

The Virtual Reality Gorilla Exhibit

Don Allison, Brian Wills, Doug Bowman,
Jean Wineman, and Larry F. Hodges
Georgia Institute of Technology

The Virtual Reality Gorilla Exhibit teaches users about gorilla behaviors and social interactions. We present techniques for building the environment and the virtual gorillas that inhabit it.

The Virtual Reality Gorilla Exhibit, an immersive virtual environment, lets students assume the persona of an adolescent gorilla, enter into one of the gorilla habitats at Zoo Atlanta in Atlanta, Georgia, and interact as part of a gorilla family unit. The exhibit combines a model of Zoo Atlanta's Gorilla Habitat 3—home of Willie B, a 439-pound male silver-back gorilla and his family group—with computer-generated gorillas whose movements and interactions are modeled as accurate representations of gorilla behaviors (see Figure 1). The VR Gorilla Exhibit's goal is to create an experiential educational tool for middle school students to learn about gorillas' interactions, vocalizations, social structures, and habitat.

Motivation

Gorillas are an endangered species. Zoos around the world have expended time and money on public education about gorillas and their plight to raise public awareness and motivate people. People are being asked to help fund conservation efforts and apply political pressure to encourage the governments of the only three countries (Uganda, Rwanda, and Zaire) to which gorillas are native to actively prosecute poachers and promote conservation.

We felt that a well-designed virtual environment could contribute to these educational efforts, augmenting them in ways not possible through normal educational media. Also, since a VR setup is transportable while an actual gorilla family is not, a virtual gorilla environment could reach people who live too far from a zoo to visit and spend several hours watching the gorillas and reading each exhibit's informational signs.

Unfortunately, students must learn many aspects of gorilla life through reading, rather than through direct observation at a zoo exhibit. For example, introducing a new gorilla to a group is done off-exhibit, so students

rarely get the chance to observe the establishment or reinforcement of the dominance hierarchy, and challenges to it. Also, for the animals' own protection from disease and because of the logistics problems it would cause the keepers, visitors normally may not observe the night quarters. Nor do they get to watch the routine involved in letting the gorillas out in the morning and bringing them in at night. On exhibit, the distance separating the gorillas from the students makes it hard to observe such mood indicators as gorilla vocalizations or facial expressions. Finally, gorillas are active in early morning and late afternoon, and sleep in the middle of the day. Because of class scheduling logistics, most middle school students visit the zoo during the sleep period.

A virtual gorilla exhibit solves these logistical problems. It lets students observe a broader set of gorilla behaviors—time-shifting behaviors that they normally would not see—and lets them visit off-limits areas.

Pedagogically, constructivist theories of education advocate that the more viewpoints presented to students, the more they learn and the better they retain what they learn. With the VR gorilla exhibit, not only do students get to explore areas normally off limits to them, they also get to assume a gorilla identity and interact with other gorillas as a peer, thus gaining a different perspective on gorillas by experiencing gorilla life from a first person point of view. By interacting with other virtual gorillas, students learn through first-hand experience a gorilla group's social structure and accepted social interactions. This first-person interaction also tends to hold students' attention longer, allowing more information to be presented by the system and retained by the student.

Background

The popular press and educational and scientific communities have discussed the impact and appropriateness of VR for educational applications. However, few actual applications of VR to education exist, and the majority of those have focused more on adult task training (such as piloting a plane or driving a tank) than on general information acquisition for middle school students.

Similarly, the computer game market has helped spark interest in building autonomous creatures and interacting with them in a virtual environment. Much of the research, though, focuses only on small pieces of the problems encountered when building such creatures, instead of examining the process as a whole and trying to build complete, interesting, and autonomous characters.

VR and education

Wickens¹ summarized research by others from a theoretical perspective and argued that VR might make doing lessons easier while reducing retention. Damarin,² on the other hand, argued that VR helps students construct new knowledge expeditiously by letting them experience a subject from multiple viewpoints and through self-directed exploration.

At the implementation level, Brelsford³ compared a VR physics simulator, which implemented simple Newtonian mechanics, with lectures on the same material. For both junior high and college students, the groups that used the VR simulation showed higher retention than those receiving the lecture. The Virtual Reality Roving Vehicle project⁴ used VR to teach students about VR. Finally, the Narrative Immersive Constructionist/Collaborative Environments (NICE) project⁵ let students interact with each other in a virtual garden that they had to plant and tend, learning about gardening and working together in the process.


Interactive virtual animals

Much simulated animal work focuses on building a single type of behavior or on using certain programming techniques (such as neural nets or genetic programming) to construct specific behavior controllers. However, the virtual whale project at Simon Fraser University (<http://fas.sfu.ca/cs/research/projects/Whales/>) attempts to visualize data on whale foraging, with the long-term goal of letting users learn about whale life and foraging behavior by entering a virtual underwater environment.

The Artificial Life Interactive Video Environment (Alive) project⁶ is a complete system in which users interact with a virtual dog, playing catch with the dog and controlling it with gestures. Geared toward entertainment, this system serves as a testbed for gesture tracking and interaction research.

Tu and Terzopoulos⁷ built a physically based fish simulation, on which they have layered a vision system and behavior controllers. Although users can't interact with the fish, the fish interact with each other differently depending on whether they are predators, prey, or pacifists.

The Oz project built a system of creatures called Woggles—autonomous agents with interesting, emotional behaviors for users to interact with.⁸ These autonomous agents also served as the basis for a user-



1 The Virtual Willie B surveying a secluded area of his habitat.

directed improvisation system (targeted to children) where users specify possible scripts to control their characters' interactions.⁹ Both systems focus more on action selection and direction and less on the interface, which is still mouse-based with a third-person point of view.

Two recent commercial projects let users interact with complete virtual animals for entertainment purposes: Fujitsu's Fin Fin (<http://www.fujitsu-interactive.com/products/finfin.html>) and PF Magic's Computer Petz (<http://www.pfmagic.com/>). Neither is immersive.

Basic gorilla


One of our project goals is to accurately simulate gorilla behavior. While many information sources describe general primate behavior, only two major observational studies of gorilla behavior in the wild exist: that of George Schaller in the late fifties and early sixties, and that of Dian Fossey in the mid-sixties through the mid-eighties. Our main written reference on gorilla behaviors was Maple's,¹⁰ which summarized Schaller and Fossey's research and included information on the behavior of captive gorillas (see the sidebar "Gorillas in the Wild").

While books helped us find out what gorillas did, seeing them do it was necessary for accurately simulating

Gorillas in the Wild

At the gross level, not many differences exist between the lifestyles of gorillas in the wild and gorillas in a well designed zoo exhibit. In both cases, gorillas have well defined territorial boundaries that they normally don't cross. Within their territories they forage for food and rest in a diurnal cycle in both cases. Gorillas have no natural enemies other than man, so the lack of predators in captivity doesn't result in major behavioral changes. Changes in group composition still occur in captivity, but are facilitated by the exhibit curators. Thus, although the virtual gorilla environment is based on the habitats at Zoo Atlanta, the information learned generally applies to gorillas in the wild as well.

2 Interior of the virtual visitors' center looking out.



their behaviors. We shot several hours of video at Zoo Atlanta and studied additional footage provided by zoo researchers, including some behind-the-scenes footage of gorilla introductions. We even convinced the zoo scientists to act out some of the more difficult-to-visualize motions, to make sure we correctly implemented the details. These sources served as a basis for constructing the gorilla models and motions, which were then reviewed by the Zoo Atlanta gorilla experts and refined iteratively based on their comments.

The lifestyle of a captive gorilla housed in a natural habitat mirrors that of a gorilla in the wild with a few modifications (mainly imposed by the restricted size of the habitats). In the morning, the gorillas eat a breakfast of primate chow in the night quarters. Then they are let out into the habitats where they forage for “browse” (slices of fruit and vegetables scattered around by the keepers ahead of time). They spend the middle of the day resting and then have an afternoon snack of more fruits and vegetables. In the evening, they return to the night quarters building, where they build nests with vegetation provided. This mirrors a gorilla in the wild’s diurnal routine. Because of the limited size of the habitats and the difficulty of replacing 100-year-old oak trees and other vegetation, gorillas are discouraged from climbing trees or eating local vegetation. However, dead tree trunks and branches are provided for them to play with, eat, or destroy as they wish.

A gorilla family group centers around the dominant male silverback, so-named because the hair on his back is gray or silver instead of black. The group is generally composed of a single silverback, several females, and possibly some blackback males, juveniles, and infants. The silverback male is usually father to most of the infants and juveniles in the group.

Just as there is a pecking order among all the gorillas in the group, a pecking order also exists among the females, with the head female getting most of the silverback’s attention. Not as much attention is paid to rank among the juveniles and infants.

Mothers of infant gorillas tend to be very protective of

their young, carrying their infants or keeping them close at hand for about the first three years. As the infants grow into juveniles, they are allowed to wander farther from their mothers. They also have more interactions with their siblings and the other adults. While infants and juveniles can be quite playful—chasing each other, climbing trees (at least in the wild), and so on—as gorillas mature the play sessions become more infrequent and tree climbing becomes much rarer.

Gorillas use sounds, gestures, and motion to establish or reinforce position in the group’s hierarchy and to interact between groups. Gestures such as ground slapping, chest beating, or charging combined with vocalizations such as grunts or hoots

establish dominance, correct disobedient youngsters, or chase off another group from a group’s territory. Sentries use sounds to warn their group, and gorillas make sounds to express contentment or alert the other gorillas of their group where they’re located.

Building the initial system

Implementing the VR Gorilla Exhibit required constructing a gorilla habitat model and gorilla models that encapsulated gorilla geometry, movements, and vocalizations. Basic VR software support was available through Georgia Institute of Technology’s Simple Virtual Environment (SVE) toolkit (<http://www.cc.gatech.edu/gvu/virtual/SVE/>). SVE provides a set of software tools for common VR actions such as head tracking, model maintenance, and locomotion.

To model the gorilla habitats at Zoo Atlanta, we started by generating a 3D triangulated irregular network (TIN) mesh for the habitats and dividing moats, using topographical data provided by the zoo. Architectural construction documents were used to model the Gorillas of Cameroon Interpretive Center (see Figure 2), which is the ending point for most real zoo visitors, but the starting point for visitors to our virtual habitat. We used texture maps liberally to keep the polygon count low for performance reasons, while always staying under the hardware texture memory limits (1 Mbyte in the original version and 4 Mbytes in the current version). Once we generated the basic terrain model of all the habitats, we modeled Habitat 3 (the one the student explores) more extensively. Foliage, trees, and rocks were accurately modeled and placed in the habitat.

To reduce the polygon count to a level that would allow interactivity while preserving model integrity, we used several optimization techniques. We rebuilt the TIN model with a reduced number of polygons by removing vertices using the criteria that their removal would not change the terrain slope by more than five degrees over a two-foot interval within areas that the user could explore, or by 10 degrees of terrain slope in areas that the user could see but not explore. We also employed a

“point of view” heuristic to delete unseen building and terrain faces. Faces that were not visible from any permitted student position were identified and removed. Curved surfaces (rocks, tree trunks, and support structures) were modeled with as few polygon faces as possible, while smoothing techniques removed the boxy look the resulting objects would normally have.

We built models of the five different gorilla types and use two (a silverback and an adult female) in the current system. Limb lengths and circumferences were based on the available literature and are anthropometrically correct. All models currently have 11 joints and 28 degrees of freedom (see Figure 3).

After building the body models, we generated gorilla motions as a series of poses. Each pose specifies desired joint angles, global body orientation, and translation offsets to be achieved at a given time. Specifying general motion parameters in relative rather than absolute terms lets us reuse a single set of poses in many situations. In our current system, we generate poses by hand, but nothing precludes generating them with a dynamic simulation. (Motion capture is probably not a viable option, however, since it’s difficult to place the reflective dots required for such a system on the gorilla’s joints. Plus, most gorillas wouldn’t stand for having things attached to their body.) Linearly interpolating between poses generates intermediate positions. To make motions look the same (only more or less smooth) within a range of frame rates, pose playback is tied to wall-clock time instead of frame rate. Sounds are associated with motion sequences where appropriate.

In the first version of our system, the computer handled low-level gorilla behaviors such as terrain following and obstacle avoidance, while higher level behaviors were selected manually by a knowledgeable gorilla behavior expert (the so-called “Wizard of Oz” control architecture). In this system, the gorilla expert watched the students’ interactions with the gorillas and pressed keys on the keyboard to select the appropriate gorilla responses.

Terrain following and obstacle avoidance employed a terrain height field created by resampling the TIN mesh on a regular grid. Off-limits areas were indicated by large negative height values, as were fixed obstacles such as trees. A similar grid for users (who could explore the interior of the building and moats as well as the habitat) constrained their height to a constant offset from the ground. Each gorilla in the system could have its own model and its own control routines, or could use one of the five



3 Willie B, (a) in the flesh and (b) virtually.

generic ones. Each gorilla was animated by a sense-act loop that sensed the environment, took care of any reflex actions such as avoiding trees or the moats, and performed any other actions specified if no reflex actions were taken. The body parts were then moved to their new positions and the gorilla was redrawn.

Students in the VE stood on a circular platform that had a handrail completely encircling them. This was partly to provide support in case they became disoriented in the virtual world and partly to keep them from wandering beyond the reach of the tracker and head-mounted display (HMD) cables. The HMD provided a biocular (both eyes see the same image) display and monaural audio to users and had a single tracker attached to it to provide head tracking (position and orientation). A subwoofer concealed beneath the circular platform provided additional audio feedback. To let users move around the virtual world, we connected the buttons on a joystick through the mouse port for “virtual walking.”

4 Testing the prototype system at Zoo Atlanta.



Preliminary user testing

Our first prototype system had two major educational goals. First, we wanted middle school students to learn experientially about social interactions between individuals in a gorilla group based on their place in the dominance hierarchy. Second, we wanted them to learn about the design of outdoor gorilla habitats for zoo exhibits. To support these goals, we defined an initial scenario to create learning opportunities while allowing students the freedom to explore and control the pace and intensity of their experience. In this scenario, students take on the role of a juvenile gorilla. This was a natural match to our target audience of middle school students, since juveniles are younger, generally more active, and haven't mastered all the social conventions of gorilla society.

Donning the HMD, students found themselves in the virtual Gorillas of Cameroon Interpretive Center. The Interpretive Center is a building with large glass windows through which visitors can view gorilla Habitat 3, the home of male silverback Willie B and his family group. Students were first encouraged to explore the Interpretive Center to become familiar with wearing the HMD and using the hand-held joystick that let them "walk" around the environment.

Once comfortable with the system, students were told they could actually walk through the large glass windows and enter the gorilla habitat. They were also told that, upon entering the gorilla habitat, they would become a juvenile gorilla and the other gorillas would react to them according to their new identity. After passing through the glass, students could explore the habitat (including parts not visible from the viewing areas) and interact with the other gorillas. For our initial testing a silverback and a female were present in the environment with the students.

Being the low gorilla in the pecking order, if the students approached one of the other gorillas in a threatening manner or stared continuously at one of them (a social faux pas), that gorilla would become annoyed. If the annoyance persisted, the gorilla would eventually

display charging and chest-beating gestures. If the students refused to back down and flee the area, they would be put in a time-out zone and then returned to the virtual interpretive center. This represented the removal and subsequent reintroduction into a different group that would be done in the case of a real gorilla.

Once we had fully implemented our prototype system, we conducted an informal usability study with students from several area schools. These students, who ranged in age from seven to fifteen, were part of an existing educational program sponsored by Zoo Atlanta and had been coming to the zoo on a regular basis to study gorilla behaviors. Since the students were already

accustomed to visiting the zoo and working with the gorilla exhibit staff, we moved an entire VR exhibit setup into the Gorillas of Cameroon Interpretive Center at Zoo Atlanta for a day (see Figure 4).

The reaction of the students who participated in testing our first prototype at the zoo was very positive. Students had fun, and they felt like they had been a gorilla. More importantly, they learned about gorilla behaviors, interactions, and group hierarchies, as evidenced in later reactions when approaching other gorillas. Initially they would just walk right up to the dominant silverback and ignore his warning coughs, and he would end up charging at them. Later in their interactions, though, they recognized the warning coughs for what they were and backed off in a submissive manner. The observed interactions—as they evolved over time—give qualitative support to the idea that immersive VEs can be used to assist students in constructing knowledge first-hand.

Since the users were free to explore as they wished with minimal guidance from one of the project staff, they could customize their VR experience to best situate their new knowledge in terms of their preexisting knowledge base. It was interesting to note that younger students spent more time exploring the environment, checking out the corners of the habitat and the moats, and trying to look in the gorilla holding building. Older students spent more time observing and interacting with the other gorillas. The students tailored their experience to their interests and level of maturity, yet everyone spent some time on each of the two learning goals.

Implementing full autonomy

Zoo Atlanta is building the Conservation Action Resource Center, a \$9 million facility to incorporate high technology into their conservation, education, and research efforts. They're interested in having virtual gorillas as a permanent exhibit. However, since the zoo docents who would be manning the exhibit are not computer or gorilla experts, we needed to eliminate the wizard in our Wizard of Oz interface (that is, our virtual

gorillas should act completely under their own control, without any human intervention).

In some ways, building an autonomous virtual gorilla is easier than building autonomous virtual humans. But it is more difficult in others, especially given our educational goals. Although primates, gorillas have a smaller repertoire of fairly stylized interactions than do humans. However, since we don't know what, if anything, gorillas are thinking when they decide to exhibit a certain behavior, we can't accurately use their motivations or internal state to help with behavior selection. Also, since our focus is educating students about gorilla behaviors, it is not enough to simply do something interesting or entertaining as it is in video games—our virtual gorillas must exhibit the same behaviors that real gorillas do when exposed to the same situations.

We took a behaviorist approach in making our virtual gorillas autonomous by making sure they would respond correctly to various external stimuli, while not worrying about the accuracy of our model of gorilla internal state. As long as the external actions of our virtual gorillas mirror those of real gorillas in the same circumstances, we will consider our simulation a success.

Behavior control architecture

We built the gorilla behaviors in three layers. The first (which existed in the first prototype) is the reflexive layer, which handles such things as avoiding running into trees. The reactive layer handles all gorilla-gorilla and gorilla-student interactions except when preempted by the reflexive layer. Since reactive behaviors can be described as stimulus-response pairs, we used a modified state machine to control behavior selection in this layer. The modified state machine lets us specify time delays before allowing transitions. It also lets several transitions be valid at one time and uses a priority scheme to select which one to make. The volitional layer of our controller selects what action to perform when the gorilla isn't interacting with others or avoiding obstacles. Since our simulation is currently set in the middle of the day when gorillas rest, this layer selects a resting activity to perform when nothing else is happening. When we implement a settable time of day parameter, this layer will expand to encompass all gorilla solitary behaviors.

The dominance hierarchy of our virtual gorillas is basically linear, which mirrors that of small gorilla groups in real life. Silverbacks dominate males, who dominate females, who dominate juveniles and infants. Within the silverbacks, males, and females, a strict linear ordering exists of who dominates whom, while the juvenile and infant groups remain amorphous. Interactions between any two gorillas can occur based on their positions in the dominance hierarchy. If a less dominant gorilla ends up too close to or stares at a more dominant one, the more dominant one progresses through an escalating series of aggressive responses. The more submissive gorilla—who knows his place in the hierarchy (as all our virtual gorillas do)—reacts in a submissive way by turning away and moving. Thus the virtual gorillas can exhibit displacement behavior (where a more dominant gorilla approaches and then sits down in the spot occupied by a subservient gorilla

Table 1. Some gorilla control parameter values.

	Silverback	Female
Personal space:		
Front approach radius	10 feet	7 feet
Side approach radius	7 feet	5 feet
Rear approach radius	5 feet	3 feet
Staring:		
Staring radius	30 feet	22 feet
Staring field of view	30 degrees	30 degrees
Transition times:		
Stare → Annoyed	5 seconds	5 seconds
Annoyed → Angry	1 second	5 seconds
Annoyed → Content	5 seconds	5 seconds

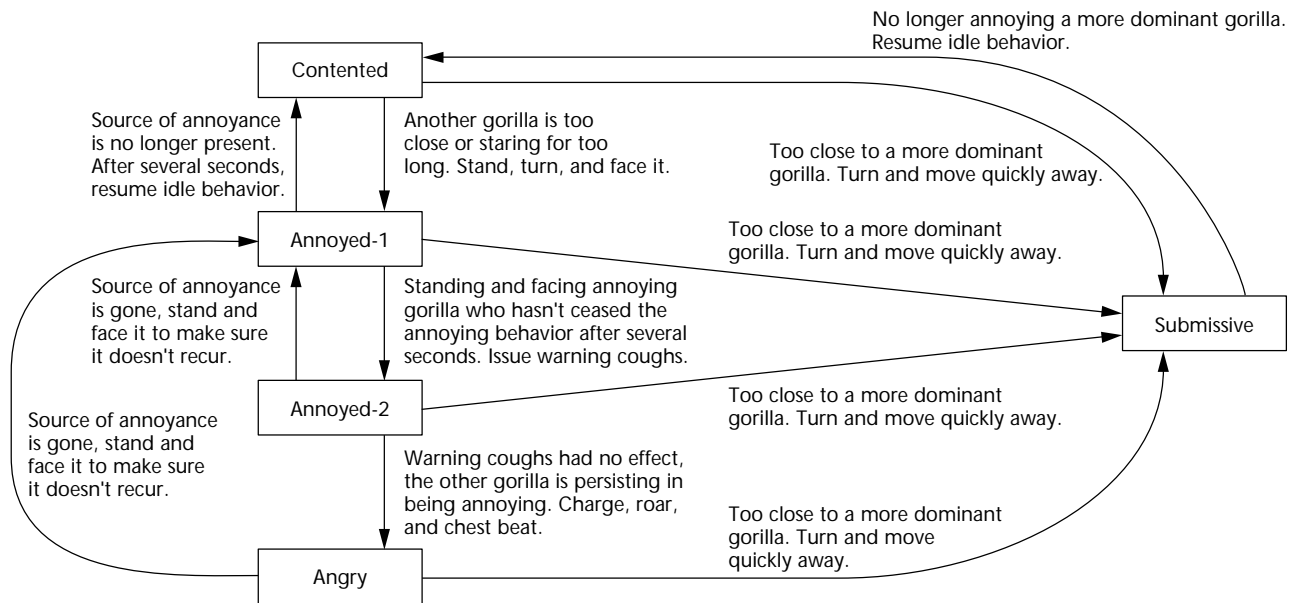
who moves a short distance) just as real gorillas do.

In cases where a specific gorilla has several others around him to interact with, submissive interactions take precedence. An example of this would be two females who are interacting—one showing dominance and the other submission—when the silverback approaches. When he gets close enough, both females switch to submissive behaviors and move. Once they are safely away from the silverback, they then resume their dominance-submission interaction. If a gorilla has several other gorillas he needs to show dominance to and none to be submissive towards, then he interacts dominantly with the closest gorilla. Thus, proximity interactions take precedence over those initiated by staring too long.

When computing interactions, such as who is staring at whom, users are assumed to be just another gorilla, with the same field of view and the same personal space as the type of gorilla they are supposed to be (which is configurable). By watching how other gorillas interact with them and each other, users can infer information about acceptable gorilla interactions and learn correct behaviors, how close they can approach other gorillas, and what kind of staring proves taboo.

The reactive behavior controller determines proximity by calculating the distance between the centers of mass of two gorillas. An additional foot and a half is added to let proximity distances take into account the body's thickness and the varying postures. If the two gorillas are closer together than this distance, a proximity violation has occurred. Since gorillas have a fairly good sense of where all the members of their group are at all times, proximity violations can occur anywhere around a gorilla, not just when the violator is within view.

Similarly, our controller determines staring by seeing if the first gorilla is within the second gorilla's field of view and if the second gorilla is in front of the first one. If so, the first gorilla observes that the second gorilla is looking at him. If this persists for longer than a predetermined time, the first gorilla decides that the second gorilla is staring rudely and takes offense. Table 1 lists some of the values used as control parameters when determining if one gorilla is too close to another or is staring annoyingly at another. Note that when users are inside the visitors' center, the moats, or other off-limits areas, they remain invisible to the other gorillas. This



Contented	In this state, the gorilla is content and engaged in solitary behavior, usually resting.
Annoyed-1, Annoyed-2	In these states, the gorilla is annoyed at a gorilla who isn't following proper gorilla etiquette. In Annoyed-1, the gorilla stands up and turns to face the source of annoyance. In Annoyed-2, the gorilla is now facing the annoyance source, and coughs and does gaze aversion.
Angry	In this state, the source of annoyance has ignored the warning coughs. The gorilla is angry and bluff charges at the source of annoyance.
Submissive	In this state, the gorilla is too close to another that is higher in the dominance hierarchy, so it turns and walks rapidly away from the more dominant gorilla.

5 Gorilla interaction state machine.

keeps the other gorillas from trying to interact with users at these times.

Once the reactive behavior controller examines all possible interactions, it sets each gorilla's mood based on the current interaction. This is then used to select the next action to take. If the action to take is the one currently occurring, it finishes before starting another instance of it. This lets associated sound files play to completion and generates more reasonable-looking motions. (It looks disconcerting to see a gorilla start to stand up and roar in preparation for a bluff charge only to drop down to all fours so it can repeat the same motion.) If our controller desires a different motion, then the transition to it from the current motion begins immediately.

Because all motions don't change realistically to all other motions, we developed a table to encode acceptable transitions. For instance, in real life when gorillas go from lying down to standing up, they sit up and then stand up. Interpolating directly between the lying and standing positions gives unrealistic-looking motion. Similarly, directly interpolating between some of the differently oriented lying positions caused the gorilla to hop vertically up into the air and then flop back down

as the constraints that the feet don't drop below ground would elevate the gorilla. To overcome this, whenever the behavior controller selects a new position or motion, it is tagged as the desired action instead of the actual next action. To determine the next pose to interpolate, the controller looks in the table for the intersection of the current action and the desired action and finds the most plausible next pose that smoothly switches from the current action to the desired action.

Dithering—where two or more actions are cycled repeatedly as a simple selection threshold is crossed and recrossed—proves problematic for simple action selection schemes. The solution is to impose some form of hysteresis in the selection mechanism. The gorillas initially exhibited this behavior when a submissive gorilla would move away from a more dominant one until its center of mass was just outside of the dominant gorilla's personal space radius. When it sat back down, the center of mass moved just inside the personal space radius, so the submissive gorilla would start to stand back up and walk further away. This action would move the center of mass outside the radius, so the gorilla would sit back down, over and over again. The end result was that the sub-

missive gorilla seemed to be scooting away on its bottom! Our solution for this (and other dithering problems) was to finish the current action when interacting with another gorilla, so that one complete walk cycle would be performed before the submissive gorilla would sit down again. This would move the submissive gorilla far enough away so that it didn't inadvertently move back inside the dominant gorilla's personal space radius.

State machine details

The state diagram (Figure 5) defines the basics of how gorillas interact with each other and with the user. Since we worked with a small group, we assumed a linear dominance hierarchy, with the silverback dominating the females, who dominate the juveniles (including the student user). Normally, each gorilla remains in the Content state. In the Content state, the volitional layer controls a gorilla's behavior.

Let's trace the interaction sequence for a particular gorilla, say gorilla X. If a gorilla who belongs in the lower class of the dominance hierarchy enters gorilla X's personal space (as defined in Table 1), or if the less dominant gorilla stares at gorilla X from a position in his field of view for more than a specified time, gorilla X switches to the Annoyed-1 state. In this state, gorilla X stands up (if he wasn't already) and turns to face his source of annoyance. If the annoyance disappears before gorilla X stands and faces it directly, gorilla X switches back to the Content state after a suitable cool-down period. If the annoyance still exists, however, gorilla X then switches to the Annoyed-2 state. In this state, gorilla X coughs at the source of annoyance and performs gaze aversion (that is, looks at another gorilla with momentary glances). If these gestures don't dissuade the annoyance source, then gorilla X finally switches to the Angry state and proceeds to bluff charge at the annoyance source while chest beating. Otherwise, gorilla X switches to the Annoyed-1 state where he cools down before switching back to Content.

At any time during this process if a more dominant gorilla approaches gorilla X, gorilla X switches to the Submissive state and turns and walks away from the more dominant gorilla. Once gorilla X is at a safe distance, he switches back to the Content state and then on to other states as necessary.

What next?

Having done a trial run with real middle school students using our initial prototype system, we now have a better idea of the types of questions we need to answer when building a VR system for educational purposes. We are addressing these issues in the latest version of the system. However, even the results of our first trials seem to indicate that it is possible to use VR as a general educational tool for middle school students. VR sys-



6 A student interacting with the virtual Willie B while a gorilla expert provides assistance.

tems lets students experience the real world from viewpoints other than their own, and they can learn from first-hand experience in environments that would normally be too dangerous or impossible for them to experience in the real world. Given a rich but accurate environment in which to interact, students can personalize their experiences and internalize the content presented through first-person interactions.

We intend to make several improvements to our system before conducting more extensive user testing. We observed that even though our student testers had been watching gorillas for several weeks as part of the zoo's "gorilla squad," there were times during their interaction with the system where they were overwhelmed. What was happening in these interactions was so far beyond what they had already learned that they were unable to connect the new information with what they already knew in any meaningful way. (For example, this happened when they heard some of the gorilla vocalizations, which visitors are normally too far away to hear.) During these times, it proved helpful to have one of the zoo's gorilla experts there to explain what the students were seeing and hearing. We are currently encapsulating the assistance provided by the gorilla experts and incorporating it in audio form into our VE to assist future students in constructing new knowledge more easily (see Figure 6).

Also, students expressed a desire to have a peer to play with in the environment. Since juvenile play behavior is another important part of gorilla social interactions, we're investigating ways to add explicit play behaviors to the system. A lot of play behavior, though, involves touching and interacting with objects in the environment and physical contact with other gorillas, something that is difficult to simulate effectively with our current system.

Finally, once the construction of the Conservation Action Resource Center is complete and the virtual gorilla environment is ensconced there, we will conduct more thorough user studies to validate the observations we made during our pilot user testing. This will help determine more specifically what type of VE proves con-

ducive for rapid, efficacious learning and long-term knowledge retention. ■

Acknowledgments

We would like to thank Zoo Atlanta, its director, Terry Maple, and gorilla experts Lori Perkins, Kristen Lukas, and Kyle Burks for their assistance with and support for this project. Additional support for this project was provided by the Center for Educational Technology at Georgia Institute of Technology and the National Science Foundation.

References

1. C.D. Wickens, "Virtual Reality and Education," *IEEE Int'l Conf. on Systems, Man, and Cybernetics*, IEEE Press, Piscataway, N.J., Oct. 1992, pp. 842-847.
2. S.K. Damarin, "Schooling and Situated Knowledge: Travel or Tourism," *Educational Technology*, March 1993, pp. 27-32.
3. J.W. Belsford, "Physics Education in a Virtual Environment," *Proc. Human Factors and Ergonomics Soc. 37th Annual Meeting*, Human Factors and Ergonomics Society, Santa Monica, Calif., 1993, pp. 1286-1290.
4. W. Winn, "The Virtual Reality Roving Vehicle Project," *T.H.E. Journal*, Dec. 1995, pp. 70-74.
5. M. Roussos et al., "NICE: Combining Constructionism, Narrative, and Collaboration in a Virtual Learning Environment," *Computer Graphics*, Vol. 31, No. 3, Aug. 1997, pp. 62-63.
6. P. Maes et al., "The Alive System: Full-Body Interaction with Autonomous Agents," *Computer Animation 95 Proc.*, IEEE Computer Society Press, Los Alamitos, Calif., 1995, pp. 11-18.
7. X. Tu and D. Terzopoulos, "Artificial Fishes: Physics, Locomotion, Perception, and Behavior," *Proc. Siggraph 94*, A. Glassner, ed., ACM Press, New York, July 1994, pp. 43-50.
8. J. Bates, "The Role of Emotion in Believable Agents," *Comm. ACM*, July 1994, pp. 122-125.
9. B. Hayes-Roth, L. Brownston, and E. Sincoff, *Directed Improvisation by Computer Characters*, Tech. Report KSL-95-04, Stanford Univ. Knowledge Systems Laboratory, Stanford Univ., Palo Alto, Calif., 1995.
10. T.L. Maple and M.P. Hoff, *Gorilla Behavior*, Van Nostrand Reinhold, New York, 1982.



Don Allison is a PhD student in the Graphics, Visualization, and Usability Center and the College of Computing at Georgia Institute of Technology. He has a BS from Bethany Nazarene College in math and physics (1979), an MS in math from the University of Illinois (1981), and another MS in computer engineering from National Technological University, headquartered in Ft. Collins, Colorado. His research interests include VR, behavioral animation, and AI. He is a member of the IEEE, IEEE Computer Society, ACM, Siggraph, and the Mathematical Association of America.



Brian Wills is a research scientist jointly appointed by the College of Architecture and the Graphics, Visualization, and Usability Center at Georgia Institute of Technology. He has a BS and MA in architecture from Georgia Institute of Technology. He is currently working on VR environments and developing digital tools for architectural design processes.



Doug Bowman is a PhD student in the Graphics, Visualization, and Usability Center and the College of Computing at Georgia Institute of Technology. He received a BS in mathematics and computer science from Emory University in 1994, and an MS in computer science from Georgia Institute of Technology in 1997. His research interests include human-computer interaction issues in immersive virtual environments and applications of virtual environments for education, information visualization, and design.



Jean Wineman is an architectural researcher at the College of Architecture at Georgia Institute of Technology. She received a BA in psychology and sociology from Wellesley College in 1971. She graduated as a Wellesley College Scholar Master of Urban Planning and received her PhD in architecture at the University of Michigan in 1973 and a PhD in 1977 with a specialization in aging from the Institute of Gerontology at the University of Michigan. Her research interests include museum and exhibit design, the design of work settings to enhance productivity, and design for special populations. She is a member of the Environmental Design Research Association.



Larry F. Hodges is an associate professor in the College of Computing and Associate Director for Industrial Relations of the Graphics, Visualization, and Usability Center at Georgia Institute of Technology. He received his PhD in computer engineering from North Carolina State University in 1988. He is on the editorial board of the journal *Presence: Teleoperators and Virtual Environments* and is the conference chair for the 1998 IEEE Virtual Reality Annual International Symposium (VRAIS 98). He is a member of IEEE Computer Society.

Contact Allison at the Graphics, Visualization, and Usability Center, Georgia Institute of Technology, Atlanta, GA 30332-0280, e-mail don@cc.gatech.edu.