

Physically interactive story environments

by C. S. Pinhanez A. D. Wilson
J. W. Davis A. F. Bobick
S. Intille B. Blumberg
M. P. Johnson

Most interactive stories, such as hypertext narratives and interactive movies achieve an interactive “feel” by allowing the user to choose among multiple story paths. In this paper we discuss physically interactive environments with narrative structure in which the ability to choose among multiple story lines is replaced with having users, first, interact with the story characters in small, local “windows” of the narrative and, second, actively engage their bodies in movement. In particular, we found that compelling interactive narrative story systems can be perceived as highly responsive, engaging, and interactive even when the overall story has a single-path structure, in what we call a “less-choice, more-responsiveness” approach to the design of story-based interactive environments. We have also observed that unencumbering, rich sensor technology can facilitate user immersion in the experience as the story progresses—users can act as they typically would without worrying about manipulating a computer interface. To support these arguments, the paper describes the physical setup, the interactive story, the technology, and the user experience of four projects developed at the MIT Media Laboratory: KidsRoom, It/I, Personal Aerobics Trainer, and Swamped!

Following Myron Kruger’s pioneering work in the 1980s on computer-controlled interactive spaces,¹ the beginning of the 1990s saw an explosion in the creation of physically interactive environments^{2–4} for entertainment, where users could explore an environment or interact with a character. Initially confined to research laboratories, physically interactive environments are now commonly available in arcades and museums. At the same time, video games moved from simple shoot-and-kill scenarios to increasingly complex

narratives in multicharacter stories such as *Myst*, role-playing games such as *Tomb Raider*, or strategic games such as *The Age of the Empires*.

The four projects developed at the MIT Media Laboratory described in this paper created interactive environments that physically engage their users as characters in a story, in an attempt to merge the compelling interaction of physically interactive environments with the engagement and suspension of disbelief of complex stories. *KidsRoom*,⁵ *It/I*,⁶ *Personal Aerobics Trainer*,⁷ and *Swamped!*⁸ immerse their users in physical experiences—each with a beginning, development, and an end, exploring the narrative notions of climax and catharsis in the interactive realm. We call such interactive story-based computer-controlled environments *physically interactive story environments*.

The goal of this paper is to examine and discuss the technological and narrative mechanisms used to combine physical action with interactive stories in these four projects. In particular, these projects do not follow current conventions of interactive storytelling and computer games. First, no cumbersome devices such as joysticks or head-mounted displays are employed. Second, unlike many other virtual reality (VR) interactive environments, realistic rendering is not used.

©Copyright 2000 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

The most distinctive characteristic of these four projects is that the participants in these environments have very limited control over the overall development of the story. Although the characters and scene environments are responsive to the actions of the users at any given moment during the story, all users experience the same basic story. Whereas traditional interactive stories have tried to make the user feel the story responsive by providing control over the story development through some form of choice among possible story paths,⁹ the projects described in this paper do not use this choice mechanism. Nevertheless, our users feel that our environments are interactive. They appear to experience a sense of control over the story. We believe that this responsive feeling is conveyed because (1) our characters respond in compelling ways to small decisions and actions users make while interacting with them, and (2) because the users can engage in natural, pleasing physical movements during the experience. This constitutes what we call the “less-choice, more-responsiveness” design principle for interactive story environments.

These issues are initially discussed in the next section of the paper in the context of previous work on physical story-based environments and interactive stories. Then, in four separate sections, we describe the four projects we have developed at the Media Laboratory. We conclude by comparing the four interactive experiences and discussing possible reasons for the success of our less-choice, more-responsiveness approach to designing physically interactive story environments.

Physically interactive stories

Since ancient times children have been playing games in which they pretend to be characters living real and fantastic stories. Similarly, role-playing¹⁰ has been part of human ritual for millennia in religious ceremonies and theatrical performances. Role-playing combines the emotional component of narrative with the physical activity of the body to create a powerful sense of corporeal immersion in a character, environment, or communal act. The sensation of immersion in such situations is many times enhanced by the use of scenarios, costumes, and ritual objects as well as by the presence of professional performers portraying some of the characters.

New technologies have been employed throughout the ages to increase the user’s feeling of physical immersion in stories. The advent of computers made

it possible to create compelling story-based environments with realistic imagery and sound where computer characters are always available—24 hours a day—to help play out the user’s fantasy.

Physical story-based environments. Physical realizations of stories have been part of human culture for centuries (e.g., theater and religious rituals). The *panoramas* of the nineteenth century were one of the first examples of capturing stories in environments controlled by machines. A panorama consists of an observation platform surrounded by a painted canvas. To create the illusion of reality, the edges of the canvas are hidden by an umbrella-shaped roof that covers the platform and by a “false terrain” projecting from the observation platform. Panoramas were extremely popular throughout the nineteenth century. They depicted landscapes, battle scenes, and journeys (see Oettermann¹¹ for a history of panoramas).

Although mechanical “haunted houses” have populated the carnivals and fairs of the twentieth century, the development of Disneyland** and other theme parks pushed the limits of technology in terms of creating vivid physical renditions of characters and stories with machines. Disneyland pioneered the use of *animatronic puppets*—sophisticated robotic devices with lifelike movement. These puppets are used to endlessly recreate a physical realization of a story where the characters display a fairly complex set of actions.

However, the traditional theme park ride lacks interactivity. Visitors exist in the story space but their actions are never actually reflected in the development of the story or the life of the characters. Although many theme park rides move, shake, and play with the participants’ sense of equilibrium, quite commonly the users’ physical activity is severely restricted.

About the same time that animatronics became popular, tiny, extremely reactive characters started to populate the screens of video games in arcades. With simple devices such as buttons and joysticks, it became possible to interact with such characters in environments that, though responsive, were very limited in their ability to foster user immersion: the user’s connection to that environment was restricted to a joystick and a video display. Since then the interfaces of arcade games have advanced considerably, enabling full-body action and sensing. Examples include many skiing and cycling games and the

recent Japanese “dance” game Dance Dance Revolution.¹²

There are only a few cases where the full-body arcade interactivity has been combined with more complex story and characters. A good example is the ride Aladdin, developed and tested at Disneyworld,¹³ where four users wearing VR helmets loosely control a voyage through a city on a flying carpet. Although most of the user control is restricted to deciding what to see, the presence of the users affects the behavior of the characters in the city.

Interactive stories. Most of the academic discussion about interactive stories has been concentrated in the field of literature, which paradoxically uses one of the least interactive media: paper. Murray⁹ examines the characteristics of interactive narrative, particularly in the cases where the story develops as a reaction to the reader’s action. Normally this action consists of choosing the next development in the story from a small set of options proposed by the author.

Many “narrative” video games, such as Dragon’s Lair, are also based on a “choice” structure, although in this case the user’s goal is typically to discover patterns of choice that minimize the traversal of all possible paths. Similarly, there have been attempts to create *interactive movies* where the audience decides, by voting, among different paths for the characters. The most widely known example, Mr. Payback,¹⁴ was coldly received by the critics and was not commercially successful.

However, most studies of interactive stories (e.g., Murray⁹) neglect the rich experience of role-playing in dramatic games and human rituals. In spite of all the interaction among role-playing participants, the choices made by the participants tend to keep the game inside of the limitation of the “game reality” (which can encompass much fantasy). Role-players normally avoid the full exploration of the tree of multiple paths (including its uncontrolled dead ends), and seem to extract most of the pleasure from the portions of rewarding mutual interaction and character discovery that happen during the game play through fortuitous encounters.

Less choice, more responsiveness. The two sections above herald the main hypotheses of this paper that in physically interactive stories responsiveness is likely to be more important than choice. We have reached this conclusion based on the fact that we


have designed and constructed engaging experiences that feel highly interactive without providing to the participants any real control over the story. These interactive environments, described later in this paper, clearly show that story choice is not a prerequisite for interactive stories. It could be argued that it is possible to structure the interactivity in a story-based environment around mechanisms of choice. However, so far we have not seen a successful implementation of a physically interactive story environment based on extensive story choice.

Galyean¹⁵ proposed a *water-skier* model for interactive experiences, similar to the less-choice, more-responsiveness paradigm proposed here. In the water-skier model, the user is compared to a water-skier who is unable to determine the direction of the pulling boat (the story) but who has some freedom to explore his or her current situation in the story. This model was employed by Blumberg and Galyean in the Dogmatic¹⁶ VR experience, where the user’s avatar encounters a dog, Lucky, in a desert town and, ultimately, is led by the dog to her own death. As users look around and decide where to go, they involuntarily became part of the story. Most of the pleasure associated with Dogmatic seems to be derived from the process of understanding the story. Although the water-skier model advocates limited story choice, it fails to observe that compensation for this limitation on story control can be made with an increase in responsiveness and local interaction.

Unlike Dogmatic, the projects described here focus on creating rewarding experiences in the process of interacting with the characters in the story and in the physical aspects of this interaction. That is, as in many physical rituals and theatrical games, the satisfaction in such interactive stories comes from the pleasure of small, localized interactions with the other characters.

It is possible to keep the users and the characters in the context of a well-structured and interesting story by concentrating story development on local interactions instead of providing multiple story paths. A key problem with multiple-path stories is that some paths are considerably weaker than others. As described by Marinelli,¹⁷ choice-based interactive stories are like season tickets to hockey games: the experience involves some good games, some boring games, and hopefully a truly remarkable evening that becomes a story to be told to our kids. A great story is a very special, fortunate, and rare conjunction of ideas, events, and characters.

Figure 1 Physical setup of the KidsRoom



By developing a system that is locally responsive to user actions as the user progresses through a single-threaded story, we can assure that our users always receive the full impact of the best possible story (as handcrafted by its author) without losing the impression that the story unfolds as a consequence of the participants' actions. To illustrate these ideas, we now examine four projects developed at the MIT Media Laboratory from 1996 to 1999.

KidsRoom

The KidsRoom⁵ project aims to create a child's bedroom where children can enact a simple fantasy story and interact with computer-controlled cartoonish monsters. It is a multiuser experience for children between 6 and 12 years of age where the main action happens in the physical space of the room and not "behind the computer screen," as in most video games. Furthermore, the children are not encumbered by sensing devices, so they can actively walk,

run, and move their bodies. A detailed description of the project can be found in Reference 5.

The physical setup. KidsRoom is built in a space 24 by 18 feet with a wire-grid ceiling 27 feet high. Two of the bedroom walls resemble real walls of a child's room, complete with real furniture and decoration. The other two walls are large video projection screens, where images are back-projected from outside of the room. Behind the screens there is a computer cluster with six machines that control the interactive space. Computer-controlled theatrical colored lights on the ceiling illuminate the space. Four speakers, one on each wall, project sound effects and music into the space. Finally, there are three video cameras and one microphone used to sense the children's activity in the room. Figure 1 shows a view of the complete KidsRoom installation.

The room has five types of output for motivating participants: video, music, recorded voice narration,

sound effects, and lighting. Still-frame video animation is projected on the two walls. Voices of a narrator and monsters, as well as other sound effects, are directionally controlled using the four speakers. Colored lighting changes are used to mark important transitions.

The interactive story. The KidsRoom story begins in a normal-looking bedroom. Children enter after being told to discover the magic word by “asking” the talking furniture, which speaks when approached. When the children scream the magic word loudly, the room transforms into a mystical forest. There, the children have to stay in a single group and follow a defined path to a river. Along the way, they encounter roaring monsters and, to stop them, they have to quickly hide behind a bed. To guide the participants, the voice of a narrator, speaking in couplets, periodically suggests what the children can do in the current situation.

After some time walking in this forest, the children reach a river and the narrator tells them that the bed is now a magic boat that will take them on an adventure. The children can climb on the “boat” and pretend to paddle to make it “move,” while images of the flowing river appear on the screens. To avoid obstacles in the river the children have to collaboratively row on the appropriate side of the bed; if they hit the obstacles, a loud noise is heard. Often, the children add “realism” by pushing each other.

Next, the children reach the monster world. The monsters then appear on the screens and show the children some dance steps. The children have to learn these dance steps to become friends of the monsters. The monsters then mimic the children as they perform the dance moves. Finally, the children are commanded to go to bed by an insistent, motherly voice, and the adventure ends with the room transforming back to a normal bedroom.

The technology. Three cameras overlooking the KidsRoom are used for the computer vision analysis of the scene. One of the cameras (marked as the “overhead camera” in Figure 1) is used for tracking the children and the bed in the space. The overhead position of the camera minimizes the possibility of one user or object occluding another. Further, lighting is assumed to remain constant during the time that the tracker is running. Standard background subtraction techniques (described by Davis and Bobick¹⁸) are used to segment objects from the background, and the foreground pixels are clustered into

2-D blob regions. The algorithm then maps each person known to be in the room to a blob in the incoming image frame. In the scene with the boat, the overhead camera is also used to detect whether the children are rowing and in which direction.

The other two cameras (marked as “recognition cameras” in Figure 1) are used to recognize the dance steps performed by the children during the last scene with the monsters. The images from these cameras are processed to segment the silhouette of the child facing the screen. Using motion-template techniques¹⁸ the system distinguishes the occurrence of four different “monster dance steps”: crouching, spinning, flapping arms, or making a “Y” figure.

The experience. KidsRoom was designed and built in the fall of 1996. The installation was experienced by dozens of children and adults during the three months it remained open (see Figure 2). A new, shortened version of the KidsRoom, the KidsRoom2, was built in 1999 in London as part of the Millennium Dome Project and is scheduled to run continuously through the year 2000.

A typical run of KidsRoom at the Media Laboratory lasts 12 minutes for children and usually slightly longer for adults. Not surprisingly, we found children to be more willing to engage in the action of the story and to follow the instructions of the narrator. Children are typically very active when they are in the space, running from place to place, dancing, and acting out the rowing and exploring fantasies. They interact with each other as much as they do with the virtual objects, and their exploration of the real space and the transformation of real objects (e.g., the bed) enhance the story.

From our observation of the children, there has never been a situation where the children did not understand that they are characters in a story and that they have to act out their parts to make the story flow. Occasionally the children do not understand the instructions and the experience has small breaks in its flow. However, the control software of KidsRoom is designed to always push the story forward, so such interruptions are usually overcome quickly.

The story of KidsRoom ties the physical space, the participants’ actions, and the different output media together into a coherent, rich, and immersive experience. In particular, the existence of a story seems to make people, and especially children, more likely to cooperate with the room than resist it and

test its limits. The well-crafted story also seems to make participants more likely to suspend disbelief and be more curious and less apprehensive about what will happen next.

In fact, the users of KidsRoom have absolutely no control of the overall story development and do not seem concerned at all about that. Some of the best moments of the experience, as judged by the enthusiastic reaction of the young users, are connected to simple physical activities (augmented by image and sound) such as rowing on the river, dancing with “live” cartoonish monsters, or simply running in a group toward the bed, away from the monsters, and piling on top of one another.

It/I

It/I is a theater play in which a computer system takes the part of one of the play’s two characters. The computer character, called *It*, has a nonhuman body composed of computer graphic (CG) objects projected onto rear-projection video screens. The objects are used to play with the human character, *I*, performed by a real actor on a stage. The play is a typical example of *computer theater*, a term (proposed by Pinhanez¹⁹) that refers to theatrical experiences involving computers, in direct analogy to the idea of *computer music*.

It/I was created with two goals in mind. The first goal was to design an automatic interactive computer character that could co-inhabit the stage with a human performer in front of an audience, throughout the length of a complex story. This imposes strong requirements in terms of expressiveness and reliability on the computer “actor.” The second goal was to create a space where persons could re-enact a story they have watched by taking the place of the human performer; Pinhanez refers to this space as an *immersive stage*.²⁰ A detailed description of the play and its underlying technology can be found in References 20 and 6.


The physical setup. Figure 3 depicts a diagram of the different components of the physical setup of It/I. The sensor system was composed of three cameras rigged in front of the stage. The computers controlled different output devices: two large back-projected screens, speakers connected to a MIDI (Musical Instrument Digital Interface) synthesizer, and stage lights controlled by a MIDI light board.

The interactive story. It/I depicts a story about relationships between human beings and technology.

Figure 2 Users experiencing the KidsRoom



Figure 3 Physical setup of *It / I*



The character played by the computer, called *It*, represents the technology that surrounds and many times controls us; that is, in *It/I*, the computer plays itself. *It* is, in fact, quite an unusual creature: it has a “body” composed of CG objects (representing clocks, cameras, televisions, electrical switches) projected on stage screens. *It* can “speak” through large, moving images and videos projected on the screens, through musical sounds played on stage speakers, and through the stage lights. Figure 4 depicts two scenes from the play.

The play is composed of four scenes, each being a repetition of a basic cycle: *I* is lured by *It*, *I* is played

with, *I* gets frustrated, *I* quits, and *I* is punished by *It* for quitting. For example, in the second scene a CG object similar to a photographic camera appears on a small screen and follows *I* around. When *I* accepts the game and makes a pose for the camera, the camera’s shutter opens with a burst of light. Then, on the other screen, a CG television appears, displaying a slide show composed of silhouette images “taken” by the camera. After some pictures are shown, the camera “calls” *I* to take another picture. This cycle is repeated until *I* refuses to take yet another picture (that is, the human performer decides it is a good time to finish the cycle), provoking an irate reaction from *It*, which in response throws CG-blocks

Figure 4 The human and the computer characters during scenes of *It / I*



at *I* while flickering the lights and playing threatening sounds.

The technology. The primary sensors in *It/I* are three video cameras positioned in front of the stage. Using the information from the three cameras, it is possible to segment the human character from the background using a stereo system similar to the one proposed by Ivanov et al.,²¹ independently of stage lighting and changes on the background screens.

The play was written taking into account the sensory limitations of computer vision technology. That is, the actions of *I* are restricted to those that the computer can recognize automatically through image processing. In many ways, *It's* understanding of the world reflects the state-of-the-art of real-time automatic vision: the character's reaction is mostly based on tracking *I's* movements and position and on the recognition of some specific gestures (as in Davis and Bobick¹⁸).

Unlike most interactive environments, *It/I* portrays a long and complex story that lasts for about 40 minutes. Additionally, the control system of the play has to be extremely robust to cope with the requirement of live performances in front of large audiences. Since at the time when the play was produced there were no story representational languages able to satisfy both requirements,²² it became necessary to develop a special language for representation of interactive stories. In *It/I* the control of all the sensor and actuator systems is described in an *interval script*, a


paradigm for interaction scripting based on the concept of time intervals and temporal relationships developed by Pinhanez,²⁰ based on previous work with Mase and Bobick.²² A description of the interval script paradigm is beyond the scope of this paper and can be found in Reference 20.

The experience. *It/I* was performed six times at the MIT Media Laboratory for a total audience of about 500 people. The audience clearly understood the computer character's actions and intentions, and the play managed to keep the "suspension of disbelief" throughout its 40 minutes. In particular, the sound effects played a key role in creating the illusion that *It* was alive and in conveying the mood and personality of the character. Each performance was followed by an explanation of the workings of the computer-actor. After that, audience participants were invited to go on stage and play the second scene (as described earlier), first in front of, then without, an audience (see Figure 5).

Theater scripts are clear examples of stories in which the general story structure and development is fixed, but the realization of the character interaction is left to the performers. Although actors usually have no influence on how the story unfolds, they are responsible for discovering and creating the minutia of the moment-by-moment intercharacter relations.

It/I follows this traditional structure and therefore, by design, creates an interactive computer environment based on responsiveness. Given the argument

Figure 5 Audience participants playing in the *immersive stage* of It / I after the performance of the play



above, its similarity to traditional theater makes it clearly a comfortable place for the actor. Beyond that, we observed that the audience also enjoyed engaging in this experience, where they had no control of the final outcome of the story but where play-acting was fun.

Compensation for the lack of story control in It/I is accomplished by expanding the repertoire of physical interaction. During the play, the actor (and later, the audience) is able not only to explore physical activity but also to use the body to produce and change imagery, sound, and lighting. This is certainly one of the possibilities not present in traditional role-playing that is made available by computer-mediated spaces for physically interactive stories.

Personal Aerobic Trainer

Whereas the two projects described above create stories populated by fantastic characters, the Personal Aerobic Trainer project, or PAT, is focused on the quality of the physical activity of the user. The main goal of PAT is to create a system that helps a user to work out athletically by enthusiastically pushing the user through a series of aerobic exercises while monitoring the user's activity and correcting his or her movements. A detailed description of the PAT system can be found in Reference 7.

The physical setup. The silhouetting method employed for monitoring the user in the PAT system is based on the optical blocking (or eclipsing) of infrared (IR) light rather than the use of color differ-


ences between the person and background (as done in KidsRoom) or stereo disparity (as in It/I). This is necessary because to monitor the quality of an aerobic movement it is important to have a very precise and sharp silhouette of the user. However, the space in PAT is carefully engineered to hide the IR and the sensing apparatus.

Figure 6 shows the environmental configuration of the PAT system. It consists of a room in which two of the walls are replaced by large screens on which video is back-projected. Behind one of the screens there is an array of IR emitters that evenly lights the screen. In front of the opposite wall a camera equipped with an IR filter is positioned facing the IR illuminated screen.

This configuration allows the camera to quickly and effortlessly obtain a high-quality silhouette of the user, in spite of light changes in the room and the images in the videos projected in the screens. The infrared light is not visible to the human eye and thus the user sees only the video projected on the display screens.

The interactive story. The experience in PAT starts when the user enters the space. The entrance of the user triggers the opening of a video window in the screen with the camera, portraying a virtual instructor (an army drill sergeant) welcoming the user. The instructor images are obtained from a collection of prerecorded video clips depicting a multitude of actions, attitudes, and comments spanning a reasonable range of possible reactions for the drill sergeant.

Figure 6 Physical setup of PAT



After the brief introduction, the system goes through a sequence of physical exercises. Ideally, that sequence would be personalized according to the identity and physical history of the user. For each exercise, the instructor executes the moves and accompanies the user during the performance. The drill instructor gives feedback, often humorous, based on how well the user is performing the exercises. During each exercise, background music is synchronized to the user movements (unlike workout videos, which make the user follow the pace of the music). After the workout is complete, the instructor congratulates the user. If the user prematurely leaves the space, the instructor stops and leaves the screen.

The technology. The PAT system employs the same real-time computer vision methods for recognizing large-scale body movements proposed by Davis and Bobick.¹⁸ The method is based on the layering of participant silhouettes over time onto a single template and measuring shape properties of that template to recognize various aerobic exercise (and other) movements in real time.

Figure 7 presents templates generated from the infrared silhouettes for the movements of left-arm-rise (left-arm stretch) and fan-up-both-arms (deep-

breathing exercise stretch). The motion of the user is encoded in the varying gray levels within the template. For recognition of these moves, statistical pattern recognition techniques are applied to moment-based feature descriptors of the templates. The system employs user training to get a measure of the variation that results from participation of different people.¹⁸

The experience. The first prototype for the Personal Aerobics Trainer was set up at the Villers Facility of the MIT Media Laboratory in December of 1997 and was experienced by many users over a period of three months. Figure 8 displays a model of the interaction of the PAT system, where a user is exercising in front of a television. Users can easily understand the structure of the interaction and become naturally engaged in the activity. The comments of the drill sergeant seem to help to make the routine more personal, as well as creating a sense of responsibility and accomplishment that is lacking in traditional workout videotapes. Although the PAT system does not have a traditional story line like the three other projects described in this paper, the experience is clearly structured as a narrative. Indeed, the drill sergeant character it employs interacts with the user, who assumes the role of an army private.

Figure 7 Motion templates used by PAT for recognizing “left-arm stretch” and “deep-breathing stretch”



PAT exemplifies a temporally structured human activity where immediate response to instructions is more important than narrative choice. In fact, the critical aspect of a workout experience is to make the user correctly perform a sequence of physical tasks by managing the user's natural desire to quit. In other words, the system strives *to prevent choice*. Although the adaptation of the system to the user's pace is important, making the user persevere is achieved in PAT mostly by the personification of the control system, which creates the feeling that the user is being watched and stimulated by another person—the drill sergeant. Moreover, in PAT the physical activity is at the center of the interaction and its health benefits constitute the basic source of reward.

Swamped!


Swamped!⁸ is an interactive system developed at the Media Laboratory, in which the participant assumes the role of a chicken that is trying to protect its eggs from a hungry raccoon in a barnyard setting. Unlike the full-body interaction of the previous projects, in Swamped! the user controls a character by manipulating a plush doll representing that character. Moreover, one of the main goals of the project is to explore how a manipulative interface can be used—not to explicitly control the character's body movements (as in most shoot-and-kill video games), but instead to suggest a line of action for a character. A detailed description of the project can be found in Reference 8.

The physical setup. In Swamped! the user stands in front of a projection screen showing the virtual world and the virtual chicken, while holding a plush doll similar to the chicken. The user can direct the chicken by making appropriate gestures with the doll. For example, wobbling the doll back and forth makes the virtual chicken walk; flapping the doll's wings makes it fly. The participant's attention is meant to focus on the interactions in the virtual world and not on the doll itself. Figure 9 shows a user in front of the screen holding the “controller” doll while watching the unfolding saga between the chicken and the raccoon.

The interactive story. At the start of the interaction, the user discovers that he or she is playing the role of a chicken trying to protect its eggs from a raccoon. Figure 10 shows a picture from a typical run. The chicken displays various behaviors such as squawking to get the raccoon's attention and make it angry, scratching its own head, kicking the raccoon, and setting a trap for the raccoon. As described above, these behaviors are selected through the doll's movements and the story context. The raccoon is fully autonomous, choosing what actions to take based on its desires, perceptions, and emotional state.¹⁶

In a normal interaction, the raccoon's attempts to get the eggs are blocked by the user-manipulated chicken. However, the raccoon eventually gets one egg and then runs away with it. Next, it stops to examine the egg on a giant bull's-eye painted on the

Figure 8 Typical interaction of the PAT system



ground. Guess what? When the raccoon looks up, a heavy weight descends from the sky and smashes the raccoon.

The technology. The physical doll used to control the chicken character is fabricated to match the virtual character. An armature made of plastic, brass tubing, and wire holds a sensor package and provides an articulated structure (see Figure 11). The sensor package inside the doll includes an array of 13 sensors: two pitch and roll sensors, one gyroscope sensing roll velocity, three orthogonally mounted magnetometers sensing orientation with respect to magnetic north, two flexion sensors for wing position, three squeeze sensors embedded in the body and beak, and one potentiometer to sense head rotation about the neck.

Raw data from the doll are processed in real time on the host computer to recognize gestures that are taught to the system in a learning phase. The system can detect a variety of actions of the doll under user control, such as *walk*, *run*, *fly*, *squeeze-belly*, *hop*, *kick*, and *back flip*. Each of these action primitives is learned off line and recognized using hidden Markov models²³ (HMMs).

The chicken's behavior system treats the output of each HMM as a sensory input to a corresponding consummatory behavior, using a reactive behavior system similar to the one proposed by Blumberg.²⁴ For example, when the user flaps the chicken's wings, the HMM for the flying gesture surpasses its threshold and stimulates the flying behavior. If this is the most appropriate behavior at the time, the flying behav-

Figure 9 User in front of the screen of Swamped! controlling the computer graphics character by manipulating the plush doll



ior becomes active, which causes the virtual chicken to begin flying.

The experience. Over 400 users interacted with the *Swamped!* installation in the Enhanced Realities exhibit at SIGGRAPH 98. Users were given the cartoon scenario and were told that the goal was to keep the raccoon busy by activating the chicken's various behaviors.

We found three categories of users: *teachable*, *ideal*, and *skeptical* (in decreasing order of approximate group size). The ideal users are often children who pick up the doll, start manipulating it, and immediately understand the concept of the interface. The teachable users are by far the largest group. The typical member of this group picks up the doll and tries to manipulate one part, such as one wing or a foot, expecting a direct mapping. After a walking gesture is made and the "voodoo doll" metaphor is demonstrated, many of these users can quickly learn to use

the doll and enjoy the experience. Several users, however, never understand how the doll controls the character and are even skeptical about connections between the doll's moves and the character's behavior.

Although the story line of *Swamped!* is quite simple, it sets up clear goals to guide the user interaction. Moreover, the raccoon character fulfills the dramatic role of pushing the story ahead, culminating, pathetically, in its own smashing. However, there is no doubt that involving the user in a story has been critical for setting clear objectives for the manipulation of the doll. It avoids an "experimentation phase" in which the user tries everything, making it very hard for the system to recognize a particular intention of the user. For instance, the situation in the story might require the "chicken" to either run or to fight against the "raccoon," so the system can concentrate on distinguishing between both gestures.

Figure 10 Snapshot of scene of *Swamped!*



Although it can be argued that the user is making choices in terms of story whenever the user decides for a particular line of action for the chicken, it is important to recognize that most of the pleasure of the installation seems to come from the manipulation of the doll per se. By combining motor activity and character behavior, *Swamped!* is able to produce a novel kind of user immersion that we doubt could have taken place if, for instance, the user shouted what to do at the chicken.

Discussion

The main argument of this paper, based on our experiences in the four projects described, is that in physically interactive stories, user's immersion and satisfaction can be achieved by creating very responsive physical interactions between the user and the characters, without the use of choice among multi-

ple story paths. This conclusion was reached by observing users while they were experiencing the environments described above. We have not conducted formal studies on user satisfaction, mostly because of logistical issues, but also because of the lack of established, widely accepted methodologies to measure user pleasure or engagement in a story. User testing has been centered on measuring satisfaction toward the completion of tasks, and there is little literature on how to determine how engaging a story is or how pleasing a physical interaction is. Simple methods such as counting laughter during the sessions or making the users answer questionnaires were discarded, since they seem to be able to capture only small fractions of the meaning of "having fun" or "being emotionally engaged."

In the projects described, we experimented with three types of story participation: the user can pre-

Figure 11 Sensor package installed in the physical doll interface used in *Swamped!*



tend to be a character, as in *KidsRoom*, in *PAT*, or when the audience played with *It/I*; the user can be the performer of a character, like the actor in *It/I*; and finally, the user can control (puppeteer) a character, as in *Swamped!*

We have not worked on projects in which the user becomes the master of the story and is able to control the many different characters. We have also not worked on projects in which the user participates in a story as the audience. An example of the latter case is the interaction between the audience and the computer graphics character presented at the SIGGRAPH 99²⁵ electronic theater in 1999.

Although it is not clear whether our argument holds in those “story master” situations, it probably holds in the case of audience participation. In fact, audi-

ence participation in theater has been mostly successful exactly when the performers manage the audience responses in order to keep the story going in a particular direction. The many attempts in theater, especially in the 1960s, of letting the audience control the development and the direction of a play did not prove popular with audiences.

Given that, it is important to examine possible reasons why physical interaction and responsiveness is able to complement the experience in such a way that control of the story seems unimportant. A possible explanation relates to the kinesthetic pleasure associated with moving the body (like the pleasure of dancing). Our conjecture is that by making the body an interface, we can extract aesthetic pleasure and engagement from the variety of responses coming from the muscle actuators and skin sensors to the point that they subsume the need for intellectual pleasure related to determining the developments of a story.

It is important to differentiate this notion of bodily pleasure from the achievement of high-level hand-eye coordination, the foundation of many video games. In our projects, we never had to resort to skill mastery as a way to reward the participants (except, in a different way, in the *PAT* project). In physically interactive stories, the participant’s goal should be to immerse him- or herself as much as possible in the story, i.e., to “live” the story.

The projects described here show how physical immersion can be greatly enhanced by using nonencumbering sensing mechanisms. The spontaneity of the movements of the children in the *KidsRoom* could not be achieved if they had wires or head-mounted displays attached to them. Devices can become a constant (and heavy) reminder that the story is just an illusion, making the “suspension of disbelief” harder to achieve. Furthermore, the gear can interfere with the pleasure of moving which, in our opinion, is a major source of reward in those experiences. Of course, given the precarious nature of both computer vision and movement detection technology, it is necessary to carefully craft the stories so that they can accommodate the limitations of current sensing systems. In fact, we found this to be a prevailing issue when developing the four projects described in this paper.

Although we believe that responsiveness is probably a more important factor than story control in physically interactive environments, we found that

framing our experiences within a story was extremely positive. Stories seem to give a sense of purpose and discovery to the participants, considerably enriching their interactive experiences. In particular, we found a good story resolution, after a story climax, to be very rewarding for the users. Unlike the messages in video games that say “game over,” and thrive on frustrating the user “outside” of the story structure, we found it important to keep the participant “inside” the story up to the last moment and to make the end a natural development of the experience.

Conclusion

“Stories do not require us to do anything except to pay attention as they are told,” claims Murray.²⁶ In this paper, on the contrary, we examined four projects in physically interactive stories designed and built to be experienced by people. Based on those experiences, we made a set of observations that, besides being useful to the design of future interactive stories, seem to run against some common beliefs in the literature of interactive narratives.

First, all the environments were built based on complex sensory mechanisms designed to make the interaction as natural as possible, completely avoiding cumbersome sensing apparatuses such as head-mounted displays or body-tracking suits. By doing so, it is possible to explore using kinesthetic pleasure as an element to reward the participant.

Second, all the environments do not rely on realistic computer graphics or simulations. Most of our characters are cartoonish and in some cases use either low image refresh rates or nonhuman bodies. However, the characters seem to have been perceived as responsive, intentional, and goal-oriented, largely as the result of the combination of their responsiveness and the context of the story.

Third, unlike most hypertext interactive stories, the feeling of interaction is not based on explicit mechanisms of choice but on making the environments and the characters that inhabited them extremely responsive. The sensation of immersion is, in fact, mostly created by the extensive physical activity of the user.

Finally, the four projects realize complete narratives that take the participants through a clear path with an introduction, character and story development, and a climatic end. Our experience suggests that a

well-structured story has the power to engage the users effectively in meaningful interaction.

Acknowledgments

All four projects described in this paper were sponsored by the MIT Media Laboratory. The research projects reported here were conducted while the authors of this paper were members of the MIT Media Laboratory. KidsRoom was developed by Aaron Bobick, Stephen Intille, James Davis, Freedom Baird, Claudio Pinhanez, Lee Campbell, Yuri Ivanov, Arjan Schutte, and Andrew Wilson. It/I was written and directed by Claudio Pinhanez, and produced by Aaron Bobick; the crew was composed by John Liu, Chris Bentzel, Raquel Coelho, Leslie Bondaryk, Freedom Baird, Richard Marcus, Monica Pinhanez, Nathalie van Bockstaele, and the actor Joshua Pritchard. PAT was developed by James Davis and Aaron Bobick, with actor Andrew Lippman. Swamped! was developed by Bruce Blumberg, Michael P. Johnson, Michal Hlavac, Christopher Kline, Ken Russell, Bill Tomlinson, Song-Yee Yoon, Andrew Wilson, Teresa Marrin, Aaron Bobick, Joe Paradiso, Jed Wahl, Zoe Teegarden, and Dan Stiehl.

Claudio Pinhanez was partially supported by a scholarship from CNPq, process number 20.3117/89.1.

**Trademark or registered trademark of The Walt Disney Company.

Cited references and note

1. M. W. Krueger, *Artificial Reality II*, Addison-Wesley Publishing Co., Reading, MA (1990).
2. P. Maes, T. Darrell, B. Blumberg, and A. Pentland, “The ALIVE System: Full-Body Interaction with Autonomous Agents,” *Proceedings of the Computer Animation'95 Conference*, Geneva, Switzerland, IEEE, New York (1995).
3. C. Sommerer and L. Mignonneau, “Art as a Living System,” *Leonardo* **30**, No. 5 (1997).
4. N. Tosa, H. Hashimoto, K. Sezaki, Y. Kunii, T. Yamada, K. Sabe, R. Nishino, H. Harashima, and F. Harashima, “Network-Based Neuro-Baby with Robotic Hand,” *Proceedings of IJCAI'95 Workshop on Entertainment and AI/Alife*, Montreal, Canada (August 1995).
5. A. Bobick, S. Intille, J. Davis, F. Baird, C. Pinhanez, L. Campbell, Y. Ivanov, A. Schutte, and A. Wilson, “The KidsRoom: A Perceptually-Based Interactive Immersive Story Environment,” *PRESENCE: Teleoperators and Virtual Environments* **8**, No. 4, 367–391 (1999).
6. C. S. Pinhanez and A. F. Bobick, “‘It/I’: A Theater Play Featuring an Autonomous Computer Graphics Character,” *Proceedings of the ACM Multimedia'98 Workshop on Technologies for Interactive Movies*, Bristol, England, ACM, New York (1998), pp. 22–29.
7. J. W. Davis and A. F. Bobick, “Virtual PAT: A Virtual Personal Aerobics Trainer,” *Proceedings of Workshop on Percep-*

- tual User Interfaces (PUI'98)*, San Francisco, CA, IEEE, New York (1998), pp. 13–18.
8. M. Johnson, A. Wilson, C. Kline, B. Blumberg, and A. Bobick, "Sympathetic Interfaces: Using a Plush Toy to Direct Synthetic Characters," *Proceedings of the CHI'99 Conference on Human Factors in Computing Systems*, Pittsburgh, PA, ACM, New York (1999), pp. 152–158.
 9. J. Murray, *Hamlet on the Holodeck: The Future of Narrative in Cyberspace*, The Free Press, Simon & Schuster, New York (1997).
 10. R. Schechner, *Performance Theory*, Routledge, London, England (1988).
 11. S. Oettermann, *The Panorama: History of a Mass Medium*, Zone Books, New York (1997).
 12. See http://www.konami.co.jp/am/AM_English/DDR_E/index.html.
 13. R. Pausch, J. Snoddy, R. Taylor, S. Watson, and E. Haseltine, "Disney's Aladdin: First Steps Toward Storytelling in Virtual Reality," *Proceedings of SIGGRAPH 96*, ACM, New York (1996), pp. 193–203.
 14. *Mr. Payback*, Interfilm Technologies (1995).
 15. T. A. Galyean, *Narrative Guidance of Interactivity*, Ph.D. thesis, MIT, Cambridge, MA (1995).
 16. B. M. Blumberg and T. A. Galyean, "Multi-Level Direction of Autonomous Agents for Real-Time Virtual Environments," *Proceedings of SIGGRAPH 95*, ACM, New York (1995), pp. 47–54.
 17. D. Marinelli, "Synthetic Interviews: The Art of Creating a Dyad Between Humans and Machine-Based Characters," *ACM Multimedia '98 Workshop on Technologies for Interactive Movies*, Bristol, England, ACM, New York (1998).
 18. J. W. Davis and A. Bobick, "The Representation and Recognition of Human Movement Using Temporal Templates," *Proceedings of 1997 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 97)*, IEEE, New York (1997).
 19. C. S. Pinhanez, "Computer Theater," *Proceedings of the Eighth International Symposium on Electronic Arts (ISEA'97)*, Chicago, IL (September 1997).
 20. C. S. Pinhanez, *Representation and Recognition of Action in Interactive Spaces*, Ph.D. thesis, MIT, Cambridge, MA (1999).
 21. Y. Ivanov, A. Bobick, and J. Liu, "Fast Lighting Independent Background Subtraction," *Proceedings of the IEEE Workshop on Visual Surveillance (VS'98)*, Bombay, India, IEEE, New York (1998), pp. 49–55.
 22. C. S. Pinhanez, K. Mase, and A. F. Bobick, "Interval Scripts: A Design Paradigm for Story-Based Interactive Systems," *Proceedings of CHI'97*, Atlanta, GA, ACM, New York (1997), pp. 287–294.
 23. A. Wilson, A. F. Bobick, and J. Cassell, "Temporal Classification of Natural Gesture and Application to Video Coding," *Proceedings of 1997 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 97)*, Puerto Rico, IEEE, New York (1997), pp. 948–954.
 24. B. Blumberg, *Old Tricks, New Dogs: Ethology and Interactive Creatures*, Ph.D. thesis, MIT, Cambridge, MA (1996).
 25. B. Blau, M. Butler, and A. Goldberg, "M.C. Leon," *SIGGRAPH 99 Catalog on Electronic Art and Animation*, ACM, New York (1999), p. 108.
 26. Reference 9, p. 140.

Accepted for publication May 11, 2000.

Claudio S. Pinhanez IBM Research Division, Thomas J. Watson Research Center, 30 Saw Mill River Road, Hawthorne, New York

10532 (electronic mail: pinhanez@us.ibm.com). Dr. Pinhanez is a research staff member at the IBM T. J. Watson Research Center and a media artist. He received a bachelor's degree in mathematics in 1985, a master's degree in computer science in 1989 from the University of Sao Paulo, and a Ph.D. degree in media arts and sciences from the MIT Media Laboratory in 1999. His current work focuses on the convergence of computers and arts and entertainment, and on the design and technology of physically interactive spaces.

James W. Davis Ohio State University, Department of Computer and Information Science, 2015 Neil Avenue, Columbus, Ohio 43210-1277 (electronic mail: jwdavis@cis.ohio-state.edu). Dr. Davis is currently an assistant professor of computer and information science at Ohio State University. He received his M.S. and Ph.D. degrees in computational vision at the MIT Media Laboratory.

Stephen Intille MIT Home of the Future Laboratory, 265 Massachusetts Avenue, Cambridge, Massachusetts 02139 (electronic mail: intille@mit.edu). Dr. Intille is a research scientist at the multidisciplinary MIT Home of the Future Laboratory, where he works on computational sensing and future computer interfaces. He received a B.S.E. degree in computer science and engineering from the University of Pennsylvania in 1992, an S.M. degree in media arts and sciences from MIT in 1994, and a Ph.D. degree in media arts and sciences from MIT in 1999.

Michael Patrick Johnson MIT Media Laboratory, 20 Ames Street, Cambridge, Massachusetts 02139-4307 (electronic mail: aries@media.mit.edu). Mr. Johnson received his S.B. degree in computer science from MIT in 1993. He then enrolled at the MIT Media Lab where he received his S.M. degree in media arts and sciences in 1995. Currently, he is a Ph.D. candidate at the MIT Media Lab. His research interests include physical interfaces, character animation algorithms, and artificial life.

Andrew D. Wilson MIT Media Laboratory, 20 Ames Street, Cambridge, Massachusetts 02139-4307 (electronic mail: drew@media.mit.edu). Mr. Wilson received his B.A. degree in computer science from Cornell University, Ithaca, New York, in 1993, and his M.S. degree in media arts and sciences from MIT in 1995. He is currently a member of the Vision and Modeling Group and a Ph.D. candidate at the MIT Media Laboratory. His research activities include work on models for the representation of gesture and human motion, machine learning, and real-time computer vision.

Aaron F. Bobick Georgia Institute of Technology, College of Computing, 801 Atlantic Drive, Atlanta, Georgia 30332 (electronic mail: afb@cc.gatech.edu). Dr. Bobick received his B.S. degree in mathematics and computer science from MIT in 1981 and his Ph.D. degree in cognitive science from MIT in 1987. After serving as research scientist at the Systems Research Institute Artificial Intelligence Center and as professor at the MIT Media Lab, Dr. Bobick has recently joined the Future Computing Environments group at Georgia Tech, where he is an associate professor, an associate director of the Gvu Center, and the Director of the Aware Home Research Initiative. His current research focuses on understanding people's activities from visual input.

Bruce Blumberg MIT Media Laboratory, 20 Ames Street, Cambridge, Massachusetts 02139-4307 (electronic mail: bruce@media.mit.edu). Dr. Blumberg is an assistant professor at the Media Lab

and head of the Synthetic Characters Group. He was the creator of Silas T. Dog, an autonomous animated dog for virtual environments, and one of the chief architects of the ALIVE project. His doctoral dissertation *Old Tricks, New Dogs: Ethology and Interactive Creatures* was done at the Media Lab under the direction of Pattie Maes. Prior to coming to the Media Lab, he held positions with Apple Computer, Inc., and NeXT Software, Inc. He received an M.S. degree from MIT's Sloan School in 1981 and a B.A. degree from Amherst College in 1977.