Animated Pedagogical Agents and Problem-Solving Effectiveness: A Large-Scale Empirical Evaluation

James C. Lester
Department of Computer Science

Sharolyn A. Converse
Department of Psychology

Brian A. Stone
Department of Computer Science

Susan E. Kahler
Department of Psychology

S. Todd Barlow
Department of Psychology

North Carolina State University
Raleigh, NC 27695-8206 USA

Email: lester@csc.ncsu.edu

Abstract

Animated pedagogical agents offer great promise for learning environments. In addition to their significant motivational benefits, these intriguing lifelike characters can also play an important pedagogical role. Introduced into a rich learning environment, animated pedagogical agents could increase problem-solving effectiveness by providing students with customized advice. To evaluate the effects of animated pedagogical agents on students' problem solving, we conducted a large-scale formal study with 100 middle school students interacting with an animated pedagogical agent that inhabits Design-A-Plant, a design-centered learning environment for botanical anatomy and physiology. The lifelike agent advises students as they design plants to survive in different environmental conditions. In the study, students interacted with one of five different clones of the agent. The results of the study demonstrate that animated pedagogical agents can yield important educational benefits in the form of improved problem solving, particularly for complex problems. They also show that animated pedagogical agents that provide multiple levels of advice with multiple modalities yield greater improvements in problem solving than less expressive animated pedagogical agents.

1 Introduction

Endowing animated agents with believable, lifelike qualities has been the subject of a growing body of research [1, 2, 3, 4]. Because of the prospect of combining key feedback functionalities with a strong visual presence, *animated pedagogical agents* [11, 12] that provide clear explanations to students hold much appeal. Introduced immersively into learning environments, animated pedagogical agents could observe students' progress and provide them with visually contextualized problem-solving advice. By creating the illusion of life, animated pedagogical agents have the potential to significantly improve the effectiveness of learning environments.

To design the most effective learning environments, it is essential to understand how animated pedagogical agents can contribute to students' learning experiences. Because recent advances in affordable graphics hardware are beginning to make the widespread distribution of animated pedagogical agents a reality, we have reached a critical juncture in assessing their potential benefits. Although much interesting research has examined the social aspects of human-computer interaction and users' anthropomorphization of software [9], little is known about the role of animated agents in learning environments. Our work addresses these problems by designing and implementing animated pedagogical agents, introducing them into learning environments, and empirically evaluating students' interactions with them. In previously reported work, we discovered the *persona effect*, which is that the presence of a lifelike character in an interactive learning environment can have a strong positive effect on students' perception of their learning experience [6].

Because the motivational benefits derived from the persona effect are compelling, it has become increasingly important to evaluate animated pedagogical agents' effects on students' problem solving. In particular, it would be useful to obtain evidence for answering each of the following questions. Can animated pedagogical agents contribute to students' problem solving or do their activities interfere with problem solving? Which type of advice (*principle-based* advice, which is abstract, or *task-specific* advice, which is direct and concrete) or combinations of types of advice could agents employ to improve problem solving? Which modality (visual or auditory) or combinations of modalities could agents employ to improve problem solving?

To investigate these issues, we have constructed an animated pedagogical agent, Herman the Bug [12] (Figure 1) and introduced it into DESIGN-A-PLANT [8], a learning environment for the domain of botanical anatomy and physiology.[1] The agent exhibits strikingly lifelike behaviors as it provides advice to students solving problems. We created five *clones* of the agent, each with different behavioral characteristics combining different levels of advice and different modalities. As a control, one of the clones was mute, i.e., it provided no advice. Finally, we conducted a large-scale formal study with 100 middle school students, each of whom interacted with one of the clones. During the interactions, the learning environment logged all of their problem-solving activities, and the students were administered rigorous pre-tests and post-tests. The results of the study are three-fold:

- *Problem-Solving Effectiveness:* Animated pedagogical agents can yield important educational benefits by improving students' problem solving.
- *Multi-Level, Multi-Modality Effects:* Animated pedagogical agents that provide multiple levels of advice combining multiple modalities yield greater improvements in problem solving than less expressive agents.
- *Complexity Benefits:* The benefits of animated pedagogical agents increase with problem-solving complexity. As students are faced with more complex problems, the positive effects

[1]All of the animated pedagogical agent's behaviors were designed, modeled, and rendered on Macintoshes and SGIs by a large multi-disciplinary team of computer scientists, graphic designers, and animators.

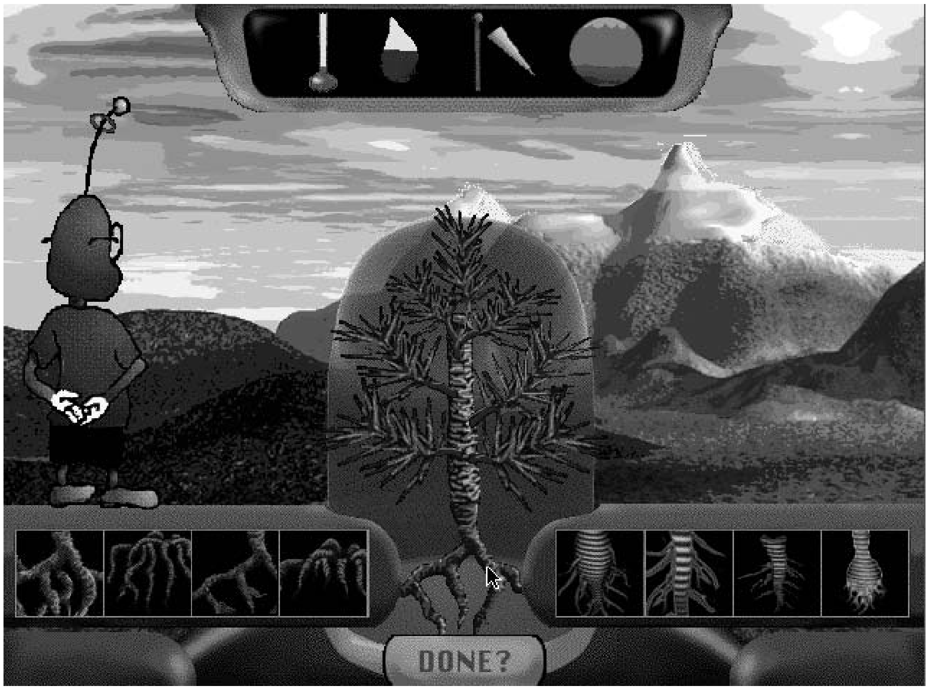

Figure 1: The animated pedagogical agent, Herman the Bug, in the DESIGN-A-PLANT learning environment

of animated pedagogical agents on problem solving are more pronounced.

2 Animated Pedagogical Agents and Learning Environments

Conducting a formal empirical study of animated pedagogical agents' effects on students' problem solving requires full-scale implementations of (1) a *behavior sequencing engine* that dynamically assembles an animated pedagogical agent's behaviors, (2) a *learning environment* in which students solve problems and interact with the agent. We briefly overview each of these in turn.

2.1 The Animated Pedagogical Agent Behavior Sequencing Engine

To create a lifelike presence for animated pedagogical agents, we have developed the *coherence-structured behavior space* framework [12] for sequencing their behaviors. Applying the coherence-structured behavior space framework 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 [10] as a starting point, the design of agents' behaviors builds on previous computational models of classical animation principles [5] to specify behaviors.

Behavior spaces contain: *manipulative* behaviors (e.g., the agent picks up and manipulates objects), *visual 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 an artifact-based task model [7] to select and assemble advisory behaviors from its behavior space.

To illustrate, the animated agent of DESIGN-A-PLANT is a talkative, quirky insect with a propensity to fly about the screen and dive into a 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. 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.

2.2 Design-Centered Learning Environments

To create a challenging testbed for studying animated pedagogical agents in a constructivist setting, we have developed a *design-centered* microworld in which students learn by designing artifacts that satisfy given environmental specifications. In DESIGN-A-PLANT [8], students explore the physiological and environmental considerations that govern plants' survival. Design-centered problem solving revolves around a carefully orchestrated series of design episodes.

DESIGN-A-PLANT's design 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.

3 Experimental Design

To evaluate animated pedagogical agents' effects on students' problem solving, we designed an experiment in which a large number of students interacted with different clones of our implemented agent in a controlled learning environment. Participants in the evaluation were 100 students (50 females; 50 males) who were enrolled at a local middle school. Students were assigned to interact with one of five clones of the agent (described below). These assignments were random except for the fact that an equal number of males and females were assigned to each of the five clones. Data from four participants were eliminated due to technical difficulties. The study was conducted in a classroom at the middle school at which the participants were enrolled.

To investigate how different characteristics of animated pedagogical agents affect students' problem solving, we developed five clones of the agent and introduced each into identical copies of the Design-A-Plant learning environment. There are three types of communicative behaviors that an agent can exhibit. One type of behavior is a short principle-based, "high-level" animated segment which combines animations of an object in the domain (e.g., the plant) and spoken descriptions by the agent to convey principle-based advice about the object. Students must then operationalize this advice in their problem-solving activities. The second type of communicative behavior is principle-based "high-level" advice spoken by the agent. This type of advice is similar to that provided by the advisory animations, but it is conveyed without the benefit of the accompanying animations. For example, Herman might say, "Remember that small leaves are struck by less sunlight." The final type of advisory behavior is a direct, task-specific suggestion spoken by the agent about what action the student should take. This advice is immediately operationalizable. For example, Herman might say, "Choose a long stem so the leaves can get plenty of sunlight in this dim environment." Each of the five clones differed from their siblings with respect to their modes of expression and the level of advice they offered in response to students' problem-solving activities:

- **Muted**: This agent can provide no advice at all about the plant components that are affected by the environmental factors. It serves as an experimental control.

- **Task-Specific Verbal**: This agent can only provide task-specific verbal advice.

- **Principle-Based Verbal**: This agent can only provide principle-based verbal advice.

- **Principle-Based Animated/Verbal**: This agent is limited to providing only principle-based animated advice accompanied by the spoken descriptions. It may not employ either the abstract audio-only advice as a reminder of previously seen animations nor may it offer direct verbal-only advice.

- **Fully Expressive**: This agent exhibits all of the three types of communicative behaviors. For example, it may give principle-based animated advice to challenge the student, or may employ the task-specific audio advice if the student is having difficulty.

Despite these differences, the clones are identical in all other respects. All are identical in appearance and in vocal qualities. Moreover, they all exhibit identical "non-advisory" behaviors.

Data were collected over the course of eight days. In each data collection session, four students came to the classroom. Each student was assigned to a researcher, who accompanied his or her student to one of the four workstations. Each data collection session lasted from one and one half to two hours. After completing a demographic questionnaire, each data collection session proceeded in four distinct phases: *Pre-testing*, *Agent Interaction*, *Post-Testing*, and *System/Agent Assessment*. To measure the student's knowledge of botanical anatomy and physiology before and after interacting with the learning environment, paper-and-pencil pre- and post-tests were administered. The pre- and post-tests consisted of 13 identical multiple-choice questions, but the order of the questions was varied between the pre- and post-tests. Each test question was carefully constructed to evaluate students' knowledge of botanical anatomy and physiology. For example, to test students' knowledge of the roots, one question asks, "If the water table is low, the plant most likely to survive will have which of the following types of roots: branching, not branching, deep, or shallow?" All of the questions on the pre- and post-tests were non-complex, i.e., they required students to consider only a single environmental factor.

When the student had completed the pre-test, he or she was taken to the computer workstation that provided an introductory animated presentation about botanical anatomy and physiology and information about how to select plant parts during the plant assembly task. The students then interacted with the learning environment, which required them to design eight plants for survival in four different environments. The first four problems were designed to teach the student about a single environmental constraint and its impact on the plant's probability of survival. The last four problems were more complex and required students to work with multiple constraints.

Students worked through the problems in the same order, starting with single constraint problems for each of the four environments, then advancing to multiple constraint problems for each of the environments. The problems and the order in which they were presented was identical, regardless of the type of agent clone inhabiting the environment. The learning environment logged every problem-solving action taken by the student (both correct and incorrect) as they designed a plant for each environment. A five-minute break was given to the students when they had finished interacting with the DESIGN-A-PLANT learning environment. Following the break, they completed the post-test. Finally, they were asked to complete a questionnaire indicating their opinions of the agent.[2]

4 Results and Statistical Analysis

Because of the magnitude of the study, it produced an enormous body of data bearing on animated pedagogical agents' impact on students' problem solving. We conducted three analyses of the data: (1) an analysis of the difficulty levels of the problems in DESIGN-A-PLANT; (2) an analysis of students' problem-solving performance as a function of the type of agent with which they interacted; and (3) an analysis of students' pre-test/post-test performance as a function of the type of agent with which they interacted.

First, we sought to determine whether the DESIGN-A-PLANT problems that were labeled as *complex* (because they involved a larger number of constraints) posed a greater challenge to students than problems that were labeled as *simple*. An analysis of students' problem-solving errors confirmed the difficulty ratings: students' problem-solving errors committed while attempting to solve the complex problems were statistically significantly greater than those committed while attempting to solve the simple problems.[3]

Second, we sought to determine agents' effects on students' problem-solving performance. We endeavored to ascertain: (1) whether agents improved or impeded problem solving, and (2) which levels of advice and which modalities (if any) could yield the most positive improvements. A summary of the problem-solving logs generated by the learning environments as students' interacted with them is shown in Table 1. There was a significant interaction between agent type and the

[2]The data gathered in the final stage revealed the *persona effect*, which is reported in [6].

[3]All data were analyzed by a three-way mixed model Multivariate Analysis of Covariance (MANCOVA).

Agent Type	*Mean Student Problem-Solving Errors for Simple Problems*	*Mean Student Problem-Solving Errors for Complex Problems*
Muted	1.11	6.79
Task-Specific Verbal	0.75	4.06
Principle-Based Verbal	0.88	4.86
Principle-Based Animated/Verbal	1.07	5.21
Fully Expressive	0.73	3.93

Table 1: Effect of agent type on students' problem-solving performance

problem-solving performance of students with respect to difficulty of problems [$F(4,91) = 5.65, p < 0.0004$)]. Although for simple problems there were no statistically significant differences between different types of agents' effects on students' problem-solving performance, for complex problems there was a strong difference produced by different types of agents. For complex problems, the fully expressive agent and the task-specific-verbal agent produced the best problem-solving performance in students. These agents produced levels of problem-solving performance that exceeded all of the others and were statistically significantly higher. The principle-based animated/verbal and the principle-based-verbal produced levels of performance that were not as high as the fully expressive and task-specific-verbal agents but were statistically significantly higher than the muted (control) agent.

Finally, we sought to determine agents' effects on students' performance using a traditional (i.e., non-interactive) measurement of learning. To this end, we analyzed the results of the pre-tests administered to students prior to their interaction with the software and the post-tests administered immediately afterward. A summary of the pre-test scores, post-test scores, and pre/post differences is shown in Table 2.[4] Although the differences between agents was not statistically significant, because the pre-tests and post-tests employed only simple problems, these results are in accord with the results of the analysis of agents' impact on problem-solving. Critically, every agent produced statistically significant *improvements* from pre-tests ($m = 4.85$) to post-tests ($m = 7.74$). The muted agent produced the smallest improvements, and the fully expressive agent produced the greatest improvements.

5 Conclusions

The study demonstrates that animated pedagogical agents can yield important educational benefits. Its findings are three-fold. First, it shows that animated pedagogical agents can improve students' performance. All of the active agents produced better student problem-solving perfor-

[4]That there was no significant difference between groups on pre-test scores indicates that students of equal prior botanical knowledge interacted with each type of agent. Note that the fourth column of Table 2 shows the means of the improvements. In contrast to the differences of the means, these means of the differences represent the means of the improvements between pre- and post-tests for each student.

Agent Type	*Mean Student Pre-Test Scores*	*Mean Student Post-Test Scores*	*Mean Student Improvement*
Muted	5.05	7.05	2.00
Task-Specific Verbal	5.50	8.35	2.85
Principle-Based Verbal	4.25	7.45	3.20
Principle-Based Animated/Verbal	5.15	8.30	3.25
Expressive	4.30	7.55	3.75

Table 2: Effect of agent type on test scores

mance in the learning environment than the muted agent. This finding is particularly noteworthy because the agents provide advice to students after the students committed a problem-solving error rather than prior to it. Hence, agents' advice delivered in response to earlier problems improves students' performance on later problems. Second, agents that provide multiple levels of advice and employ multiple modalities produce the best problem-solving performance. While all active agents improve interactive problem solving, those providing both principle-based and task-specific advice that employs both animated and auditory media yielded the best performance. Finally, the problem-solving benefits provided by agents are greatest for complex problems. While all agents produced significant gains from pre-tests to post-tests that consisted of simple problems, students most benefited from agents when they were faced with complex problems. As the level of problem difficulty increases, the magnitude of animated pedagogical agents' effects on students' problem-solving performance also increases. Furthermore, it is revealing that for complex problems, the expressive agent's ability to increase students' problem-solving performance was particularly strong relative to the other agents

These findings have many important implications for the design of educational software. However, it is important to consider two caveats. First, generalizing the findings to other age groups and domains must be done with great caution. Second, the long term effects of students' interactions with agents were not explored in this study. Bearing these in mind, three conclusions may be drawn from the study:

- *Animated pedagogical agents can improve problem solving.* While they are almost certainly not appropriate for all age groups and domains, animated pedagogical agents can have a positive effect on students' problem solving, particularly for complex problems.

- *Designers of learning environments for children should give careful consideration to including an animated pedagogical agent.* Coupled with previous findings that agents have a strong positive effect on student's perception of their learning experience, the study suggests that animated pedagogical agents can yield benefits in both motivation and performance.

- *In designing animated pedagogical agents, it is critical to endow them with the ability to provide principle-based advice, task-specific advice, animated advice, and verbal advice.* Because animated pedagogical agents that communicate with multiple levels of detail through multiple modalities yield greater gains in students' problem-solving performance, designers of learning environments should strive to create agents that are as expressive as possible.

In short, animated pedagogical agents hold much promise for significantly improving children's learning experiences. This large-scale empirical evaluation constitutes a first step toward identifying the contributions that agents can make to learning environments. The encouraging results call for further studies to pinpoint the precise age groups and domains that can most benefit from animated pedagogical agents.

Acknowledgements

Thanks to: the animation team which was lead by Patrick FitzGerald of the North Carolina State University School of Design; the students in the Intelligent Multimedia Communication, Multimedia Interface Design, and Knowledge-Based Multimedia Learning Environments seminars. For facilitating and participating in the evaluation we wish to thank: Karen Banks and Tony Lanier of the Wake County School System; Vice Principal Jim Palermo and the students of West Lake Middle School; and the Wake County Chapter of the Women in Science Mentoring Program. Support for this work was provided by the IntelliMedia Initiative of North Carolina State University, Novell, and equipment donations from Apple and IBM.

References

[1] Elisabeth André and Thomas Rist. Coping with temporal constraints in multimedia presentation planning. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, pages 142–147, 1996.

[2] Joseph Bates. The role of emotion in believable agents. *Communications of the ACM*, 37(7):122–125, 1994.

[3] John P. Granieri, Welton Becket, Barry D. Reich, Jonathan Crabtree, and Norman I. Badler. Behavioral control for real-time simulated human agents. In *Proceedings of the 1995 Symposium on Interactive 3D Graphics*, pages 173–180, 1995.

[4] David Kurlander and Daniel T. Ling. Planning-based control of interface animation. In *Proceedings of CHI '95*, pages 472–479, 1995.

[5] John Lasseter. Principles of traditional animation applied to 3D computer animation. In *Proceedings of SIGRAPH '87*, pages 35–44, 1987.

[6] James C. Lester, Sharolyn A. Converse, Susan E. Kahler, S. Todd Barlow, Brian A. Stone, and Ravinder Bhogal. The persona effect: Affective impact of animated pedagogical agents. In *Proceedings of CHI'97 (Human Factors in Computing Systems)*, pages 359–366, 1997.

[7] James C. Lester, Patrick J. FitzGerald, and Brian A. Stone. The pedagogical design studio: Exploiting artifact-based task models for constructivist learning. In *Proceedings of the Third International Conference on Intelligent User Interfaces*, pages 155–162, 1997.

[8] James C. Lester, Brian A. Stone, Michael A. O'Leary, and Robert B. Stevenson. Focusing problem solving in design-centered learning environments. In *Proceedings of the Third International Conference on Intelligent Tutoring Systems*, pages 475–483, 1996.

[9] Clifford Nass, Youngme Moon, B. J. Fogg, Byron Reeves, and D. Christopher Dryer. Can computer personalities be human personalities. *International Journal of Human-Computer Studies*, 43:223–239, 1995.

[10] Roger Noake. *Animation Techniques.* Chartwell, London, 1988.

[11] Jeff Rickel and Lewis Johnson. Integrating pedagogical capabilities in a virtual environment agent. In *Proceedings of the First International Conference on Autonomous Agents*, pages 30–38, 1997.

[12] Brian A. Stone and James C. Lester. Dynamically sequencing an animated pedagogical agent. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, pages 424–431, 1996.