

The Persona Effect: Affective Impact of Animated Pedagogical Agents

James C. Lester[†]
lester@adm.csc.ncsu.edu

Sharolyn A. Converse[‡]
sharolyn@poe.coe.ncsu.edu

Susan E. Kahler[‡]
sekahler@unity.ncsu.edu

S. Todd Barlow[‡]
stevent@unity.ncsu.edu

Brian A. Stone[†]
bastone@eos.ncsu.edu

Ravinder S. Bhogal[§]
b.ravinder@rca.ac.uk

[†]Department of Computer Science
North Carolina State University
Raleigh, NC 27695-8206
(USA)

[‡]Department of Psychology
North Carolina State University
Raleigh, NC 27695-7801
(USA)

[§]London Institute
Royal College of Art
Kensington Gore
London SW7 (UK)

ABSTRACT

Animated pedagogical agents that inhabit interactive learning environments can exhibit strikingly lifelike behaviors. In addition to providing problem-solving advice in response to students' activities in the learning environment, these agents may also be able to play a powerful motivational role. To design the most effective agent-based learning environment software, it is essential to understand how students perceive an animated pedagogical agent with regard to affective dimensions such as encouragement, utility, credibility, and clarity. This paper describes a study of the affective impact of animated pedagogical agents on students' learning experiences. One hundred middle school students interacted with animated pedagogical agents to assess their perception of agents' affective characteristics. The study revealed the *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. The study also demonstrates the interesting effect of multiple types of explanatory behaviors on both affective perception and learning performance.

KEYWORDS: Educational applications, intelligent systems, children, agents, empirical studies.

INTRODUCTION

Animated agents offer great promise for delivering sophisticated, realtime problem-solving advice with strong visual appeal. Because of their lifelike behaviors, the prospect of introducing these agents into educational software is especially appealing. In addition to the possibility of increasing students' learning effectiveness with customized feedback, *animated pedagogical agents* [19] may provide another im-

portant benefit: motivation. By creating the illusion of life, the captivating presence of the agents can motivate students to interact more frequently with agent-based educational software. This in turn has the potential to produce significant cumulative increases in the quality of a child's education over periods of months and years.

As a result of rapid advances in animated agent technology [1, 2, 3, 7, 9, 14, 18, 19, 21] and the growing availability of affordable graphics accelerators, the technical barriers preventing the introduction of animated agents into educational software are quickly being eliminated. Because deploying animated pedagogical agents on a broad scale will soon be possible, it is important to understand the nature of the contributions that they can make to educational software. While much interesting research has examined the social aspects of human-computer interaction and users' anthropomorphization of software [15, 16, 17], no large-scale formal empirical study has been conducted to assess the potential impact of lifelike animated pedagogical agents on students using interactive learning environments. To design the most effective learning environment software, it is essential to understand how students perceive an animated pedagogical agent that provides realtime problem-solving advice. How much credibility would students give such an agent? To what extent would students find an agent's advice helpful and clear? Given the opportunity to be assisted by an agent, would they prefer to have one or be left alone?

To study these issues, we conducted a formal controlled study of the affective impact of animated pedagogical agents on students using an interactive learning environment. Empirically studying animated pedagogical agents requires three entities:

- A fully functional *Animated Pedagogical Agent*: Students interacted with Herman the Bug (Figure 1), a lifelike agent developed in our laboratory by a large, multidisciplinary team of computer scientists, graphic designers, and animators. Because the agent's animated and verbal behaviors are assembled by a realtime behavior sequencing engine [19] in response to problem-solving actions taken by students, it can provide them with advice and explanations

that are individually tailored to their idiosyncratic learning strategies.

- An *Interactive Learning Environment*: In the study, students solved problems in a constructivist learning environment, DESIGN-A-PLANT [13, 11], which is also being developed in our laboratory. DESIGN-A-PLANT is a “design centered” learning environment in which students learn about botanical anatomy and physiology by designing plants that will thrive in given environmental conditions.
- A *Pool of Students*: One hundred middle school students participated in the study, which was conducted onsite at a middle school over the course of eight days. Each student interacted with the agent for approximately an hour.

To study the explanatory behaviors that might influence students’ perception of human-agent interaction, we developed five clones of the agent, each of which interacted with 20 students. Each clone was identical to the others in appearance, but it communicated with different explanatory behaviors. Some were more visually expressive while others were more verbally expressive; some provided high-level, principle-based advice while others provided low-level, task-specific advice; one provided no advice at all.

The study reveals that well-designed lifelike personae interacting with students using learning environments are perceived as being very helpful, credible, and entertaining. Surprisingly, this *persona effect* holds for highly muted lifelike agents. The study also reveals an important synergistic effect of multiple types of explanatory behaviors on students’ perception of agents: agents that are more expressive (both in modes of communication and in levels of advice) are perceived as having greater utility and clarity. Because of the combined force of the persona effect and its *agent expressivity corollary*, we believe these results have important implications for the design of educational software.

The paper is structured as follows. We first describe the interactive learning environment and the animated pedagogical agent. After describing the experimental method, we present the findings, which is followed by a discussion of the design implications for interactive learning environments.

DESIGN-A-PLANT

Because manipulable simulations offer experiences that are qualitatively different from more didactic approaches, recent years have experienced a resurgence of interest in *microworlds* [20, 8, 4, 10, 6]. DESIGN-A-PLANT [13, 11] is a design-centered microworld that provides students with the opportunity to explore the physiological and environmental considerations that govern plants’ survival. Design-centered problem solving revolves around a carefully orchestrated series of design episodes.

In each DESIGN-A-PLANT design episode, a student is 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 branching factor. 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.

Constraints relating environmental factors to artifact structures govern the composition of artifacts. For example, a design-centered learning environment for botanical anatomy and physiology might include the constraint that a low incidence of sunlight requires large leaves. Hence, in the course of designing artifacts for a variety of environments, students acquire an understanding of the (possibly complex) effects of the environment on artifact functionalities. By continuously designing and redesigning artifacts until they satisfy the given specifications, students gradually bridge the conceptual gap that separates specific environmental factors from specific artifact components.

THE ANIMATED PEDAGOGICAL AGENT

Students interacted with the animated pedagogical agent shown in (Figure 1). **Herman the Bug** is a lifelike agent whose visual and verbal actions are controlled by a realtime *behavior sequencing engine* [19, 12] in response to changing problem-solving contexts. It 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. Its behaviors include 30 animated segments 160 audio clips, and several songs. The animations were designed, modeled, and rendered on SGIs and Macintoshes by twelve graphic artists and animators. Twenty of them are in the 20–30 second range, and ten are in the 1–2 minute range. Herman’s actions are accompanied by a dynamically sequenced score which is assembled from a large library of runtime-mixable, soundtrack elements.

The behavior sequencing engine is based on the *coherence-structured behavior space* framework [19]. Applying this framework to create an agent entails constructing a behavior space, imposing a coherence structure on it, and developing a behavior sequencing engine that dynamically selects and assembles behaviors. A *behavior space* contains animated segments of the agent performing a variety of actions, as well as audio clips of the agent’s utterances. The behavior space is structured by (1) a tripartite index of ontological, intentional, and rhetorical indices, (2) a pedagogically appropriate prerequisite ordering, and (3) behavior links annotated with distances computed with a *visual continuity metric*. At runtime, the behavior sequencing engine creates global behaviors in response to the changing problem-solving context by exploiting the coherence structure of the behavior space. The sequencing engine selects the agent’s actions by navigating coherent paths through the behavior space and assembling them dynamically.

Throughout the learning session, Herman remains onscreen, standing in the “plant design studio” when he is inactive and

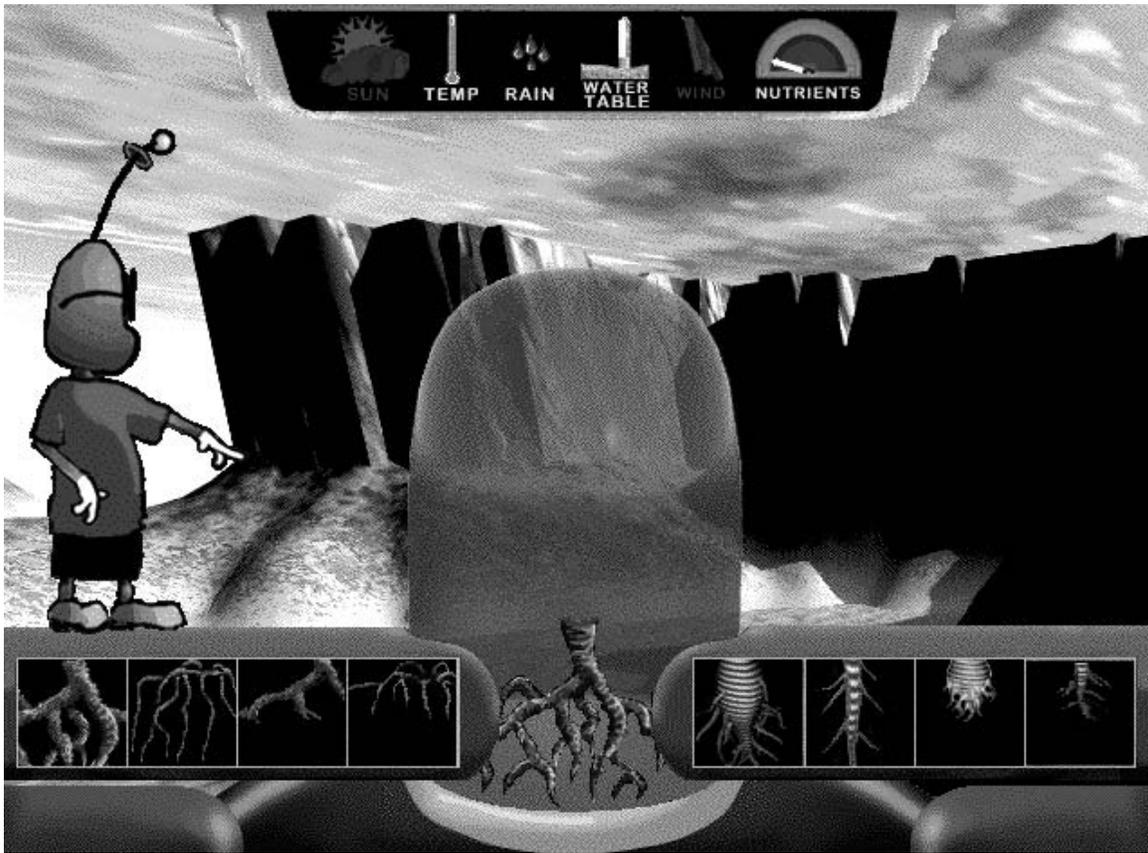


Figure 1: An animated pedagogical agent

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 its behaviors are sequenced on a highend Power Macintosh.

There are three types of communicative behaviors that the agent can exhibit. One type of behavior is a short 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 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."

To study how different classes of explanatory behaviors influence students' perception of human-agent interaction, we developed five Herman "clones" and introduced each one into

a copy of the DESIGN-A-PLANT learning environment. Each of the five clones differs from their siblings with respect to their modes of expression and in the level of advice they offered in response to students' problem-solving activities:

- **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 employ the task-specific audio advice if the student is having difficulty.
- **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.
- **Principle-Based Verbal:** This agent can only provide principle-based verbal advice.
- **Task-Specific Verbal:** This agent can only provide task-specific verbal advice.
- **Muted:** This agent can provide no advice at all about the plant components that are affected by the environmental factors.

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. For example, at the beginning of a learning session, each clone introduces himself and his spaceship and tells the students what types of tasks they will be requested

to perform. Also, each of the clones verbally introduces new problems by pointing at and then describing the environmental factors which define the problem, and it offers encouragement by expressing joy and congratulating the student when correct choices are made. For plant components which are not affected by the current environmental factors, it tells the students which to choose. Finally, it interacts with the environment in a visually intriguing way once the problem has been successfully completed. For example, it skis down mountains and bungee jumps.

METHOD

To empirically investigate students' perception of animated pedagogical agents, we sought to (1) expose a large number of students to agents in controlled learning experiences and (2) obtain students' assessment of the agents on a number of affective dimensions including helpfulness, clarity, and desirability. We also sought to determine the pedagogical effects of the agents on learning effectiveness.

Participants and Setting

Participants in the evaluation were 100 students (50 females; 50 males) who were enrolled at a local middle school. The average age of the students was 12 years. Students were recruited by their teachers who asked interested students to obtain a consent form from a parent or legal guardian. Students were assigned to interact with one of the five Herman clones. Assignment to clone was 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 are enrolled. The classroom contained four high-end Macintoshes, each with 80 MB, a one-button mouse, and color monitors. The software ran at a resolution of 640x480. Audio output was delivered in stereo on headphones.

Materials and Procedure

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. Of this time, students interacted with the agent for approximately one hour on average. Data were collected over the course of eight days. The students completed a consent form upon their arrival in the lab. This consent form assured them of the anonymity and confidentiality of their responses. They were then asked to complete a paper-and-pencil demographic questionnaire to provide the research team with information about the student's computer usage, including whether or not the student was comfortable using a computer, frequency of the student's computer usage, where the student used a computer, type of computer(s) the student used, type(s) of mouse(s) the student used, and tasks the students had performed on a computer. In addition, participants were asked to provide their age.

After the initial activities, each data collection session with a given student proceeded in four distinct phases: *Pre-testing*, *Agent Interaction*, *Post-Testing*, and *System/Agent Assess-*

ment (Figure 2). 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. However, the order of the 13 questions was different between the pre-test and the post-test. The questions were each constructed to evaluate the students' knowledge of the relationship between specific plant characteristics (e.g., large leaves) and the likelihood that a plant containing this characteristic would survive in a specific type of environment.

When the student had completed the pre-test, he or she was taken to the computer workstation where a computerized learning environment training module was viewed by the student. The training module first instructed the student how to interact with *Design-A-Plant's* interface. Concepts explained included the icons employed in the interface and how to select plant parts during the plant assembly task. The training module then advised the student that they would first be given information about plant parts and their relationship to survival in specific types of environments. Students were also told that they would be asked to design plants that would be likely to survive in certain environments once they had completed the initial tutorial. Students were then taught how to access the help from the animated pedagogical agent in the event that they made errors while designing their plants.

After completing the training module, the students interacted with the learning environment. The learning environment required the student 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. 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.

Students were then asked to complete a system/agent assessment evaluation. To reduce response bias, student were strongly encouraged to record their responses because the researchers wanted to use their responses to improve the software. The researchers also gave the students privacy as they responded. The system/agent assessment form consisted of 18 questions (Figure 3), each of which was presented on a Likert scale. Sixteen of the Likert scale questions contained five response categories which used terms that were associated with ratings such as *very good* (5), *good* (4), *neither good nor bad* (3), *poor* (2), and *very poor* (1). Two of the rating scales contained only four response categories. The form also asked students to record free-form responses. At the conclusion of the session, students were thanked and given the opportunity to ask any questions they may have had about

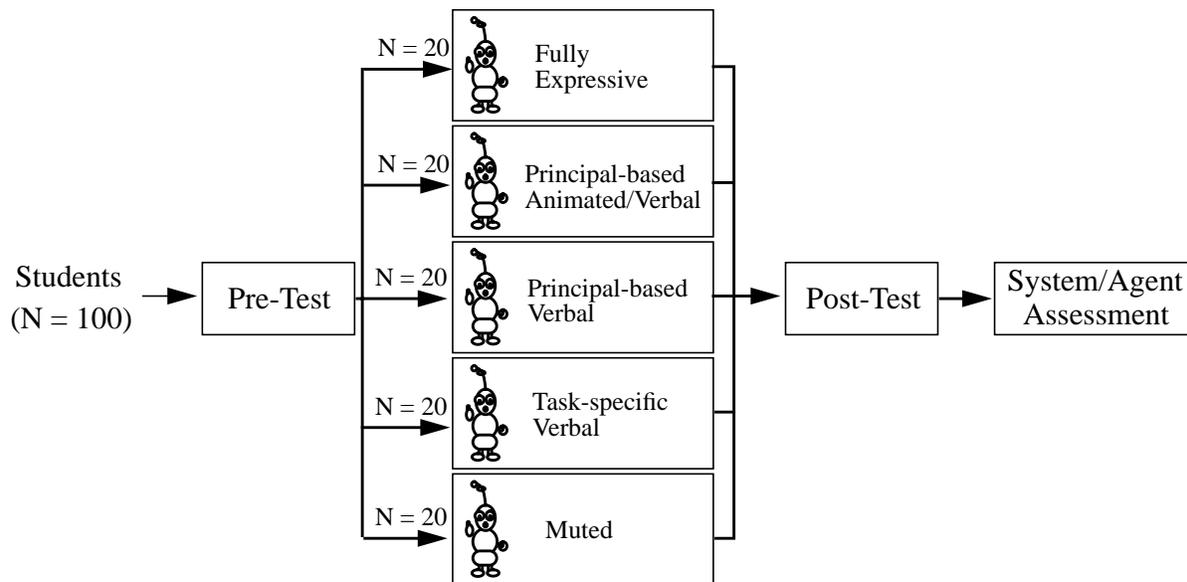


Figure 2: Evaluation Procedure

the study.

DATA ANALYSIS AND RESULTS

Three issues were evaluated in the analyses: potential differences between students' prior knowledge of the domain between clones types; overall affective impact of the agents on students' perception of their learning experiences; and effect of clone type on different affective dimensions.

The results of a one-way between-subjects Analysis of Variance (ANOVA) in which pre-test scores were analyzed indicated no significant differences between groups, [$F(4, 96) = 1.40, p > .05$].

To determine the pedagogical effectiveness of the agents, pre- and post-test scores for each clone type condition were analyzed in an ANOVA identical to that described above. Tukey's-Hsd Post hoc analyses were performed on all significant effects. Study results indicated that post-test scores were significantly higher than the pre-test scores [$F(1, 99) = 2.88, p < .05$], and that clone type affected the magnitude of the difference scores [$F(4, 96) = 3.01, p < .05$]. The greatest magnitude was obtained for students in the Fully Expressive, Principle-Based Animated/Verbal, and Principle-Based Verbal clone conditions. There were no other significant differences found.

Subjective assessment data for each combination of question and clone type were submitted to a mixed-model ANOVA in which question was a repeated measure while clone type was a between-participant measure. Results indicated a significant difference between question ratings [$F(17, 99) = 5.85, p < .05$]. Mean ratings for questions ranged from 3.0 to 4.6, with higher ratings indicating more acceptable performance (Table 1). The question receiving the best overall rating asked if the agent provided help when students committed an error. Other questions receiving relatively high ratings concerned: how believable was the advice the agent provided;

relationship between the agent's feedback and the student's progress; utility of the agent's advice; and if the student would like the agent to help with homework. The lowest ratings were given for questions asking if the agent knew more about plants than science teachers, the helpfulness of the agent in comparison to a science teacher, and if the agent became more helpful as the student progressed through the design exercises. A significant interaction was obtained for question and clone type [$F(68, 99) = 2.29, p < .05$]. For the Fully Expressive clone, students not only gave high ratings to those dimensions mentioned above, but also rated questions concerning the utility of the program's advice and the utility of the agent's encouragement particularly high. There were no differences between ratings for questions for the remaining clone types, and there was no effect of clone type on the means of the lowest rated questions.

DISCUSSION AND DESIGN IMPLICATIONS

The findings of the study are encouraging. First, they establish the critical "baseline" behavior that one would expect of animated pedagogical agents, namely, that students exhibit performance gains after interacting with the agent and the learning environment. In the absence of this effect, any results concerning the agent's positive impact on affective factors would be of limited use. While we hypothesize that the difference in the magnitude of pre- and post test scores is due to the students' interaction with the agent, an alternative explanation of the difference scores effect is that they are due to time-on-task or practice. However, in light of the fact that there were no differences in the pre-test scores between clone types, the significant effect of clone type on the magnitude of the post-test difference scores suggests that with regard to increasing learners' performance, the Task-Specific Verbal and the Mute clones were not as effective as the remaining clone types. This finding supports the notion that factors other than time-on-task influenced the difference between the pre- and post-test scores. During data collection, numerous types of

1. How entertaining was the educational program?	10. Would you like Herman the Bug to help you with your homework?
2. Regardless of how difficult it was to design the appropriate plants, how did you find using the educational program?	11. Do you think Herman the Bug knows more about plants than science teachers?
3. How well do you think the feedback from the educational program adapted to you?	12. When you made a mistake, did Herman the Bug become irritated?
4. Was the program's advice useful to you?	13. When you made a mistake, did Herman the Bug become concerned?
5. Did Herman the Bug encourage you?	14. Was Herman the Bug's advice useful to you?
6. Did you believe the advice you got from Herman the Bug?	15. As you progressed in the educational program, did Herman the Bug become more helpful?
7. Did Herman the Bug help you when you were having difficulties designing plants?	16. How was the feedback from Herman the Bug with respect to your progress?
8. In comparison to a science teacher, how helpful was Herman the Bug?	17. How entertaining or boring was working with Herman the Bug?
9. How clear was Herman the Bug about what he was saying?	18. How talkative was Herman the Bug?

Figure 3: System/Agent Assessment Form

data were collected, including error rates and task performance times for each combination of trial block and clone type. These data are currently being analyzed and should allow us to differentiate differences in pre- and post-test scores that are due to time on task and due to effectiveness of the pedagogical agent. While the current study focuses on the comparison of agents with different characteristics (including a muted agent), an interesting follow-on study involves the evaluation of an "agent-less" learning environment. The follow-on study will also provide important insights about the contributions of agents.

Second, the very presence of an animated agent 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. We refer to this as the *persona effect*. While this study was not designed to discover the specific way in which the animated pedagogical agent enhances learning, it is interesting to speculate how this may occur. There are two potential effects of agents on learning. First, there may be a direct cognitive effect in superior knowledge acquisition. This is consistent with the self-explanation effect [5]. Because agents can more actively engage students in learning, agents may well stimulate reflection and self-explanation. Second, there may be a motivation effect, which may be even more pronounced. Because lifelike characters have such an enchanting presence, they may significantly increase students' positive perceptions of their learning experiences. We hypothesize that lifelike characters create enthusiastic reactions in large part because of their believability [2] and human's innate responses to psycho-social stimuli. Clearly, studies that manipulate motivation levels (via pay-

offs or instructions) are needed to provide further data to distinguish between direct cognitive effects and the effects of increased motivation on learning performance.

Third, the significantly higher mean ratings for the Fully Expressive agent relative to the other agents suggests that multiple types of advisory behaviors may interact to improve positive affective impact. This *agent expressivity corollary* to the persona effect suggests that, in addition to the potential benefits in learning effectiveness that more expressive agents provide, their perception by students is also more positive. This is corollary is evidenced by higher means on many (but not all) of the dimensions studied. A possible explanation for the finding is that the combination of the media with which the advice was delivered and the presentation of two types of advice interacted synergistically.

The strength of the persona effect was evidenced by its impact with all of the clones. Students' perception of the agent's concern for them, the high degree of credibility they ascribed to it, and their perception of its utility and entertainment value all point toward the powerful influence of the persona effect. Even students interacting with the muted clone, whose advisory behaviors were non-existent, perceived the agent in a very positive light. This finding is indicative of a fundamental benefit provided by animated pedagogical agents, perhaps even if they are not optimally designed.

The positive ratings for the agents were not likely to have been caused by response bias due to the three precautions that were taken to reduce response bias: (1) anonymity of participants was preserved, and they were informed of this;

#	Dimension	Mean	Clone				
			Expressive	Principle-Based Verbal	Task-Specific Verbal	Muted	Principle-Based Anim/Verbal
1	Entertainment value (Program)	4.2	4.3	4.1	4.2	4.2	4.0
2	Ease of use (Program)	3.8	3.3	3.8	4.0	4.0	3.8
3	Adaptation of program's feedback	4.2	4.2	4.2	4.2	4.1	4.2
4	Utility of program's advice	4.2	4.6	4.1	4.1	4.2	4.1
5	Herman's encouragement	3.9	4.3	3.7	3.6	3.8	3.9
6	Believability of Herman's advice	4.5	4.8	4.6	3.9	4.4	4.7
7	Herman's helpfulness with errors	4.6	4.8	4.7	4.4	4.6	4.6
8	Helpfulness compared to science teacher	3.5	3.8	3.7	3.5	3.3	3.1
9	Clarity of Herman's statements	4.2	4.4	4.3	3.8	4.1	4.2
10	Desire for Herman's help with homework	4.3	4.5	4.3	4.0	4.3	4.3
11	Herman's knowledge compared to science teacher	3.4	3.6	3.2	3.6	3.1	3.4
12	Herman's irritation after mistakes	3.0	3.2	2.9	3.1	3.6	2.0
13	Herman's concern when errors occur	3.8	4.2	4.0	3.9	3.4	3.4
14	Utility of Herman's advice	4.4	4.5	4.4	4.1	4.5	4.5
15	Herman's helpfulness over time	3.5	3.6	3.4	3.5	3.5	3.5
16	Appropriateness of Herman's feedback	4.4	4.4	4.5	4.2	4.6	4.3
17	Entertainment value of Herman	4.0	4.3	4.1	3.8	4.1	3.9
18	Talkativeness of Herman	4.1	4.1	4.1	4.2	3.9	4.2

Table 1: Means for system/agent assessment form

(2) subjects were strongly encouraged to provide honest opinions in order to improve the software; and (3) privacy was granted during data collection. In addition, the significant effect for question number and for the question and clone type interaction cannot be explained by response bias.

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. The study examined only one age group in one domain. Second, the long term effects of interacting with agents were not explored in this study. Third, the potential negative impact of agents on users is well known. Agents that behave too proactively quickly become intrusive and irritating. Bearing these considerations in mind, three recommendations are suggested:

1. *Designers of interactive learning environments should give serious consideration to including an animated pedagogical agent in their software.* Our study provides significant evidence that the presence of animated pedagogical agent has a strong, positive impact on students' perception of their learning experiences. These benefits should be weighed against the possible disadvantages such as student distraction and increased development costs.
2. *Even if the designers of interactive learning environments cannot or do not wish to create an agent that provides advice proactively, they should still give serious consideration to including an animated pedagogical agent in their software.* Because of the finding that even muted agents are perceived in a very positive light, providing students with an agent—even an “adviceless” agent—may prove

worthwhile.

3. *Designers of interactive learning environments that can create an expressive agent for their domain should give particularly serious consideration to including such an agent in their software.* Because of the combination of the pedagogical benefits and the positive affective impact of expressive animated pedagogical agents, introducing an expressive agent into a learning environment warrants serious consideration.

CONCLUSION

As a result of rapid advances in animated agent technology, the prospect of deploying animated pedagogical agents on a broad scale is quickly becoming a reality. Because these agents can provide students with customized advice in response to their problem-solving activities, their potential to increase learning effectiveness is significant. In addition, however, these agents can also play a critical motivational role as they interact with students. As a result, students may choose to use interactive learning environments frequently and for longer periods of time.

To investigate the affective impact of animated pedagogical agents on students' perception of their learning experiences, we undertook an empirical study with 100 middle school students. The study revealed that well crafted lifelike agents have an exceptionally positive impact on students. Students perceived the agents as being very helpful, credible, and entertaining. This persona effect held strong even for an agent whose communicative behaviors were muted. The study also found that combinations of types of advice can (1) increase students' positive perception of the agent and (2) increase learning performance.

This work represents a promising first step toward developing an understanding of the impact that animated pedagogical agents can have on children's learning experiences. Perhaps the greatest challenge lies in determining precisely which characteristics of these agents are most effective for particular age groups, domains, and learning contexts. We will be investigating these factors in our future research.

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. Thanks also to the anonymous reviewers for their insightful suggestions.

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. Bruce Blumberg and Tinsley Galyean. Multi-level direction of autonomous creatures for real-time virtual environments. In *Computer Graphics Proceedings*, pages 47–54, 1995.
4. E. Cauzinille-Marmeche and J. Mathieu. Experimental data for the design of a microworld-based system for algebra. In Heinz Mandl and Alan Lesgold, editors, *Learning Issues for Intelligent Tutoring Systems*, pages 278–286. Springer-Verlag, New York, 1988.
5. M. Chi, M. Bassok, M. Lewis, P. Reimann, and R. Glaser. Self-explanations: How students study and use examples in learning to solve problems. *Cognitive Science*, 13:145–1182, 1989.
6. Christopher R. Elliot and Beverly P. Woolf. A simulation-based tutor that reasons about multiple agents. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, pages 409–415, 1996.
7. 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.
8. James D. Hollan, Edwin L. Hutchins, and Louis M. Weitzman. STEAMER: An interactive, inspectable, simulation-based training system. In Greg Kearsley, editor, *Artificial Intelligence and Instruction: Applications and Methods*, pages 113–134. Addison-Wesley, Reading, MA, 1987.
9. David Kurlander and Daniel T. Ling. Planning-based control of interface animation. In *Proceedings of CHI '95*, pages 472–479, 1995.
10. Robert W. Lawler and Gretchen P. Lawler. Computer microworlds and reading: An analysis for their systematic application. In Robert W. Lawler and Masoud Yazdani, editors, *Artificial Intelligence and Education*, volume 1, pages 95–115. Ablex, Norwood, NJ, 1987.
11. 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*, 1997.
12. James C. Lester and Brian A. Stone. Increasing believability in animated pedagogical agents. In *Proceedings of the First International Conference on Autonomous Agents*, 1997.
13. 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.
14. P. Maes, T. Darrell, B. Blumberg, and A. Pentland. The ALIVE system: Full-body interaction with autonomous agents. In *Proceedings of the Computer Animation '95 Conference*, 1995.
15. C. Nass, J. Steuer, L. Henriksen, and H. Reeder. Anthropomorphism, agency and ethopoeia: Computers as social actors. In *Proceedings of the International CHI Conference*, 1993.
16. 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.
17. Byron Reeves and Clifford Nass. Information characteristics of media technologies that give a sense of "being there". In *Annual Meeting of the International Communication Association*, 1992.
18. Jeff Rickel and Lewis Johnson. Integrating pedagogical capabilities in a virtual environment agent. In *Proceedings of the First International Conference on Autonomous Agents*, 1997.
19. 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.
20. Patrick W. Thompson. Mathematical microworlds and intelligent computer-assisted instruction. In Greg Kearsley, editor, *Artificial Intelligence and Instruction: Applications and Methods*, pages 83–109. Addison-Wesley, Reading, MA, 1987.
21. Xiaoyuan Tu and Demetri Terzopoulos. Artificial fishes: Physics, locomotion, perception, and behavior. In *Computer Graphics Proceedings*, pages 43–50, 1994.