

Glorianna Davenport
1001 Electronic Story Nights: *Interactivity and the Language of Storytelling*

Thank you, Michael. Your remarks fell on live ears. I would also like to thank Cathy Robinson and Sally Dray. Ever since one of my graduate students traveled to Alice Springs to experience the eclipse, I have fantasized about coming on my first trip to Australia. You three, as well as your staff and sponsors, have made it possible.

This conference focuses on interactivity. I have worked with interactive cinematic projects since 1980. In this talk, I will discuss some of my current thinking about "the language of interactivity," and show you some of the recent work we have been doing at the MIT Media Lab. I am not concerned that you understand every detail about the inner workings of these pieces -- some of them will be available out in the lobby later for your examination. In discussing these examples, I will emphasize general features and concepts. If any of you have burning questions during this "show and tell," wave your hand around and I'll try to take the occasional question.

We're here to celebrate change and the new opportunities for expression which are arising out of new, enabling technologies. As we witness the evolution of these technologies, we see that is that they are moving us toward systems which can learn. Several ancillary technologies are also especially important, in that they measure and can manipulate user input. For example, new sensor technology will help a system sense the presence and activities of an audience without requiring them to actually handle an input device. One type of sensor responds to the presence of the small electrical currents which typically circulate through the human body. This type of input can provide a feedback loop amongst an audience, story materials, and a sequencing or storytelling engine.

In this slide^{pic1}, we see Associate Professor Pattie Maes, whose research focuses on computational agents, sitting in the middle of a virtual space chatting with a dog. The dog, Silas, is an autonomous agent; he has been programmed in 3D to exhibit high level autonomous behaviors: Silas feels hungry, Silas searches for food; Silas drinks some water, Silas needs to pee, identifies an upright architectural feature, and takes a leak. We will revisit this creation of Ph.D. candidate Bruce Blumberg later, in the context of a particular story.¹

For the moment, we use this example of an interaction between the dog, the virtual world, and Pattie to discuss a vision to which we must aspire when discussing the "language of interactivity." Ron Evans, a native American storyteller, clarifies this vision in a story he tells about the chief of an African tribe and a missionary who

pic1 Pattie Maes and ALIVE, Autonomous Agents Group, MIT Media Lab, 1994. Copyright @ MIT Media Lab

¹Bruce M. Blumberg and Tinsley A. Galyean, "Multi-Level Direction of Autonomous Creatures for Real-Time Virtual Environments, Computer Graphics Proceedings, Annual Conference Series, 1995. pp. 47-54.

visits with them from time to time. On one such visit, the missionary brought along a television as a gift to the chief of the tribe. When he arrived at the village, he presented the gift in a ceremony which was accompanied with appropriate pomp and circumstance. The gift generated great excitement in the village. Every night, the chief turned on the television set and the whole village stood around it and watched the stories that were coming out of the set. (We need not concern ourselves with the source of electricity, and we can imagine that they were able to tune into a transmission from a well-positioned satellite.) Several days later, the missionary bid farewell to the chief and continued on his journey. Six months later, the missionary returned to the village, where he discovered the television set was nowhere to be seen. The people of the village no longer gathered around the television each evening; instead, they gathered around their tribal storyteller. The missionary, somewhat baffled and hurt, went to visit the chief. With agitated voice and gesture, he asked the chief what had become of the television. The chief calmly replied, "I listen to my storyteller; he tells many stories." The missionary pressed the point, "But the television set, it too has many stories." The chief nodded wisely and responded, "Ah, but my storyteller knows me." ²

Until the machine can understand story and synthesize new story elements, effective personalization will require close attention on the part of the story's author. To achieve personalized delivery, the human storyteller must invent a fabric rich enough to accommodate many pathways through a particular story space. These pathways reflect the personal interests and attention of diverse audiences. As Michael [Hill] has already mentioned, publishers often fear that the diversity of extensible storytelling greatly escalates the task of production. Like many cinematic producers before them, these publishers may trade-off timely completion against expressive invention.

Yesterday, John Collette hosted my visit to the Australian Film, Television and Radio School -- I've heard about this school for many years, and found the facilities quite remarkable! However, what I most appreciated was John's enthusiasm concerning the new media. No chance for the old media to grow stale with this level of energy! John dreams of bringing all kinds of computers into the school. As we talked about digital production, John and I discussed his reservations about building this new medium on top of an older, well-established medium. Michael [Hill] has made reference to this as well. Many years ago, I named my group Interactive Cinema. At the time, the low bandwidth media -- text, still picture, and sound -- were being rapidly assimilated into the computational language. My concern at that time, which continues today, was how to bring this computational approach -- which enables personalization -- into the high bandwidth arena of cinematic storytelling.

How do we work in this new medium, which seems so different? This medium is participatory and democratic; as makers, we must respond to these attributes. The medium supports distributed connectivity, which changes the demographics and the very experience of audience. As John and I considered the topic for today's discussion, I was inspired to reformulate the story which Steven Hawking tells in the opening of *A Brief History of Time*. As Hawking tells it:

² Glorianna Davenport and Brian Bradley, 1001 Electronic Storytelling Nights, in manuscript

A well-known scientist (some say it was Bertrand Russell) once gave a public lecture on astronomy. He described how the earth orbits around the sun and how the sun, in turn, orbits around the center of a vast collection of stars called our galaxy. At the end of the lecture, a little old lady at the back of the room got up and said: "What you have told us is rubbish. The world is really a flat plate supported on the back of a giant tortoise." The scientist gave a superior smile before replying, "What is the tortoise standing on?" "You are very clever, young man, very clever," said the old lady. "But it's turtles, turtles, turtles all the way down!" ³

In response to John's anxiety about mixing the new and the old, I suggest that the old lady was not far off. As digital storytelling emerges, we will discover that we are not dealing with galaxies of bits, but rather that we stand on the shoulders of great storytellers through the ages, and that from today forward it will be story, story, story all the way down!

Last year, this conference concerned itself with Narrative. I believe that we must understand the premise of narrative before we can examine the "language of interactivity." From an historical perspective, major advances in civilization may be attributed to the need and desire of human beings to share stories. Spoken language was perhaps the first, truly revolutionary advance of civilization. To begin, any particular set of utterances must have been shared with only a few people who were situated at a single, very small geographic location. Language flourished with the "glory that was Rome." By decree of the Roman emperors, all inhabitants of conquered lands were required to adopt Latin as their official language. As common language spread, travelers were able to share stories which contained valuable news -- the location of a drought, war, disease. The content of these stories enabled a limited capability for prediction and was often critical to the survival of the audience. Today, spoken language remains the most powerful force in securing and extending community.

While spoken language helped build community, spoken communication lacked durability and extensibility. Without a common language, how did the traveler share his message with the community into which he had wandered? Some stories were communicated by drawing on walls; other stories were shared by means of pictorially expressive gestures combined with what was probably some unfamiliar sound sets. Pictograms often lasted longer than the spoken word: generation after generation could visit the rock cave and view the durable images, which carried meaning served some special need. Slowly, pictorial signs were codified into symbols. These symbols, eventually, were formalized into an alphabet system. The abstraction of language provided a collection of signs which could be recombined, according to rules of adjacency, into groups of letters representing spoken words -- an interlinking of thought, vision, and sound. With the invention of written language, verbal communication gained a permanence and an interoperability. Written language enabled the poets, priests, scribes, and others who acquired the

³ Stephen Hawking, *A Brief History of Time: An Interactive Adventure*, Voyager

facility to write down their stories. Permanence insured an enabling access to myths and moral tales which was not limited to generational or parochial continuities.

It took thousands of years for the written language to be embodied in a mechanical device for the purpose of making many copies. Movable type was first invented by the Chinese over one thousand years ago, but it was the Europeans, not the Chinese, who applied the technology of movable type to the printing press, allowing the printer's design of a written artifact to be reproduced and distributed in lots of hundreds or thousands of copies. Centuries later, the word processor placed the capability of the typesetter onto the desktop of offices and homes.

A 19th-century invention, photography, introduced a new paradigm for replication. In the case of the photograph, the image-making engine -- the camera -- is itself a mechanical device. The human mind, eye, and body work together to propose and capture the intended image. The photograph mirrors the world, not as we perceive it, but as it has been composed, framed for posterity. Our rereading of the image is governed by its composition, its form, and the size of the display. Generally speaking, the size of the display governs our immersion in the image.

The printing press changed the way in which the storyteller perceived her audience. As soon as many copies could be printed, the idea of the "mass" audience took hold. the press-owning entrepreneur had to grapple with the notion of distribution channels. As the commercial enterprise of print grew, the client of the press was often a commercial or political entity, rather than an individual reader: advertising increasingly subsidized the cost of newspaper printing, and soon, the expensive private newsletter soon gave way to the "penny press," affordable to all. However, distribution of identical copies of news on a daily basis required stylistic and formal coherence in the production process. The invention of radio deepened the rift between authors and audience. Whether equated to copies or to transmission, stories for the "mass" audience required a distribution channel which was generally owned by someone other than the person making the media content.

As the "mass media" of print, radio and television grew and took hold, the telephone was an innovation which stood apart. Invented in the late 18th century, it took 40 years for the telephone to become recognized as a household technology. In its infancy, nobody believed how person-to-person conversations over shorter or longer distances would revolutionize the world. The thrust of a personalized medium appearing amidst the mass media will be truly understood only in the future. For the moment, using ourselves as subjects, we occasionally and unscientifically have discerned moments -- times of political stress or in the creation of media idols -- when the worlds of mass media and personal communication came into a profound (and profitable) symbiosis.

Interactivity brings pressure to bear on the channel. In order to make the telephone efficient, switches had to be automated. Today, we must design into our media objects a feedback-channel which can elicit and make use of signals from the audience -- these signals from the audience will control the "automated switches" of future storytelling systems to some extent. In the case of early video-on-demand trials, the back-channel allowed each viewer to request a particular movie from a fixed menu of choices. In this case, the back-channel was minimal. The wider the

bandwidth of the back-channel, and the more distributed the system, the more the audience can contribute to the program. Again, it was interesting to talk to John Collette yesterday, because John does not believe that interactivity is about information. I can buy that, because I believe that successful interactivity is really about story. However, I believe, that this is an issue of semantics and emphasis. Ultimately, interactivity engages us in the assembly and construction of story, but in the process the information bits -- whether these are the program proper or a trail of user activity -- are central to the endeavor.

Even as we are trying to understand it, the digital universe is changing. Distinguishing information from story will become more and more difficult as the lines between storage, program, content, medium, and interface become more transparent. Take, for example, the problem of text and interactivity. Today, a limited form of hypertext has taken off. The "hot link" brings us to more information. Often, this results in our getting lost and forgetting why we were reading to begin with. This raises two problems: one is the issue of effective presentation, and the other is the problem of the memory trace ("Where have we been, where are we going, and who has been here before us?"). Text poses a difficult problem because there is no standard temporal dimension to text. We are used to the page format, a 2-dimensional expanse of text; however, reading large amounts of text on an electronic screen is not a particularly pleasant experience. Over the years at the Media Laboratory, we have dreamed of a flexible surface for electronic text, one that you could carry with you in your back pocket, hold however you choose, and read at your leisure. The goal of reusable, flexible digital "paper" has recently been taken up by Joe Jacobson, a young physicist on the faculty of the Media Laboratory⁴. Current progress includes the invention of an ink substrate which supports a heat-reversible process, where particles turn from black to white or white to black at certain temperatures. We imagine a future in which you can stop at a kiosk on your way to work and down-load the next section of the news onto your own personal "super paper." This is the sort of radical project which challenges many of our underlying assumptions about the electronic interface.

Before delving in to the relationship of story to interactivity, we need to examine the notion of interactivity itself. What does interactivity really mean? The fact that I can take my piece of "super paper" and recycle it, fill it with news which is meaningful to me, poses a physical as well as an electronic model of interactivity. At the physical level, I pull the sheet of "super paper" from my back pocket; I introduce it to the system; I retrieve it from the system; I hold it up and scan its contents. These interactions with the object provide a practical method for achieving my goal, catching up with this morning's news. However, interactive in this context also includes the way in which electronic sub-systems interact: the information and the layout can be affected by my personalized user profile. Once the paper can track our eye movements, we can go even further in personalizing the selection of material and its layout. In addition, in a networked electronic world, the "super paper" can actively link together a larger society of audience, networked communities which share observations with other people and with programs.

⁴Peter J. Howe, "MIT book would bind computer 'ink' to paper, Boston Globe, March 17, 1996, p.31.

In 1992, the idea of community led me to extend the scenario for interaction beyond the notion of a single user on a single machine. I began exploring a scenario which encompassed two users: one acting as an Explorer and the other as a Guide. In fiction, many stories juxtapose the journey of a main character with advice from a guide or with difficulties imposed by an obstacle character. Guide characters help the main character forge ahead, while obstacle characters inhibit her progress, often by requiring some level of accountability. As we explored the interface, we focused on several considerations: the scale of the interaction, the relation of the audience to a story goal, and the narrative structures which would act as an invitation to the audience.

I was contemplating the problem of how to move the interactive experience out of the desktop video "box" and into the room as a direct, human-scale experience. I invited Larry Friedlander (a Professor of Literature and Theater at Stanford University and an accomplished Shakespearean actor) to take a sabbatical and join me in my work at the MIT Media Laboratory. In 1992, we co-taught a workshop in which students and faculty collaborated on an creation of an installation piece: a series of large-scale, walk-through participatory environments which we called "The Wheel of Life."⁵

The "Wheel of Life" was developed for and installed in "The Cube," a 60 x60 x50 foot open room in the center of the Media Laboratory. The exploration spaces, inspired by the rich symbology of the Buddhist Mandalas, were implemented on the scale of a small theme park. The Guide stations invited visitors to collaborate with the Explorers' efforts by sending non-verbal messages to them or by solving some related, parallel puzzle. The work was open to the public for 10 days in January 1993.

As I have worked with students over the past decade, I have evolved a collaborative workshop method which allows us to create works which push the edges of expression and the new digital technologies. In this work, I focus less on longevity and issues of distribution than on general principles, with the practical goal of prototyping work in a short period of time. The "Wheel of Life" was particularly impractical in regard to distribution. It was built by some 25 students, faculty, and staff in one 13-week semester, which was a challenge in and of itself. The spaces had to be large. For *Water* and *Earth*, we hung huge, sculpted scrims from the ceiling to define and enclose their spaces. *Air* invited the audience to walk into a large mylar balloon, the size and shape of a Quonset hut. Each space was constructed with its own cosmology of experience. As you walked into the *Water* space, a short movie rear-projected onto an overhead screen conveyed the impression that you were being released from a huge fist into an expanse of water; immediately, schools of colorful fish began to swim by on monitors around you. Very quickly, you discovered that you were sharing this space with a large whale whose florescent pink throat you could walk into. The Guide sent you messages, such as images of a child's fingers pointing and an accompanying whisper, "move left," "move right," "speak to me." The goal in this space was to get the Explorer to

⁵ Glorianna Davenport and Larry Friedlander, "Interactive Transformational Environments: Wheel of Life," in Contextual Media: Multimedia and Interpretation, Edward Barrett and Marie Redmond, editors, MIT Press, Cambridge, Mass, 1995.

sing into the ear of the whale. Adults resisted this particular interaction because the ear was way down near the floor. If you did indeed sing into the whale's ear, you were rewarded by an outstanding light show and the memorable lines from the Tempest, "Full fathom five, my father lies..."^{pic2}

In the *Earth* space, the Guide sent a cryptic request to the Explorer, who had to decipher the request in order to move on through the space. *Earth* was inspired by the Percy Bysshe Shelley poem "Ozymandias" -- the Explorer entered the site to discover it in ruins; the Explorer had to solve three puzzles to reconstruct it and experience its former glory.

By exiting *Earth*, the Explorer entered the *Air* environment: the interior of a large, inflated mylar balloon representing a space ship in crisis. The fate of the Explorer was critically dependent on the collective work of five crew members currently in a state of drunkenness. By standing close to an individual crew member, the Explorer was able to awaken him briefly.^{pic3} However as soon as the Explorer left, the crew member fell back into stupefaction. In order to succeed, the crew members had to be revived in a specific order within a specified period of time. Since the space was quite large, this task involved considerable movement on the part of the Explorer. Meanwhile, the Guide played what was ostensibly a computer game; the task was to launch a collection of logs down a moving stream, one by one, in proper order to form a bridge. Part of the Guide's challenge was to discover exactly what this "proper" order was: once the rules became clear, the Guide could move down to the colored squares at the bottom of the screen. These squares directly controlled the lights within the spaceship. If sequenced in parallel with the moving logs, these squares showed the Explorer the order in which to revive the crew members. Thus, there was a necessary collaboration between the Guide and the Explorer.

The "Wheel of Life," while not technically perfect, was a grand success in terms of its fanciful spirit and the fun of its interactivity. The limited speed of the network and inadequate sensor technology did present some major problems. As I stated earlier, there is now a group at the Media Laboratory which is actively pursuing new technologies for sensing human interactions. One class of sensors responds to the electrical currents which flow within our bodies. We can use this technology in a cello and a violin to precisely measure the musicians' movements; through their actions, the musician can control their accompaniment, the background orchestration of music in a solo performance. Similar sensors can be used in other venues where the human wishes to control computer output. This focus on sensors has lead to significant work by Neil Gershenfeld and the Computers and Media Group at the Media Lab.⁶

^{pic2} Water Environment and Guide Station, Wheel of Life Installation, MIT Media Laboratory, January 1993.

The explorer plunged into the water world and down to the bottom of the sea. The guide sent messages encouraging the explorer to commune with creatures of the deep.

^{pic3} Air Environment, Wheel of Life Installation, MIT Media Laboratory, January 1993. The explorer enters a space capsule. To save the ship, she must trigger a change in the state of the crew.

⁶Tom Zimmerman, Joshua R. Smith, Joseph A. Paradiso, David Allport, and Neil Gershenfeld, "Applying Electric Field Sensing to Human-Computer Interfaces, Proceedings of CHI-95.

The Explorer-Guide relationship comes in a variety of flavors. In the "Wheel of Life," we were concerned with creating a theatrical experience within physical space. The Explorer might not even recognize that the Guide is present and playing a role. But the principle of such a relationship can lead us into some very rich emotional situations. Consider the story told to me by a friend (who is also a new media critic). David had become very discouraged about the limited emotional range provided by interactive experiences, and frequently complained that the medium would never succeed until we can trigger aspects of surprise and delight in a viewer. However, one day a flushed and excited David related to me how he had reached a sort of interactive epiphany. He had logged onto a MUD — does everybody know what a MUD is? Okay, I don't have to explain it? I do? In brief, a MUD (or "Multi-User Dungeon") is a text-based, networked experience which can simultaneously accommodate any number of participants. It is generally constructed using a spatial mise-en-scene: visitors can create new rooms or enter existing rooms in which other people on the network are "hanging out." Participants visit MUD's in search of semi-anonymous social interaction, where they are completely free to be themselves -- or to be someone else, playing a fantasy role. Interestingly, MUD's are often peopled by programs which simulate humans, as well as by actual humans. MUD's are slowly evolving into graphical worlds where a participant will be represented by an avatar, a sort of audiovisual surrogate for yourself which dwells within the synthetic space and acts under your remote control. Perhaps you have heard of them. What is most important about a MUD is that it provides a social venue for a community of players. Returning to the story, my friend, David, entered a MUD on the day in question. He journeyed into a bar. A group of other people were standing around the bar talking, and there was a robot in the room, standing all alone. The robot seemed sad and lonely because nobody was talking to him. David, being a sociable fellow, went over to meet this robot. Instantly, the robot seemed to cheer up and began to share his stories with David. In retelling this experience, David observed that at some point he was surprised to notice that he had placed his arm around the robot's shoulders. Shortly after this, the robot asked David to tie him up. What would you do if a robot in a MUD asked you to tie him up? David considered the proposition and concluded that he had established a trusting relationship with this character and that, in order to be true to their friendship, he would honor the robot's wishes. However, no sooner had he finished tying the robot up, then the robot turned and furiously berated him: "How could betray our trust, how could you capture me and tie me up!" (All of this was transpiring in text, of course). David was stunned. He claims that he did not know what to do next. Everyone in the bar had followed David's recent interaction with the robot. They had all witnessed the robot's temper tantrum. They sympathized with the robot, and unanimously turned on David. David felt at once betrayed and emotionally lost: he claims that the feeling stayed with him for days. I know this sounds ridiculous, but it shows that the person-to-person connectivity over the network -- and particularly, the identification with yourself as a first-person character or player -- introduces a the potential for a truly intensive emotional interaction.

Now, I'd like to show you a video. We will need to turn the lights down, but first let me set the stage. What you will be looking at is the visual portion of the Ph.D. thesis of Tinsley Galyean, one of my students at the Media Laboratory. The story, "Dogmatic," runs on a SGI Reality Engine. It is a one-on-one experience in a 3D virtual world. You control your point-of-view by rotating a mouse in front of you in

3-space. Now, I need everyone in the audience to feel that you have the mouse in your hand and you are looking around the space. The reason that this part is so jiggly and uncinematic is that the active user is learning that she can control where she looks within the synthetic scene by moving the mouse. In order to have any sense of the experience, you must pretend that you are controlling the view shown in the videotape. You're looking around the environment. You hear a car; you try to keep the car in sight. The dog jumps out of the car and pads toward the cactus. If you had taken your eyes off the dog, he would have come over and barked at you to get your attention before he takes a pee. The movie continues: the dog brings a severed human arm over to you: clenched in its fist is a note, "Lucky Strikes." Later, there is a fight; you hear the car screeching out of control, coming towards you, you black out and come to. You see the world on its side.^{pict6}

This *film noir* is experienced in a virtual environment with a virtual character, the dog, who exhibits some high-level autonomous behaviors. In an interactive story, the actions of the user must have consequence. In this experience, you turn toward the sound of the car; this triggers the action which results in your catastrophic downfall. "Dogmatic" highlights a range of research issues for interactive narrative. The piece challenges the continuous flow of space and time which has become the tradition of virtual reality experiences. I think "Dogmatic" shows us that cinematic language can enhance a virtual, first-person story. In this piece, the cuts are all executed under programmatic control: the program detects where you are looking and determines when to cut away, and for how long. For some short period of time after the cut, we take away the real-time control over point of view; this allows us to normalize your new view of the world. The effect at the end of the experience startles the participant viewer. Almost everybody who experiences this work in first-person mode is absolutely wiped out when they realize, after the crash, that they no longer have physical control over the world. They cannot sit up. The emotional response to this ending is similar in effect to David's rejection in the MUD, which I mentioned earlier. Both scenarios reveal the real consequences of your actions, which in turn generate emotional investment and reaction.

The irony of the ending is nicely revealed. The fact that the viewer can look around the world gives the participant viewer the false impression that she can control the world. In reality, the viewer has very little control. The program has a fixed content which may be triggered by the viewer's actions; however, the viewer is not aware of the exact relationship between their act of looking around and story. Most viewers do not realize that they cannot walk down the street or enter one of the venues of the desert town. From the author's viewpoint, constraining the viewer's movement placed realistic constraints on the author. We did not have to juggle the implications of multiple *noir* endings: you can only die once.

"Dogmatic" required a musical accompaniment which could adapt to the pacing of the participant viewer's interaction. The music is written in a somewhat ditsy, plinky voice because it must be terminated or extended depending on viewer action. Dialog is totaling missing from the piece. If we are going to invite dialog, then the storytelling system needs to have some understanding of what the viewer might say

^{pict6} Tinsley Galyean, The set for Dogmatic, a virtual film noir experience. The participant, as main character, views the world and is confronted by an unusual antagonist.

and respond appropriately. This is a big research area, and a comprehensive solution is unlikely to appear in the immediate future. By pointing out certain limitations, I would like to argue that the storyteller system is not a magic box which already exists; rather, it is a system which is designed for (and around) particular interactive stories. In order to make such a magic box, there needs to be a plan, a taxonomy, a knowledge of relevant issues and past solutions.

This brings me to another point related to the language of interactivity. Can we define types of interactive engagement? If so, how might we embed these activities into a narrative experience? In short, how can we invite the viewer to participate? Jean Piaget, a Swiss researcher, studied children's game play in the 1960's and was able to identify four discrete types of activity: practice, pretend, playing by the rules, and construction. These activities are often combined. For instance, hopscotch combines practice with a rule base, *Doom* combines practice with a goal-driven pretend experience. Today, the designation of the "twitch" or "reaction-timer" type of computer game pejoratively refers to a particular class of practice games. These games require manual dexterity and concentration at a task level; success is unambiguous. The shoot-em-up actions which dominate practice-based computer games today can not easily be replaced by complex-force models; to the extent that there are rules, they only weakly associated with a modeled world. The main impediment to transitioning to a richer content-base is the public itself. Publishers depend on sales; shoot-em-up games are startlingly popular, particularly for males in the 12- to 18-year-old range.

Pretend scenarios allow us to develop effective interpersonal skills. Stories teach us about temporal structures, moral frameworks, language, and communication. Children often create imaginary worlds to play in. In a "good" movie, the creators sculpt an experience which invites empathetic transference by the audience. A critical attribute of computer-based virtual reality environments is the immersive nature of the experience. Actions are executed by the viewer as a first-person character. This brings us to the subject of voice and interaction; interaction is intrinsically a first-person act. Last year, we created an interactive story, "Lurker." The story is played out over the network to and by an audience of six people. Prior to the start of each story, six people sign up as Lurkers (<http://lurker.www.media.mit.edu/registration>). In the act of subscribing, our society of audience accepts a pretend role of being a Lurker. Lurkers assist the hackers. The hackers fall into tense times when one of their group disappears. The Lurkers are called in to assist. As part of your assistance mission, you must practice some difficult computer skills.

Games with rules constitute a particularly rich class. Many of us grew up playing action games, such as hopscotch and hide-and-seek, as well as board games such as Monopoly and Scrabble. Sports are also rule-based games which incorporate the practice of forming and executing strategies. In sports, the goal and the basic skills are known quantities; the challenge lies in how you accomplish this goal. Certain computer activities -- most notably MUDs -- combine rule-based game play with an important element of pretend or role-playing. In a MUD, you play at being a character; the pretend element is minimized when you play yourself, but posturing and selective self-revelation are commonplace pretend elements even then. As with

practice games, rule-based games are often designed around winning and losing. MUDs suggest a new direction in that they focus on socializing.

Finally, interactivity is built into construction-based activities. For example, the activity of movie editing is constructionist by nature. For many years, Seymour Papert, who worked in the 1950's with Jean Piaget, has been actively engaged in embedding the activity of construction as a mode of learning in schools around the world. He champions what he calls "hands on and heads in" learning, which includes the social aspects of learning. This combination of active doing and social interaction effectively defines the potential of new media.

In post-modernist theories of literature, critics concern themselves with the way in which the audience constructs meaning. In reading a book or watching a movie, the audience first deconstructs or parses their experience into the smallest possible elements of meaning. Allowing for a certain latency, the audience then reconstructs these elements in higher-level groupings. This process of granular reconstruction mimics the activity of the original creator. One of our recent projects is designed around the idea of construction. Our subject is historical biography; the title is "Jerome B. Wiesner: a Random Walk through the 20th Century."

At one level, this research is about the nature and expression of history. Our lives are reflected in and through the fabric of multifaceted chronologies. Events which are common between cultures are also distinct to each culture. Increasingly, I feel that we, perhaps especially in America, are losing our history. As the turmoil of the 1960's fades, I watch students enter MIT who have no knowledge of the Second World War and its impact on America, let alone the Civil Rights Movement and Vietnam. They are students of the Now, but it is hard for them to put their generation in context. Although this may not inhibit their ability to write computer programs in the years to come, it cannot help but leave their world less rich. In addition, it may limit their ability to realistically evaluate strategy, occurrence, and consequence. World War II left its heavy mark across a cultural framework for several generations. Remembering the circumstances of the war and the Cold War which followed can help us formulate a context for our current history. However, if we restrict ourselves to some thin causal chronology of the times, we cannot understand the global cultural trends that have emerged. A contextual approach to history allows us to examine the actions and interactions of individuals as more or less innocent agents of change. We followed this approach in our portrait.

Jerome Wiesner was born in 1915. He served as President John F. Kennedy's Science Advisor from 1961 until Kennedy's assassination in 1963; subsequently, Wiesner became President of the Massachusetts Institute of Technology. When Wiesner died in 1994, I became interested in making a biographical piece which would reflect some of the strengths which had allowed Wiesner to affect the community of MIT, as well as the course of American -- and hence, global -- history. Five circumstances influenced my commitment. I knew Jerry Wiesner fairly well from my time at MIT, and I had filmed and edited several short pieces about him while he was still alive. I was fascinated by the story and had already collected several key pieces of visual material which featured Jerome Wiesner. I had become friendly with Cheryl Morse, his assistant at the time of his death, and she was willing to participate in this project. Jerry Wiesner had written some compelling

chapters for an incomplete autobiography; we had some access to these, but not publication rights. Finally, my own research in the development of digital story engines for documentary narratives almost certainly would benefit from the richness of this story.

The goal, then, was to develop a biographical piece around the theme of individual influence. What method of engagement, what style of thinking allowed Wiesner to become so influential? How did he formulate and negotiate his commitment to change within the society? What could his life lessons tell us about our own journey forward? In all of this, I was less interested in conveying a sense of historical causality than in probing individual consequence. The roots of this approach evolved out of my own growing-up. In my experience, conviction about the causality of events often shortchanges the type of dialog which is so necessary to learning. In moving away from a causal presentation, it was useful to present an individual portrait of action, reaction, and interaction as it occurred over time and within a distributed context of sociology and psychology of a time. I knew Jerome Wiesner well enough to understand that he respected and engaged human capability at home and abroad. What would we discover by listening to circumstances of his interactions as told by individuals who had been on the other side of the interaction? [pic7]

This project was published as a WWW site with a companion CD-ROM as part of the tenth anniversary of the Media Lab, which Jerome Wiesner co-founded. The piece presents itself to the participant viewer via a conceptual map designed by Michael Murtaugh, a graduate student in my group. The concept map provides the participant viewer with an explicit and dynamically suggestive navigational scheme. The visual structure of the map is derived from the wall of the Media Lab's atrium, which was designed by the painter Kenneth Noland. The images and text contained in this map comprise the complete set of annotations which we used to describe and enhance the video sequences and text documents contained in the portrait. Any piece of media can have multiple annotations. The interface mediates the participant viewer's journey by providing a layered visualization. The viewer can select from a period of Wiesner's life (arranged as a downward diagonal on the grid), a decade, the cast of interviewees, and a set of 8 keywords. As the viewer selects from this palette of choices, the concept map indicates other annotations which will lead the viewer to relevant or related content of interest. The system jumps immediately to content when this is the only possibility for closure. This method of narrative guidance through annotations encourages meaningful choices and some amount of narrative continuity and coherence in the viewer's experience. The initial version of the project contained 53 first-person stories told by 28 of Wiesner's friends and colleagues, as well as a collection of letters and autobiographical writings.

In this experience, the viewer can begin exploring the content from any annotation, according to their knowledge and interest. Someone who has no idea who Jerome Wiesner was might begin with "Science" or with the period of his life connected to

pic7 Glorianna Davenport, Cheryl Morse, Michael Murtaugh et al, "Jerome B. Wiesner: A Random Walk through the Twentieth Century", Interactive Cinema Group WWW server, MIT Media Laboratory, 1995-6. [<http://ic.www.media.mit.edu/JBW>]

"Kennedy;" someone with a knowledge of MIT during the 1960's might begin with "November Actions," a period of tremendous upheaval on campus which included street rioting and the student's takeover of the President's office. If I select "Science," the system will highlight in red all periods of Wiesner's life, as well as many of the interviewees. As soon as I select the second period of his life, the "Research Lab for Electronics," the system discovers closure and streams a QuickTime story by Professor Jerome Lettvin. In the story, Lettvin describes an encounter he had with Jerome Wiesner at MIT's Research Lab for Electronics in the 1950's.

Red Howland and I developed an anti-wiretap device. Now they're trying to find the original documents because all of a sudden it's come back into fashion in a different way. So anyway, it worked, it was wonderful and we were so short of money at the time... it was really bad. So I got in touch with some bookies in NY and they said, "Yeah, we need this," because the nice thing about it was that you didn't have to have two... just one was enough. The speaker turned on the noise when he wasn't... speaking, and -- I really think it was interesting -- he could get the whole message coming in, but everybody including me heard only the noise. So we were delighted and the bookies said they'd give us 10,000 bucks so we said, "OK." Red and I were looking forward to it. And it was Saturday, we went Saturday to pick up the money and, ah... but I said, "Maybe you should probably tell Jerry what we're doing," so we did it. Jerry was very interested and said "That's a very intelligent way to do that, very good, very good" And then Friday afternoon, Jerry calls us in Brooklyn, there's a colonel, a major. And he said, "Would you explain this to us?" So we explained it. At the end of which time he said, "You have a document?" I said, "Oh yes." And he got out the stamp and "top secret, top secret, top secret," and he gave it to Jerry to put in the file. And Jerry said, you know, we sadly just... "I'm afraid you're gonna have to tell your clients it didn't work."⁷

We could have started anywhere. Each edited clip builds up the viewer's impression of the person and of the historical time. The viewer is able to consider the relationship between individual style and the progress of transforming societal opinion. We have incorporated a mechanism by which individual viewers can contribute to the discourse. Aimed at developing a society of audience, a firmer idea emerged as colleagues at the Media Laboratory viewed some of our early sequences. Many of our early viewers knew Dr. Wiesner and had lived through particular portions of the history. When they would watch a clip, particularly one about Vietnam, they would comment, "Oh I was there! Can I add my story?" Therefore, we incorporated a discussion component on the World Wide Web. Occasionally, we discover something about the language of interactivity from observing our mistakes. Initially, the discussion thread was open and unedited; however, this was not well-used, so we introduced an editorial function. This form still feels stiff and does not elicit the type of commentary which we had hoped for. We need to continue to explore the design of community-generated commentary.

⁷WWW Java Interface Design by Michael Murtaugh for "Jerome B. Wiesner: A Random Walk through the Twentieth Century", Interactive Cinema Group, MIT Media Laboratory, 1995-6 [<http://ic.www.media.mit.edu/JBW>]

Working backwards from the end-user application of Jerome Wiesner, we can look at a demonstration program which offers a compelling visualization of a content base annotated by keywords. "ConTour,"⁸ a visualization program written by Master's Candidate Mike Murtaugh, provides a mechanism for exploring a media database, given a keyword approach to annotating media objects. In this program, the story engine dynamically sequences video clips based on a dynamic weighting of what has already been played. This mechanism insures the appearance of continuity. Using the video included in the Jerome B. Wiesner project, we can begin by selecting the descriptor "Personal Style," for example. The relevant descriptors grow dynamically with the playout of a segment. Therefore, if we link "Personal Style" and the "Washington Years," the system presents us with a segment from Amar Bose. From here, the system will look for the combination of descriptors or, failing that, one of the two descriptors. This might cause the "Washington" descriptor to shrink because it no longer applies. "Personal Style," on the other hand, remains large. At any point the viewer can "steer the story" by selecting a keyword. The dynamics of this interface permits a critical visualization for content; it invites rather than requires our participation, while emphasizing content-based continuity.

In addition to exploration, interaction can promote emotional transference and role-playing by the audience. In a recent piece, "Lurker," the audience is situated as a supporting character in the story. In cyberspace, a Lurker is to a hacker, as, in the music world, a groupie is to a star. "Lurker" is run on our World Wide Web server once 6 people have volunteered to participate in the story. Shortly thereafter, the Lurkers are given a call to arms: a hacker has disappeared; they can help find her. The Lurkers work collaboratively as a "society of audience."

Written in 1995 by a former student, Lee Morgenroth, "Lurker" is an example of a larger genre which we call "Thinkies;" a "Thinkie" challenges the audience to think in a style which is appropriate to their circumstance in the narrative. Through this experience, we believe that the audience can gain new skills for problem-solving. Normally, this project takes five days to play out; events and postings of materials are released in a timely way, controlled by a clock. After the audience/player registers for the game, she is assigned a pseudonym and invited to explore the Toad Sexer's Pad; one section of the pad contains all the WWW home pages of the characters. Simultaneously, they begin to receive e-mail from the hackers and from other Lurkers. Soon they receive a call to arms from the hackers. One of the hackers, Shira, has disappeared; the Lurkers are asked to help the hackers find her. Because Bippy wears a head-mounted camera, the hackers are able to post a video which shows Shira disappearing down a catwalk. In the process of unraveling the situation, the audience members perform a series of more or less difficult processing tasks on the computer. For instance, in the original version, you install and learn to use a PGP encryption code on your computer. Installing it requires a "hackerly" mindset. By using this code to unencrypt a picture in the context of an engaging story, you acquire some thinking and doing skills. I believe that Thinkies are

⁸Glorianna Davenport and Michael Murtaugh, "Context: Towards the Evolving Documentary", Proceedings of Multimedia'95, ACM, 1995. [<http://ic.www.media.mit.edu/icPublications/acm-context/acm-context.html>]

applicable to many educational contexts. The simulation running here will provide the attendees with a taste of the complete experience.

Moving beyond these demonstrations, I would like to consider the future of interactivity. The astonishing growth of activity on the WWW suggests that a new medium, which has been trying to birth itself for the past 15-20 years, is finally at hand. What narrative directions will be most appropriate to future media? What modes of interactivity will compel us to participate? How will the society of audience be shaped as the technology moves forward?

This year we began a new project which we call the World Wide Movie Map. In this project, we hope to engage an international community to participate with us in creating a repository for personal explorations of place. If the project succeeds we should be able to hand shake from place to place around the world by the millennium. Ultimately, we hope that you will all participate with us to create portraits of Sydney, Melbourne, Cairns and so forth. This idea of individuals mapping the terrain of the world grew out of an old exploration in simulated travel, the Aspen Movie Map, which was constructed at the MIT Architecture Machine Group in 1979-81. The World Wide Movie Map will emphasize the role of the amateur in building this sort of simulation; as a research project, our interests are split between understanding community storytelling in an electronic environment and in understanding the tools which are needed to increase access to the WWW as a distributed publishing and distribution enterprise.

In this project, we will provide a navigational substructure, perhaps built on the Argus Map database, and some tools for submitting material. Our tools will help you to publish content and attach keyword annotations to your submissions. Image processing tools might include stabilizing shaky video shots, translating video into still imagery using "salient still"⁹ technology developed at the Lab and, perhaps, searching for similar pictures. This project introduces a new level of collaborative and playful construction and should eventually incorporate all of the modes of interactivity which Jean Piaget observed in children's game-play.

Another graduate student, Kevin Brooks, is approaching dynamic movie playout by designing "Agent Stories." The environment contains agents which understand how to use story parts, such as character introduction, story introduction, diversion etc. This work has allowed us to think about a new kind of soap opera. Instead of having a fixed set of characters whose lives move forward *en masse*, we might choose to build a distributed character set whose paths cross infrequently but in profound ways. In a daytime show based on this premise, a new character could be introduced each week. Characters could be developed more or less "thickly" and could potentially be written by different writers. The audience's expectation could be woven into a serendipitous *déjà vu* which would occur as the audience happened to be watching the other side of one of those fleeting interactions.

Working in video, I find myself constantly battling the bandwidth issue. In order for dynamic, interactive narrative to become a widespread future form, we need to

⁹Laura Teodosio and Walter Bender, Salient Stills from Video, Proceedings of the ACM Multimedia '93 Conference, 1993.

allay fears about limitations in the bandwidth of the delivery system. For this reason, we are collaborating with Media Lab researchers who are working on the problem of model-based or structured video. Using this approach, the temporal stream is constructed by assembling lower-bandwidth objects on-the-fly into a 3D representation. This method offers the potential for extreme compression as it does away with a frame-based representation. Rather than shipping the whole frame once every 30 seconds, only changes in the frame data need to be delivered to the presentation device.

All of these projects incorporate strategies for using a two-way channel to personalize and globalize communication. They provide examples of how the language of interactivity may affect story structure and, by emphasizing story, celebrate the diversity of human endeavor. Today, I have the sense that our culture has finally become deeply democratic; this is reflected more in the changing nature of our communications systems than it ever was in the framework of liberal politics. Perhaps we can trace the "home page" to Andy Warhol, who suggested that everyone in the world had the right to 15 minutes of media fame. Of course, if I can make something and a million -- or a billion -- people want to see it, that should also be OK, too.

I think Piaget's taxonomy provides a valuable benchmark which we can use to describe approaches to the interactive story experience. Both designers and audience can measure their interest in incorporating some or all of these interaction types: practice, pretend, rules of the game, and construction. However, these interactions must be meaningfully mapped into the story environment. Because the medium is very young, we do not always understand what we are making. The wonder of being an "interactive multimedia" author today lies in the discovery, the effective surprise of the creation. The author only has a limited ability to previsualize a project before it becomes a functioning system. I encourage designers to throw convention to the wind, not to dismiss visual intensity of a presentation, but rather to disregard -- in the beginning -- decisions about buttons and mouse clicks and menus and branching, and concentrate on finding an underlying structure which can be driven procedurally and presented with a rich dimensionality. The design, which grows procedurally out of this underlying structure, can be extensible as well as memorable.

With that I end my talk and invite questions.

Question: I just wanted to ask ... I didn't quite get "Dogmatic." I found it entertaining, but I wanted to climb inside of it or something, There was something, I'm not quite sure what you're objective was with it.

Answer: The research objective in that case was to make an interactive story which offered apparent freedom to the viewer, who acts in first person, but which was controlled by a story engine. The story was constructed using principles of narrative guidance. The experience incorporated various image, sound, scene and act dynamics to insure that the story adjusted itself to the behavior of the viewer. There were elements -- the character of the dog, the music, the cuts -- which were rendered dynamically during the experience. As the viewer, you sense that the dog is an autonomous character, that he is responding to you. This piece was created in 1994-

5, so what you're seeing is already in the literature; however, it conveys a moment when we have created a subtlety which goes beyond a hard-wired experience. What is the nature of storytelling when your characters and your editor are dynamic programs?

Question: I was just wondering if you could ... what are your views of Michael's opening comments that interactive media may not necessarily need a story in the same way that a game of pool doesn't.

Answer: Right. A lot of people talk about that. I think we constantly make stories all the time in our minds. I would say even a game of pool is a story. The story is associated with how we work within rules and arrive at a goal. A game of pool is not a story in the traditional Aristotelian sense, which "Dogmatic" is. The story I make up when I am in a pool hall has to do with the people as people and players and their interactions as well as the way in which my knowledge of geometry and calculus create expectations; in short, I have a goal in mind, and there is an obstacle to that goal. My sense is that almost everything we do -- having a conversation, dreaming, playing a sport -- are handled by the human mind as story, even though they are not traditional Aristotelian, conflict-resolution scenarios. Sports does, of course, present a conflict-resolution scenario. And then, there is always the "psychological game" aspects, which the participants enact around the table. Michael and I can perhaps argue that one out later.

Question: With "Dogmatic," I'm interested with the future development, what was happening with the physical environment. So, for example, I don't like mice, I'm really getting fed up with the mouse. So could you have a steering wheel, to direct the dog, like drive around. And, could other people sit there with another steering wheel or engine. So the dog is moving with you with a steering wheel and then someone else is... say, three or four other people interacting, moving the other characters within that space.

Answer: So what you're asking, is it extensible to multiple players? We are really working on the coding scenario and what the story engine has to do. So I think the story engine will be extensible, but let me actually give you another example of an interactive piece that I think is really marvelous, which is done by the same person in the Museum of Science and Industry in Chicago. In that particular scenario people can go down an aisle -- there are processing stations. -- and get their face photographed. They can process their face in various different ways and the faces go into the network. The aisle leads you to a room in which a single audience member can sit and peer into a world using a Fake-Space boom. Two people participate on either side by changing the time of day and the ambient sound of the passage through the virtual world. Occasionally, the people's faces that have been captured and are in the network fly into the virtual reality space and go "blop" onto a building or onto the landscape. I think that there are many ways of extending the sort of idea, the question is how much flexibility do you have to have in the experience itself, how much flexibility do you have for shortening or for lengthening it. That's what we were looking at. It's a fairly complex set of programs that runs that.

Question: I have three questions about a person you were talking about named Piaget. Who is Piaget, what has he written, and how do you spell Piaget.

Answer: The last one's easy -- "P-i-a-g-e-t." His first name is Jean, and I apologize for not establishing the context. Jean Piaget was a psychologist in Switzerland working from the 1920s onward into the 1950s and 1960s. Basically, he studied how children learn and he proposed the first taxonomy of child development. When does a child come to understand the abstract concept of volume, for example? Or that volumes of different shapes represent can represent the same amount of material. He made up a famous set of pouring experiments where you have water in a tall container and in a short squat container and the child talks tells the researcher whether or not the volumes are the same. Piaget also observed children at play and described their game-playing activities. Seymour Papert, who founded the Learning and Epistemology group in our Lab, worked with Piaget in Switzerland in the 1960's.

Question: I want to ask you about something you said fairly quickly, about "we discovered that there was an electrical charge that went through the body, and we could turn pages." You seemed to be talking about that in terms of that physical environment with the whale and I wondered if you'd worked with that in terms of avoiding the mouse in the computer environment and using touch.

Answer: What it provides us with is a way of talking to a sensor. In our theatrical space, for example, we used several different kinds of sensors to detect the Explorer; we used sensors on the floor or sensors on the walls or an IR beam that you crossed. This allows you to actually ... the system would know whether I was approaching the person on this side of that first row or whether I was approaching Cathy over at the other side of that first row. So, if you had two sensors here, they could measure my proximity and the direction of my motion within the limits of the sensing field. It is a small discovery which can have very profound implications for how we will interact with computers in the future.

Question: I'm Ted Clark from the Australian Teachers of Media. I was very interested in the thing that you were calling "Thinkies" and how they would work. You said that they were transportable across to other educational areas. How do "Thinkies" actually work?

Answer: "Thinkies" are stories which situate the audience in a "head space." When you buy into in a "Thinkie," you accept the fact that you will be manipulated by the author in a particular way. You also know that when you walk into a movie... it is an issue of expectation. When you walk into a movie, if you've read the reviews, or talked with friends, you have some prior knowledge about how you are going to be manipulated. In this case, we do it more... you have more conscious activity within that manipulation. In a "Thinkie," we could probably get you to understand what it is like to have some kind of brain damage, for instance. To do this, we have to create an environment where you live in that character's world. This might be very important if you are going to work with people who have different kinds of disabilities. You know, there's a lot of work right now with hyperactive children. I think one of the big problems, and it's sort of talked about a little bit but it's not really talked about, is that most adults don't understand hyperactivity from an

inward perspective. Equally, we could take the rain forest as a subjective discovery and make you try to think like a botanist. How does a botanist go into the rain forest, what is their thinking process? Not just what they think about what they collect, but what is their awareness of the environment. Now, the rain forest is interesting because there is a lot of sensual knowledge that you gain in visiting a rain forest. You are probably very aware of humidity; right now with a computer, we can't totally make you aware of that, but we can give you some ways of measuring that. I think that we're working toward are participatory environments that have much more soul because they have more sensory capacity. The "Thinkie" is built around a problem domain... what problem is the botanist really trying to solve? What are they really looking at, and how could you think a little bit more like they are thinking? So, that's what we're trying to do.

Question: David Jobling, I teach Interactivity at the New South Wales Writer's Centre. Since you mentioned Piaget, I thought it might be worth bringing up transactional analysis dynamic as well as Wolfenbergs normalization techniques in occupational therapy, because I think that there's a lot of things that are derived from Piaget's early work, as well as of a newer dynamics in communication that have been mapped out very clearly in those theories. Do you know them?

Answer: Yes, I agree. That's a little bit ... you can go ahead if you want to say a couple of words about them, but I think certainly Piaget has had an enormous impact; he's a very important thinker to this whole area. One of the things that I think Michael also talked about is getting to know who some of these thinkers are. That allows us to generate new ideas...

Questioner: Indeed, indeed, which is why I thought I should throw a couple of other names into the stew like Wolfenbergs (?) and transactional analysis.

Question: My name's Peter. You mentioned the speed of feedback loops. The Logo project seemed to have a very long feedback loop of some five days. Have you done much research into how this effects interactivity, like, the length of the actual feedback loop.

Answer: Actually, the loop is not five days, the speed of the loop is however fast your asynchronous mail flies across the network, because we might have somebody in Australia playing with somebody in Boston playing with somebody in San Francisco or somebody in South America. Your effect on the system in "Lurker" is, I would say, minimal. Your effect on the main story line is minimal. Your effect on other players in the system is maximal. So, you can actually really help another player solve some of the puzzles that are difficult for them if you're interested in getting into that kind of exchange dynamic with them. It's not perfect. The feedback loop actually occurs every time we ship e-mail to you, and every time you ship an e-mail back in; the potential exists for exchange. I think there are a lot of problems with "Lurker," by the way, I don't want to make light of that, but they're probably too difficult to get into here and maybe that would be a good offline conversation.

Michael Hill: So we'll continue with offline, then, and thank Glorianna Davenport.