

# Increasing Believability in Animated Pedagogical Agents\*

James C. Lester and Brian A. Stone

Multimedia Laboratory  
Department of Computer Science  
North Carolina State University  
Raleigh, NC 27695-8206  
{lester,bastone}@eos.ncsu.edu

## Abstract

Animated pedagogical agents offer great promise for knowledge-based learning environments. In addition to coupling feedback capabilities with a strong visual presence, these agents play a critical role in motivating students. The extent to which they exhibit life-like behaviors strongly increases their motivational impact, but these behaviors must always complement and never interfere with students' problem solving.

To address this problem, we have developed a framework for dynamically sequencing animated pedagogical agents' believability-enhancing behaviors. By monitoring a student's problem-solving history and the agent's past activities, a *competition-based* behavior sequencing engine produces realtime life-like character animations that are pedagogically appropriate. Behaviors in the agent's repertoire compete with one another. At each moment, the strongest eligible behavior is heuristically selected as the winner and is exhibited. We have implemented this framework in Herman the Bug, an animated pedagogical agent that inhabits a knowledge-based learning environment for the domain of botanical anatomy and physiology.

## Introduction

Dynamically animated agents offer great promise for knowledge-based learning environments. Because of the immediate and deep affinity that children seem to develop for these interactive life-like characters, the direct pedagogical benefits these agents provide are per-

haps exceeded by their motivational benefits. By creating the illusion of life, dynamically animated agents have the potential to significantly increase the time that children seek to spend with educational software, and recent advances in affordable graphics hardware are beginning to make the widespread distribution of realtime animation technology a reality.

Endowing animated agents with believable, life-like qualities has been the subject of much recent research (Bates 1994; Tu & Terzopoulos 1994; Granieri *et al.* 1995; Blumberg & Galyean 1995; Kurlander & Ling 1995; Maes *et al.* 1995). Believability is a key feature of animated agents for learning environments, and experiences with students interacting with an *animated pedagogical agent* (Figure 1) developed in our laboratory (Stone & Lester 1996) have led us to conclude that increasing believability will yield significant rewards in student's motivation as they interact with learning environments.

However, achieving believability in animated pedagogical agents poses three major challenges. First, the primary goal of pedagogical agents is to promote learning, and any agent behaviors that would interfere with students' problem-solving—no matter how much these behaviors might contribute to believability—would be inappropriate. For example, if the agent were to cartwheel across the screen when the student was grappling with a difficult problem, the student's concentration would be immediately broken. Second, believability-enhancing behaviors must complement (and somehow be dynamically interleaved with) the advisory and explanatory behaviors that pedagogical agents perform. Third, if observers see that an agent is acting like a simple automaton, believability is either substantially diminished or eliminated altogether.

To address these issues, we have developed a *competition-based* framework for dynamically sequencing animated pedagogical agents' *believability-enhancing* behaviors. Throughout learning sessions, an animated agent's pedagogical behaviors are sequenced by a pedagogical sequencing engine (Stone & Lester 1996). The pedagogical sequencing engine

---

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



Figure 1: An animated pedagogical agent: Herman the Bug

is complemented by a believability-enhancing behavior sequencing engine that enables the agent to perform a large repertoire of actions such as visual focusing (e.g., motion-attracted head movements), re-orientation (e.g., standing up, lying down), locomotion (e.g., walking across the scene), body movements (e.g., back scratching, head scratching), restlessness (e.g., toe tapping, body shifting), and prop-based movements (e.g., glasses cleaning). Believability-enhancing behaviors compete with one another for the right to be exhibited. At each moment, the strongest eligible behavior is heuristically selected as the winner and is exhibited. The net result of the ongoing competition is that the agent behaves in a manner that significantly increases its believability without sacrificing pedagogical effectiveness. This framework has been used to implement a realtime competition-based believability-enhancing behaviors for an animated pedagogical agent.

### Believability in Pedagogical Agents

We define *believability* as the extent to which users interacting with an agent come to believe that they are observing a sentient being with its own beliefs, desires, and personality. Recent years have witnessed a surge of interest in believability, including investigations into the role of emotion in animated agents (Bates 1994), methods for achieving realistic movement for both individuals and groups of autonomous artificial animals (Tu & Terzopoulos 1994), techniques

for providing full-body user-agent interaction (Maes *et al.* 1995), and methods for controlling the behavior of simulated human agents that are rendered in near real-time (Granieri *et al.* 1995; Webber *et al.* 1995). Work has even begun on tools for facilitating the integration of interactive animations into user interface management systems (Kurlander & Ling 1995) and for providing user-directed control of semi-autonomous agents (Blumberg & Galyean 1995).

One of the most promising arenas for believability is education. Introduced immersively into a learning environment, an animated pedagogical agent can observe students' progress and provide them with visually contextualized problem-solving advice. In addition to improving students' problem solving, however, animated pedagogical agents that are *believable* can play an important role in increasing students' motivation. During the past year, we have conducted two observational studies of more than twenty middle school students interacting with "Herman the Bug," the pedagogical agent inhabiting the DESIGN-A-PLANT (Lester *et al.* 1996) learning environment.<sup>1</sup> Because these studies suggest that students' interest is greatly increased by

<sup>1</sup>DESIGN-A-PLANT is a knowledge-based learning environment for the domain of botanical anatomy and physiology. Given a set of environmental conditions, children graphically assemble customized plants that can thrive in the specified environments. At the time of these studies, the agent's repertoire included 30 animated behaviors and 160 verbal behaviors which were designed, modeled, and rendered by a team of 12 graphic artists and animators on SGIs and Macintoshes.

an agent’s life-like presence, we undertook a concerted effort to develop techniques for improving believability.

Believability in animated agents is a product of two forces: the visual qualities of the agent, and the computational properties of the sequencing engine that schedules its behaviors in response to evolving interactions with the user. The behavior cannon of the animated film (Jones 1989; Lenburg 1993) has much to say about aesthetics, movement, and character development, and the pedagogical goals of learning environments impose additional requirements on character behaviors. In particular, techniques for increasing the believability of animated pedagogical agents should satisfy the following criteria:

- *Situated Liveness*: Throughout problem-solving sessions, agents should remain “alive” in a situated manner (Suchman 1987) by continuing to exhibit behaviors that indicate their alertness, e.g., through visually tracking students’ activities, and by providing anticipatory cues (Thomas & Johnston 1981) to signal their upcoming actions.
- *Controlled Visual Impact*: Some behaviors such as moving from one location to another have high visual impact, while others, such as small head movements, have low visual impact. In general, the higher the visual impact, the more interesting a behavior will be, but agents must control the visual impact of their behaviors in such a manner that they do not divert students’ attention at critical junctures.
- *Complex Behavior Patterns*: Because students will interact with animated pedagogical agents over extended periods of time, it is critical that agents’ behavior patterns be sufficiently complex that they cannot be quickly induced. Easily recognized behavior patterns significantly reduce believability.

### Dynamically Sequencing Believability-Enhancing Behaviors

The believability requirements call for an approach to animated pedagogical agents that enables them to “stay alive” by exhibiting complex patterns of behaviors which are pedagogically appropriate. In contrast to the more “planful” advisory activities of pedagogical agents, believability-enhancing behaviors are performed to satisfy the situated liveness criterion. To this end, we have developed a framework for animating pedagogical agents that treats liveness as an emergent property arising from the continuous exhibition of believability-enhancing behaviors. In this framework, the agent’s believability-enhancing behaviors are dynamically sequenced by a *competition-based* sequencing engine (Figure 2).

Believability-enhancing behaviors compete with one another for the right to be exhibited. Associated with each behavior is an *exhibition strength*, a measure of the behavior’s current degree of appropriateness for presentation in the current context. Behaviors earn the

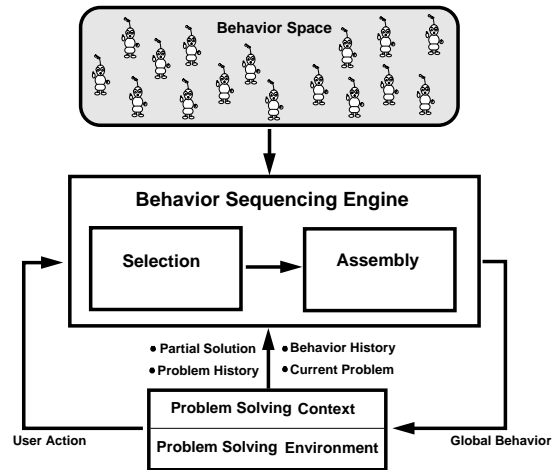


Figure 2: Behavior Sequencing Engine

right to be exhibited by increasing their strengths. At each moment in the competition, their *growth rate* may be either increased or decreased by the behavior scheduler. Behaviors that are appropriate for the student’s problem-solving activities and the agent’s behavior history are dynamically rewarded with higher growth rates. Because their strength increases more quickly, they win the competition more frequently. Less desirable behaviors are punished by being assigned lower growth rates, and they are therefore exhibited less frequently.

The realtime behavior scheduler oversees the competition in the following way: At every tick of the clock, it first rewards or punishes each behavior according to its appropriateness for the current context, and then it assesses the relative strengths of all of the behaviors. It next determines which behaviors are eligible for exhibition by evaluating their exhibition constraints with respect to the student’s problem-solving history and the agent’s behavior history. The strongest eligible behavior is permitted to exhibit itself, it is marked as exhibited, its strength is zeroed out, and the competition continues. Each round of the competition proceeds according to the following algorithm.

Let  $B$  be the set of all believability-enhancing behaviors. After zeroing out the strength  $s_i$  of each behavior  $b_i \in B$ , and after initializing all growth rates  $g_i$  to a non-zero value, the competition is conducted as follows:

1. **Compute behavior growth rates.** Determine which *growth rate effectors* are active by pattern matching their preconditions against representations of the student’s problem-solving history. For each behavior  $b_i \in B$ , execute the active growth rate effectors to compute a new  $g_i$  for  $b_i$ .
2. **Evaluate behavior strengths.** For each behavior  $b_i \in B$ , update  $b_i$ ’s *strength* by increasing  $s_i$  at the rate of  $g_i$ .

3. **Identify exhibitable behaviors.** For each behavior  $b_i \in B$ , evaluate  $b_i$ 's *exhibition constraints*. If  $b_i$ 's exhibition constraints evaluate to *true*, add  $b_i$  to  $E$ , the set of exhibitable behaviors.
4. **Determine strongest exhibitable behaviors.** Compute  $S \subseteq E$ , the set of exhibitable behaviors that are the strongest.  $S$  need not be a singleton because ties may occur.
5. **Select and display winning behavior.** Select  $b_S \in S$  such that  $b_S$  occupies the highest stratum in the *visual impact spectrum* of all behaviors in  $S$ . If more than one behavior satisfies this requirement, randomly select one. Mark  $b_S$  as displayed, and insert it in the exhibition queue.

It is important to note two aspects of the global behaviors produced by the competitive scheduling algorithm. First, even though behaviors are continuously being scheduled, the agent does not act frenetically. This is accomplished by including the null behavior,  $b_\epsilon$ , in the pool of behaviors. Because  $b_\epsilon$  occupies the lowest position in the visual impact spectrum, and because it can be repeated frequently with relatively little penalty, the agent can remain in the final pose of the preceding animation each time  $b_\epsilon$  is scheduled. Exploiting  $b_\epsilon$  permits a uniform treatment of all behaviors. Second, the competition is interrupted by three types of events, any of which may change the agent's behavior: (1) the student may request advice; (2) the student may perform a problem-solving action; or (3) the student's problem-solving idle time may exceed the allotted interval. In each case, the pedagogical behavior scheduler (Stone & Lester 1996) will take an appropriate action, update the agent's state, update the problem-solving history, and return control to the believability sequencing engine.

This iterative process of rewarded growth and assessment is conducted continuously throughout problem-solving sessions. Behaviors that are appropriate experience accelerated growth, and each round of the competition produces a new winner which is permitted to exhibit itself (Figure 3). The process plays out in real-time, with re-evaluations of each behavior's strength occurring many times each second. The net result of the ongoing competition is that the agent remains alive throughout the interaction, its visual impact is pedagogically appropriate, and its patterns of behavior are complex.

Applying this framework to create a believable pedagogical agent consists of three steps: populating a behavior space with believability-enhancing behaviors and assigning them to strata of a visual impact spectrum; encoding exhibition constraints for each behavior; and representing growth rate effectors.

### Visual Impact Spectrum

The first step in creating a believable pedagogical agent is specifying and rendering believability-enhancing be-

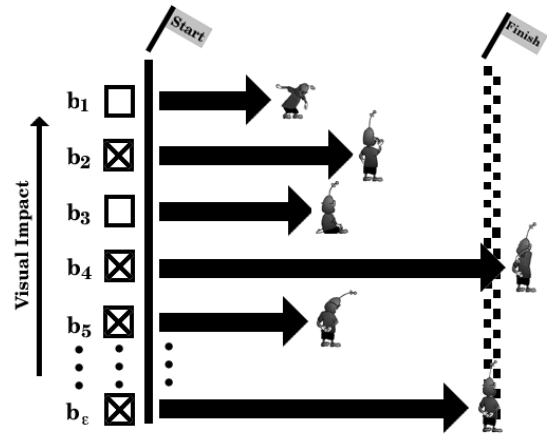


Figure 3: Competition-Based Behavior Sequencing

haviors, a collaborative process between the agent's designers and its animators. The behaviors must be kept within the strict confines of the initial character development, as defined in the character's *model sheets* (Culhane 1988). The behaviors produced in this process should cover a broad range of activities that vary considerably in visual impact. After the behaviors have been rendered, they are assigned to a stratum of the visual impact spectrum. This is accomplished by partitioning them into behavior strata  $V_1 \dots V_k$ , where behaviors in a given stratum have the same visual impact. Behaviors that exercise only a small number of body parts in a minor way typically have reduced visual impact; behaviors that signal upcoming actions have more impact; behaviors that manipulate props (e.g., glasses to be cleaned), or produce changes in orientation have even greater impact; and behaviors that involve eye-catching changes in location have the greatest impact. For example, the visual impact spectrum for Herman's believability-enhancing behaviors (shown in increasing order) is:

( $\{b_\epsilon\}$ , Visual-Focusing, Fidgeting, Anticipatory-Fidgeting, Orientational-Modifications, Prop-Based-Fidgeting, Locomotion)

By indexing behaviors by their family, the sequencing engine can efficiently locate behaviors of a given  $V_i$ .

### Exhibition Constraints

Believability-enhancing behaviors are annotated with representations of the pedagogical and physical contexts in which they may be performed. During the competition, these exhibition constraints are evaluated against (1) the student's problem-solving history, (2) the current state of the learning environment, (3) and the agent's current state. The representation of the student's problem-solving history includes: a record of the previous problems they have solved; their solution paths; and the partial solutions attempted to

date for the current problem.<sup>2</sup> The current state of the learning environment encodes: knowledge about which sub-problem is the student's current focus; a marker indicating which sub-problem the system believes is most critical; the permissible idle time, i.e., how much time the student can remain idle before the system should interject assistance; and the amount of permissible idle time that has already expired. Finally, the agent's current state is represented with a pair  $(P, O)$ , where  $P$  symbolically represents the agent's screen position and  $O$  represents its orientation. For example, the state of the agent in Figure 1 is `(mid-bar-left, lying)`. Exhibition constraints encode knowledge about conditions that must be preserved with respect to:

- *Pedagogical contexts*: For example, Herman has a constraint which preserves the condition that behaviors occupying upper strata of the visual impact spectrum cannot be performed if the idle time expired is less than one third of the maximum idle time. If he performed high visual impact behaviors when the student is most actively engaged in problem-solving, he would adversely affect learning.
- *Physical contexts*: For example, Herman has a constraint which enforces orientation changes. If his orientation is standing in one state, he cannot perform the lying down behavior; rather, he must first sit down before lying down.

At each iteration of the competition, the conjunction of exhibition constraints on a given behavior  $b_i$  is evaluated. If this evaluates to *true*, as indicated by a checked box in Figure 3, then  $b_i$  is a candidate for exhibition.

## Growth Rate Effectors

While exhibition constraints enforce conditions that can never be violated, growth rate effectors enact graduated preferences. They are defined by pre-conditions, which are predicates on problem-solving contexts, and actions, which reward behaviors that are appropriate for the problem-solving context and punish behaviors that are less appropriate. Through their evaluation of problem-solving progress, growth rate effectors contextualize three aspects of the agent's behavior:

- *Visual impact enhancement*: Behaviors inhabiting upper strata of the spectrum are rewarded when the student is addressing less critical sub-problems, e.g., when the student is solving an unimportant sub-problem, Herman is more likely to perform an interesting prop-based behavior such as cleaning his glasses or a locomotive behavior such as jumping across the screen.
- *Anticipatory signaling*: When the student's permissible idle time has almost expired, anticipatory behaviors are rewarded to indicate upcoming actions, e.g., as the student begins to stall on a problem and

<sup>2</sup>Representational details of students' problem-solving contexts may be found in (Lester *et al.* 1996).

Herman is on the verge of interjecting advice, he is likely to sit up if he is lying down.

- *Visual focusing activation*: Visual focusing behaviors are rewarded when the student attends to particular objects in the learning environment for extended periods of time, e.g., when the student lingers on a particular component (as indicated by mouse location), Herman is likely to look at that component.<sup>3</sup>

## Implementation: A Believable Pedagogical Agent

The competition-based sequencing framework for believability-enhancing behaviors has been implemented in Herman (Figure 1), the animated pedagogical agent for the DESIGN-A-PLANT learning environment. Herman is a talkative, quirky insect with a propensity to fly about the screen and dive into the plant's structures as he provides problem-solving advice to students. 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. These behaviors are all sequenced by the pedagogical behavior sequencing engine (Stone & Lester 1996). To increase Herman's believability, we implemented a competition-based sequencing engine and created a library of fifteen believability-enhancing behaviors that spanning the visual impact spectrum.<sup>4</sup>

Running on a Power Macintosh 9500/132, the competition-based sequencing engine computes a new growth rate for each behavior approximately once every 200 milliseconds. There are sometimes 2-3 winning behaviors in the queue at a given time. For example, Herman may be in the process of jumping to the opposite side of the screen, where he must then sit down, and finally turn his head to focus on a particular plant component. To ensure that the queued behaviors are exhibited seamlessly, new behaviors must begin with the same frame that terminated the previous behavior. This is accomplished by establishing a set of *home* frames indexed by the agent's position and orientation. In addition, some behaviors should return to the same frame that they began with, and some behaviors should be repeatable without constantly having to return to the home frame. The implementation therefore supports exhibiting behaviors forward and in reverse, and for looping frames within a behavior.

To illustrate the behavior of the competition-based sequencing engine, suppose a student is interacting with DESIGN-A-PLANT to construct the roots, stems,

<sup>3</sup>Our observational studies suggest that mouse position furnishes an accurate indicator of students' focus of attention in learning environments where rollover text plays an important role in the interface.

<sup>4</sup>Believability-enhancing behaviors range in duration from 8 to 30 frames. Because they are exhibited at 8-10 frames per second, each behavior lasts from 1-3 seconds.

and leaves of a plant that will survive in the desert. After pointing to the environmental icons at the top of the screen while verbally introducing the environment, Herman walks to initial home frame (standing on left of screen). The student considers what root she should choose, while the idle time slowly mounts. Within ten seconds Herman reaches up and scratches his head. Soon after, he twiddles his thumbs, then points at the empty root component and says “Don’t worry about which root to choose in this environment. Just about any root will work.” The student selects a root and Herman says “Good Job!” The student chooses to work on leaves next. Because leaves are the most important sub-problem for this environment, Herman sits down and remains still for several seconds.

Visual focusing behaviors are rewarded more than fidgets during periods of intense problem solving, so Herman remains seated and looks up at the leaf the student is considering. As the student spends an inordinate amount of time on the leaf with no progress, Herman eventually starts tapping his foot while remaining seated. When the student selects an incorrect leaf, Herman points at it while staying seated, and the pedagogical sequencing engine selects animated advice with a verbal introduction in which Herman says, “Only certain types of leaves will work in such a dim environment. Let me show you why that’s so.” He then lies down and an animated movie of Herman explaining the relationship between low sunlight and leaf choice is shown on a pop-up movie screen. When the movie is over, Herman sits up. Herman continues to exhibit appropriate believability-enhancing behaviors until the student has designed a completely correct plant.

## Conclusion

Because of their strong life-like presence, believable animated pedagogical agents can capture students’ imaginations and play a critical motivational role in keeping them deeply engaged in a learning environment’s activities. To dynamically sequence believability-enhancing behaviors, we have designed the competition-based behavior sequencing framework and used it to implement an animated pedagogical agent. It exploits the visual impact spectrum to select behaviors whose visual impact is appropriate for students’ concentration as they solve problems; it evaluates the exhibition constraints at each cycle of the competition to ensure that only behaviors that are appropriate for the current pedagogical and physical contexts can be performed; and it cyclically executes the growth rate effectors to reward desirable behaviors and punish less desirable behaviors.

The net result of the ongoing competition between behaviors is that the agent behaves in a surprisingly life-like manner that is pedagogically appropriate and complex. Although much work remains to be done on increasing the agent’s flexibility by reducing behavior granularity and enlarging its behavioral repertoire, it appears that competition-based sequencing constitutes

an effective approach to increasing the believability of animated pedagogical agents. To empirically test this hypothesis, we are currently embarking on a large-scale formal study of student-agent interaction.

## Acknowledgements

We would like to thank: the animation team, which was lead by Patrick FitzGerald of the North Carolina State University School of Design; John Myrick, the agent’s lead animator; Chris Tomasson and the students in her seventh grade class at Martin Middle School and Paula Sloan and the students of the Raleigh Chapter of the Women in Science Mentoring Program for their participation in the evaluations; and Jeff Rickel for insightful discussions on believability.

## References

- 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.
- Culhane, S. 1988. *Animation from Script to Screen*. New York: St. Martin’s Press.
- 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.
- Jones, C. 1989. *Chuck Amuck: The Life and Times of an Animated Cartoonist*. New York: Avon.
- Kurlander, D., and Ling, D. T. 1995. Planning-based control of interface animation. In *Proceedings of CHI ’95*, 472–479.
- Lenburg, J. 1993. *The Great Cartoon Directors*. New York: Da Capo Press.
- Lester, J.; Stone, B.; O’Leary, M.; and Stevenson, R. 1996. Focusing problem solving in design-centered learning environments. In *Proceedings of the Third International Conference on Intelligent Tutoring Systems*, 475–483.
- 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*.
- 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.
- Thomas, F., and Johnston, O. 1981. *The Illusion of Life: Disney Animation*. New York: Walt Disney Productions.
- Tu, X., and Terzopoulos, D. 1994. Artificial fishes: Physics, locomotion, perception, and behavior. In *Computer Graphics Proceedings*, 43–50.
- Webber, B.; Badler, N.; Eugenio, B. D.; Geib, C.; Levison, L.; and Moore, M. 1995. Instructions, intentions and expectations. *Artificial Intelligence* 73:253–269.