving in design-centered learning environments. In *Proceedings of the Third International Conference on Intelligent Tutoring Systems*, 475-483.
- Lester, J. C.; Converse, S. A.; Kahler, S. E.; Barlow, S. T.; Stone, B. A.; and Bhogal, R. in press. The persona effect: Affective impact of animated pedagogical agents. To appear in *Proceedings of CHI'97 (Conference on Human Factors in Computing Systems)*.
- Lester, J. C.; FitzGerald, P. J.; and Stone, B. A. 1997. The pedagogical design studio: Exploiting artifact-based task models for constructivist learning. In *Proceedings of the Third International Conference on Intelligent User Interfaces*, 155-162.
- Maes, P.; Darrell, T.; Blumberg, B.; and Pentland, A. 1995. The ALIVE system: Full-body interaction with autonomous agents. In *Proceedings of the Computer Animation '95 Conference*.
- Noake, R. 1988. *Animation Techniques*. London: Chartwell.
- Piaget, J. 1954. *The Construction of Reality in the Child*. New York: Basic Books.
- Rickel, J., and Johnson, L. Integrating pedagogical capabilities in a virtual environment agent. To appear in *Proceedings of the First International Conference on Autonomous Agents*.
- Stone, B. A., and Lester, J. C. 1996. Dynamically sequencing an animated pedagogical agent. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, 424-431.
- Suchman, L. 1987. *Plans and Situated Actions: The Problem of Human Machine Communication*. Cambridge University Press.
- Tu, X., and Terzopoulos, D. 1994. Artificial fishes: Physics, locomotion, perception, and behavior. In *Computer Graphics Proceedings*, 43-50.
- Wolf, B.; Blegan, D.; Jansen, J. H.; and Verloop, A. 1987. Teaching a complex industrial process. In Lawler, R. W., and Yazdani, M., eds., *Artificial Intelligence and Education*, volume 1. Norwood, New Jersey: Ablex. 413-427.