of a historic convergence as novelists, playwrights, and filmmakers move toward multiform stories and digital formats; computer scientists move toward the creation of fictional worlds; and the audience moves toward the virtual stage. How can we tell what is coming next? Judging from the current landscape, we can expect a continued loosening of the traditional boundaries between games and stories, between films and rides, between broadcast media (like television and radio) and archival media (like books or videotape), between narrative forms (like books) and dramatic forms (like theater or film), and even between the audience and the author. To understand the new genres and the narrative pleasures that will arise from this heady mixture, we must look beyond the formats imposed upon the computer by the older media it is so rapidly assimilating and identify those properties native to the machine itself.

Murray, Janet (1998): Hamlet on the Holodeck. Cambridge, MA: The MIT Press.

Chapter 3

From Additive to Expressive Form

Beyond "Multimedia"

The birth of cinema has long been assigned to a single night: December 28, 1895. A group of Parisians, so the legend goes, were gathered in a darkened basement room of the Grand Café on the Boulevard des Capucines when suddenly the lifelike image of a mighty locomotive began moving inexorably, astonishingly toward them. There was a moment of paralyzed horror, and then the audience ran screaming from the room, as if in fear of being crushed by an actual train. This no doubt exaggerated account is based on an actual event, the first public showing of a group of short films that included "Arrival of a Train at La Ciotat Station" by the Lumière brothers, who (like Edison in America) had just invented a reliable form of motion picture photography and projection. Film scholars have recently questioned whether the novelty-seeking crowd really panicked at all.[1] Perhaps it was only later storytellers who imagined that the first projected film image, the novelty attraction of 1895, could have carried with it the tremendous emotional force of the many thrilling films that followed after it. The legend of the Paris café is satisfying to us now because it falsely conflates the arrival of the representational technology with the arrival of the artistic medium, as if the manufacture of the camera alone gave us the movies.

As in the case of the printing press, the invention of the camera led to a period of incunabula, of "cradle films." In the first three decades of the twentieth century, filmmakers collectively invented the medium by inventing all the major elements of filmic storytelling, including the close-up, the chase scene, and the standard feature length. The key to this development was seizing on the unique physical properties of film: the way the camera could be moved; the way the lens could open, close, and change focus; the way the celluloid processed light; the way the strips of film could be cut up and reassembled. By aggressively exploring and exploiting these physical properties, filmmakers changed a mere recording technology into an expressive medium.

Narrative films were originally called photoplays and were at first thought of as a merely additive art form (photography plus theater) created by pointing a static camera at a stagelike set. Photoplays gave way to movies when filmmakers learned, for example, to create suspense by cutting between two separate actions (the child in the burning building and the firemen coming to the rescue); to create character and mood by visual means (the menacing villain backlit and seen from a low angle); to use a "montage" of discontinuous shots to establish a larger action (the impending massacre visible in a line of marching soldiers, an old man's frightened face, a baby carriage tottering on the brink of a stone stairway). After thirty years of energetic invention, films captured the world with such persuasive power and told such coherent and compelling stories that some critics passionately opposed the addition of sound and color as superfluous distractions.

Now, one hundred years after the arrival of the motion picture camera, we have the arrival of the modern computer, capable of hooking up to a global internet, of processing text, images, sound, and moving pictures, and of controlling a laptop display or a hundred-foot screen. Can we imagine the future of electronic narrative any more easily than Gutenberg's contemporaries could have

imagined *War and Peace* or than the Parisian novelty seekers of 1895 could have imagined *High Noon?*

One of the lessons we can learn from the history of film is that additive formulations like "photo-play" or the contemporary catchall "multimedia" are a sign that the medium is in an early stage of development and is still depending on formats derived from earlier technologies instead of exploiting its own expressive power. Today the derivative mind-set is apparent in the conception of cyberspace as a place to view "pages" of print or "clips" of moving video and of CD-ROMs as offering "extended books." The equivalent of the filmed play of the early 1900s is the multimedia scrapbook (on CD-ROM or as a "site" on the World Wide Web), which takes advantage of the novelty of computer delivery without utilizing its intrinsic properties.

For example, one early version of a Web soap about a group of friends living in New York offers diary pages of text spiced with sexually suggestive photos. The wordiness of the journal keeps us constantly scrolling through the screens, impatient for something to happen in the narrated story or for something to do, like clicking on a link to get something new. There are, in fact, clickable buttons in the journal, but instead of offering new information they merely allow us to hear (after time delays for downloading the sound clip and for installing the necessary software to play the sound file if we do not already have it) actors speaking exactly the same dialog printed on the screen. The audio snippets are amusing novelties at best, and at worst they work like so many small apologies for the limits of the printed text. Just as the photographed plays of early filmmakers were less interesting than live theater, this early Web soap continually reminds us of how much less vivid it is than the romance novels and television dramas it draws upon.

A more digitally sophisticated Web soap would exploit the archiving functions of the computer by salting each day's new episode with allusions (in the form of hot word links) to exciting previous installments. Our clicking would then be motivated not by curiosity about the media objects (show me a video clip) but by curiosity about the plot (why does she say that about him?). The computer presentation would thereby allow pleasures that are unattainable in broadcast soaps. For example, we could follow a single appealing subplot while ignoring the companion plots that may drive us crazy, or we could come in at any time in the story and review important past events in all their dramatic richness. Instead of using audio redundantly to act out dialogue in a diary entry, a sophisticated Web soap might provide the audio as an integral part of the plotline—perhaps as the wiretap of a murder threat or a political negotiation or as a phone message that carries information of hidden romantic liaisons.

Some Web stories are already using such techniques, and no doubt all of them will in time. Their adoption is part of the inevitable process of moving away from the formats of older media and toward new conventions in order to satisfy the desires aroused by the digital environment. We are now engaged in thousands of such discoveries in all the subgenres of electronic narrative, the result of which will be the development of narrative pleasures intrinsic to cyberspace itself. Therefore, if we want to see beyond the current horizon of scrapbook multimedia, it is important first to identify the essential properties of digital environments, that is, the qualities comparable to the variability of the lens, the movability of the camera, and the editability of film, that will determine the distinctive power and form of a mature electronic narrative art.

ELIZA Comes to Life

Well before the advent of multimedia, there was a moment in the history of the computer that demonstrated its representational and narrative power with the same startling immediacy as the Lumières' train did for the motion picture camera. This is the famous but often misunderstood moment in which the first completely computer-based character was created.

In 1966, Joseph Weizenbaum, a professor of computer science at MIT, created, as an experiment in natural language processing, a computer program called ELIZA that carried on a conversation by replying to typed-in statements with printed words. Since this was before the widespread use of computer screens, the program used a form of teletype device connected to one of the first time-sharing computer networks. The resulting persona, Eliza, was that of a Rogerian therapist, the kind of clinician who echoes back the concerns of the patient without interpretation. She also displayed a narrowly Freudian interest in topics like sex and the family. In a research paper describing the program, Weizenbaum included this remarkably coherent sample conversation:

> *Patient:* Well, my boyfriend made me come here.
> *Eliza:* YOUR BOYFRIEND MADE YOU COME HERE.
> *Patient:* He says I'm depressed much of the time.
> *Eliza:* I AM SORRY TO HEAR YOU ARE DEPRESSED.
> *Patient:* It's true. I am unhappy.
> *Eliza:* DO YOU THINK COMING HERE WILL HELP YOU NOT TO BE UNHAPPY?
> *Patient:* I need help, that much seems certain.
> *Eliza:* WHAT WOULD IT MEAN TO YOU IF YOU GOT SOME HELP?
> *Patient:* Perhaps I could learn to get along with my mother.
> *Eliza:* TELL ME MORE ABOUT YOUR FAMILY.[2]

Eliza proved astonishingly persuasive. Soon a creation legend started to circulate that was based on an incident involving a researcher at a Cambridge think tank who left an early version of ELIZA running on his computer, which a vice president of the company thought was typing out a communication being transmitted by his colleague from home. (Just as the Lumière train story is a staple of film history, the Eliza story in its several versions appears in many textbooks and computer science lectures, often told as if it happened to Weizenbaum himself.) The conversation takes place at nine on a Saturday morning:

Vice president: Do you think that I can use the computer this morning?

Teletype: WHY DO YOU ASK?

Vice president: I might be able to make some additional telcomp (computer system) sales.

Teletype: WHY AREN'T YOU SURE?

Vice president: My prospective customers have never seen a demonstration of the system.

Teletype: DOES THIS MEAN VERY MUCH TO YOU?

Vice president: Of course it does.

Teletype: YOU SEEM VERY POSITIVE ABOUT THAT.

Vice president: This is ridiculous.

Teletype: WHAT DO YOU MEAN THIS IS RIDICULOUS?

Vice president: Is there any reason why I can't use the computer now?

Teletype: TELL ME MORE ABOUT YOUR PROSPECTIVE CUSTOMERS.

Vice president: Please dial me up on 491-1850

But because the vice president forgets to type in a final period, the machine does not answer him. Infuriated, he calls his colleague on the phone, resulting in this predictable exchange:

Vice president: Why are you being so snotty to me?

Researcher: What do you mean why am I being so snotty to you? (*Explosion of anger*)[3]

The story has become a legend because it discharges the anxiety aroused by the fear that Weizenbaum had gone too far, that he had created a being so much like an actual person that we would no longer be able to tell when we were talking to a computer and when to a human being. This is very much like the fear that people would mistake film images for the real world.

Eliza was not just persuasive as a live conversationalist; she was also remarkably successful in sustaining her role as a therapist. To Weizenbaum's dismay, a wide range of people, including his own secretary, would "demand to be permitted to converse with the system in private, and would, after conversing with it for a time, insist, in spite of [Weizenbaum's] explanations, that the machine really understood them."[4] Even sophisticated users "who knew very well that they were conversing with a machine soon forgot that fact, just as theatergoers, in the grip of suspended disbelief, soon forget that the action they are witnessing is not 'real' " (p. 189). Weizenbaum had set out to make a clever computer program and had unwittingly created a believable character. He was so disconcerted by his achievement that he wrote a book warning of the dangers of attributing human thought to machines.

Without any aid from graphics or video, Eliza's simple textual utterances were experienced as coming from a being who was present at that moment. What was the representational force that allowed the computer to bring her so compellingly to life?

The Four Essential Properties of Digital Environments

When we stop thinking of the computer as a multimedia telephone link, we can identify its four principal properties, which separately and collectively make it a powerful vehicle for literary creation. Digital environments are procedural, participatory, spatial, and encyclopedic. The first two properties make up most of what we mean by the vaguely used word *interactive*; the remaining two properties help to make digital creations seem as explorable and extensive as the actual world, making up much of what we mean when we say that cyberspace is *immersive*.

Digital Environments Are Procedural

Eliza was brought to life by the procedural power of the computer, by its defining ability to execute a series of rules. It is surprising how often we forget that the new digital medium is intrinsically procedural. Although we may talk of an information highway and of billboards in cyberspace, in fact the computer is not fundamentally a wire or a pathway but an *engine*. It was designed not to carry static information but to embody complex, contingent behaviors. To be a computer scientist is to think in terms of algorithms and heuristics, that is, to be constantly identifying the exact or general rules of behavior that describe any process, from running a payroll to flying an airplane.

Weizenbaum stands as the earliest, and still perhaps the premier, literary artist in the computer medium because he so successfully applied procedural thinking to the behavior of a psychotherapist in a clinical interview. It is the cleverness of Weizenbaum's rules that creates the illusion that Eliza understands what is said to her and that induces the user to continue the conversation. For example, if the user says, "Everybody laughs at me," the program can apply the rule that deals with *me* statements to echo the remark as, "You say that everybody laughs at you." This general rule models the neutrality of the Rogerian therapist, who reflects the patient's statements without judgment. Or, more cleverly, the program can select the specific rule for the word *everybody*, and reply, "Who in particular are you thinking of?"

Weizenbaum made sure the program would select the more specific reply by writing a rule that gives *everybody* a higher priority as a "key word" than *me*. Eliza's language processing includes no representation of syntax or meaning. Most of the time Eliza merely listens for key words such as *everybody* or *depressed* or *father*. Yet she is more believable than many programs with a more detailed representation of linguistic structure. Her credibility is the result of Weizenbaum's adroitness in formulating rules of discourse that are based on the ways in which a therapist would behave.

Eliza is not a neutral procedural model but a comic interpretation. Just as an actor might use makeup to emphasize a large nose or an owlish look, Weizenbaum used procedural notation to give parodic emphasis to some familiar absurdities of therapy itself. Like any classic type, Eliza is predictable in a way we instantly recognize and savor.

She is a descendent of the vaudeville doctor who bungles the exam, of the comic simpleton whose role is to misunderstand whatever is happening around her, and of the enjoyably repetitive flat characters in Dickens who repeat the same tag lines or behavior whenever we see them ("I never will desert Mr. Micawber," chimes the wife of the pompous screwup in *David Copperfield*).[5] The very rigidities of the computer are exploited for comic effect because they echo the rigidities of certain psychotherapeutic methodologies. The joke is on the therapists.

By mirroring the artificial constraints of therapy talk, Eliza makes them ridiculous. Consider, for instance, the following exchange:

> *Eliza:* Tell me more about your family.
> *Patient:* My mother is very motherly.
> *Eliza:* Who else in your family is very motherly?
> *Patient:* I am very motherly.
> *Eliza:* Is it because you are very motherly that you came to see me?
> *Patient:* Do you need a mother?
> *Eliza:* We are discussing you, not me.[6]

Eliza will reflect back any "I am X" statement with a formulation like "Is it because you are X that you came to see me?" Here, the patient responds to the suggestion by playfully challenging the shrink. But, Eliza invokes her rule for *you* questions, redirecting the conversation back to the patient like any well-trained therapist. Putting these two rules together creates a coherent dramatic dialogue with a punch line. Such dialogues are easy to generate and irresistible to attempt. They amount to a kind of collaboratively written comedy skit. Few people would now perceive Eliza as a real psychotherapist.[7] But as an improv partner, Eliza is still quite popular, available on most computer networks and sometimes sending a bill by e-mail after the session.

The lesson of ELIZA is that the computer can be a compelling medium for storytelling if we can write rules for it that are recognizable as an interpretation of the world. The challenge for the future is how to make such rule writing as available to writers as musical notation is to composers.

Digital Environments Are Participatory

The energy with which people enter into dialog with Eliza is also evidence of a second core property of the computer: its participatory organization. Procedural environments are appealing to us not just because they exhibit rule-generated behavior but because we can induce the behavior. They are responsive to our input. Just as the primary representational property of the movie camera and projector is the photographic rendering of action over time, the primary representational property of the computer is the codified rendering of responsive behaviors. This is what is most often meant when we say that computers are *interactive*. We mean they create an environment that is both procedural and participatory.

Eliza's responsiveness is limited by her poor understanding of language, which makes her liable to nonsense utterances. Her direct successors are therefore mostly in research environments. It fell to another group of MIT computer scientists to develop a fictional universe that structures participation more tightly, resulting in a more sustained engagement.

A few years after the invention of ELIZA, researchers at the MIT Laboratory for Computer Science brought forth a widely popular computer-based story, the adventure game *Zork*, which is based on the *Dungeons and Dragons* tabletop game.[8] In *Zork* the computer plays the role of dungeon master by providing an invisible landscape that serves as the game board and by reporting to players on the effects of their actions. Within *Zork*'s fantasy world, players move through dungeon rooms by typing in navigational commands (north, south, east, west, up, down), look for objects that can be manipulated (by typing appropriate commands, such as "read book," "take sword," "drink potion"), solve riddles, and fight off evil trolls. The game (which, like ELIZA, is still a popular feature of university networks) begins like this:

Welcome to Zork.
West of House.
You are in an open field west of a big white house with a boarded front door. There is a small mailbox here.
Interactor> Go north.
North of House.
You are facing the north side of a white house. There is no door here, and all the windows are barred.
Interactor> East.
Behind House.
You are behind the white house. In one corner of the house there is a small window which is slightly ajar.
Interactor> Open the window.
With great effort, you open the window far enough to allow entry.
Interactor> Go in.
Kitchen.
You are in the kitchen of the white house. A table seems to have been used recently for the preparation of food. A passage leads to the west, and a dark staircase can be seen leading upward. To the east is a small window which is open. On the table is an elongated brown sack, smelling of hot peppers.
A bottle is sitting on the table.
The glass bottle contains:
A quantity of water.

In making a fantasy world that responded to typed commands, the programmers were in part celebrating their pleasure in the increasingly responsive computing environments at their disposal. Before the 1970s most complex programming was done by writing a set of commands on a piece of paper; transferring them to keypunch cards; and taking the stack of cards to a mainframe computer (in an uncomfortably chilly room dedicated to keeping the machines from overheating), from which, much later, a cumbersome paper printout would emerge. Only one person could use a machine at a time. Whenever a program crashed (which was often), the output consisted of a "core dump"—a long series of *0*'s and *1*'s arranged in eight-digit units, showing what each bit and byte in the computer memory looked like at the moment the computer quit. Debugging a program in this environment was time-consuming and tedious.

In the mid-1960s research labs began developing the current computing environment of a display device and a keyboard (originally a telex machine) linked up to a time-sharing network that let programmers send input directly to a running program and receive a response. They were also making wide use of programming languages that were interpreted rather than compiled. All programming code written in higher-level languages (with commands like "If a = 1, then print file") must be translated into machine language instructions (with commands that look a lot like the raw *0*'s and *1*'s of the bits themselves) by either a compiler or interpreter program. Compiling your code before running it is like writing a book and then hiring someone to translate it for your readers. Using an interpreter is the equivalent of giving a speech with simultaneous translation. It provides more direct feedback from the machine and a more rapid cycle of trial and revision and retrial. The particular programming language in which both ELIZA and *Zork* were written, LISP (LIst Processing Language), was developed at MIT in the 1950s for artificial intelligence. Running LISP on a time-sharing system meant that its dynamic "interpreter" could immediately "return" an "evaluation" of any coded statement you typed into it, much as a calculator immediately returns the sum of two numbers. The result was a more conversational structure between the programmer and the program, a dialogue in which the programmer could test out one function at a time and immediately receive the bafflingly inappropriate or thrillingly correct responses. Both ELIZA and *Zork* reflected this newly animated partnership.

Whereas ELIZA captured the conversational nature of the programmer–machine relationship, *Zork* transmuted the intellectual challenge and frustrations of programming into a mock-heroic quest filled with enemy trolls, maddening dead ends, vexing riddles, and rewards for strenuous problem solving. ELIZA was focused on the cleverness of the machine-created world; *Zork* was focused on the experience of the participant, the adventurer through such a clever rule system. Zork was set up to provide the player with opportunities for making decisions and to dramatically enact the results of those decisions. If you do not take the lamp, you will not see what is in the cellar, and then you will definitely be eaten by the grue. But the lamp is not enough. If you do not take water with you, you will die of thirst. But if you drink the wrong water, you will be poisoned. If you do not take weapons, you will have nothing with which to fight the trolls. But if you take too many objects, you will not be able to carry the treasure when you find it. In order to succeed, you must orchestrate your actions carefully and learn from repeated trial and error. In the early versions there was no way to save a game in midplay, and therefore a mistake meant repeating the entire correct procedure from the beginning. In a way, the computer was programming the player.

Part of the pleasure of the participant in *Zork* is in testing the limits of what the program will respond to, and the creators prided themselves on anticipating even wildly inappropriate actions. For instance, if you type in "eat buoy" when a buoy floats by on your trip up a frozen river in the magic boat, then the game will announce that it has taken it instead and will add, "I don't think that the red buoy would agree with you." If you type in "kill troll with newspaper," it will reply, "Attacking a troll with a newspaper is foolhardy." The programmers generated such clever responses not by thinking of every possible action individually but by thinking in terms of general categories, such as weapons and foods. They made the programming function associated with the command word *eat* or *kill* check the player's typed command for an appropriate object; a category violation triggers one of these sarcastic templates, with the name of the inappropriate object filled in.

Because LISP programmers were among the first to practice what is now called object-oriented software design, they were well prepared to create a magical place like the world of *Zork*. That is, it came naturally to them to create virtual objects such as swords or bottles because they were using a programming language that made it particularly easy to define new objects and categories of objects, each with its own associated properties and procedures. The programmers also exploited a programming construct known as a "demon" to make some things happen automatically without the player's explicit action; for instance, in *Zork* a magic sword begins to glow if there is danger nearby, a stealthy thief comes and goes at his own will, and a fighter troll attacks the adventurer at unpredictable times. The programmers were also prepared by research on automatons to keep track of the state of the game, which allowed them to guess at the context of commands that would otherwise be ambiguous. For instance, if a player types "attack," the program looks around for a nearby villain and a weapon; if there are two weapons, it asks which one the player wants to use. These techniques, which were taken from simulation design and artificial intelligence work, allowed the *Zork* programming team to create a dynamic fictional universe.

By contrast, more conventional programmers of the 1970s were still thinking in terms of the branching trees, fixed subroutines, and uniform data structures that go back to the early understanding of the computer as a means of encoding information purely in the form of yes/no decisions. In fact, most interactive narrative written today still follows a simple branching structure, which limits the interactor's choices to a selection of alternatives from a fixed menu of some kind. The *Zork* dungeon rooms form a branching structure, but the magical objects within the dungeon each behave according to their own set of rules. And the interactor is given a repertoire of possible behaviors that encourage a feeling of inventive collaboration. The *Zork* programmers found a procedural technology for creating enchantment.

The company they formed, Infocom, is, though long out of business, still revered by players. Many fans attribute the imaginative superiority of Infocom games to the predominance of text over graphics, just as nostalgic radio fans prefer the sightless "theater of the imagination" to television. But though the writing in its games was skillful, it was not the true secret of Infocom's success. What made the games distinctive was the sophisticated computational thinking the programmers brought to shaping the range of possible interactions.

The lesson of *Zork* is that the first step in making an enticing narrative world is to script the interactor. The *Dungeons and Dragons* adventure format provided an appropriate repertoire of actions that players could be expected to know before they entered the program. The fantasy environment provided the interactor with a familiar role and made it possible for the programmers to anticipate the interactor's behaviors. By using these literary and gaming conventions to constrain the players' behaviors to a dramatically appropriate but limited set of commands, the designers could focus their inventive powers on making the virtual world as responsive as possible to every possible combination of these commands. But if the key to compelling storytelling in a participatory medium lies in scripting the interactor, the challenge for the future is to invent scripts that are formulaic enough to be easily grasped and responded to but flexible enough to capture a wider range of human behavior than treasure hunting and troll slaughter.

Digital Environments Are Spatial

The new digital environments are characterized by their power to represent navigable space. Linear media such as books and films can portray space, either by verbal description or image, but only digital environments can present space that we can move through. Again, we can look to the 1970s as the period that made this spatial property apparent. At Xerox PARC (Palo Alto Research Center) a group of visionaries created the first graphical user interface, the image of a desktop filled with file folders, which is currently the worldwide standard for computer file management. At Atari, inventors created the first graphics-based games, first *Pong* and then *PacMan*, which established the computer as a spatial game board and paved the way for today's videogames and CD-ROMs. At MIT the Architecture Machine Group, led by Nicholas Negroponte, created the first "surrogate travel system" by making a "movie map" of Aspen, Colorado, a photographic record of the town that could be "walked through" on the computer (each press of an arrow key or mouse click indicated a step forward) in simulation of an actual visit. All of these were important steps toward the current digital environment, as was the concurrent growth of the ARPANET (Advanced Research Projects Agency NETwork), the precursor of the Internet funded by the U.S. Department of Defense, which was mapping out a virtual space that was increasingly real to those who used it. We recognize the fruit of all of these developments in our conceptualization of the digital domain as "cyberspace," an environment with its own geography in which we experience a change of documents on our screen as a visit to a distant site on a worldwide web.

Although this spatial property has been widely exploited in graphical applications, it is in fact independent of the computer's ability to display maps, pictures, or even three-dimensional models. It is also independent of its communicative function in linking geographically distant places. The computer's spatial quality is created by the interactive process of navigation. We know that we are in a particular location because when we enter a keyboard or mouse command the (text or graphic) screen display changes appropriately. We can verify the relation of one virtual space to another by retracing our steps. The text-based dungeons of *Zork*, the sequenced stills of the enchanter's isle of Myst, the flat worlds of the multilevel maze games, the crow's waterfall visible in the Placeholder VR helmet, the continuous three-dimensional world of the new videogame dreamscapes—all are realized for the interactor by the process of navigation, which is unique to the digital environment.

Let us return to *Zork* to see how this navigational creation of space lends itself to dramatic engagement:

Interactor> West.
Living Room.
You are in the living room. There is a door to the east, a wooden door with strange gothic lettering to the west, which appears to be nailed shut, and a large oriental rug in the center of the room . . .
Interactor> Look under rug.
Underneath the rug is a closed trapdoor.
Interactor> Move the rug.
With a great effort, the rug is moved to one side of the room. With the rug moved, the dusty cover of a closed trapdoor appears.
Interactor> Open trapdoor.
The door reluctantly opens to reveal a rickety staircase descending into darkness.
Interactor> Go down.
It is pitch-black. You are likely to be eaten by a grue. Your sword is glowing with a faint blue glow.
Interactor> Turn on lamp.
The lamp is now on.
Cellar.
You are in a dark and damp cellar with a narrow passageway leading east, and a crawlway to the south. On the west is the bottom of a steep metal ramp which is unclimbable. The trapdoor crashes shut, and you hear someone barring it.

You, as player/interactor, have walked into a dungeon that someone is sealing shut behind you! The moment is startling and immediate, like the firing of a prop gun on the stage of a theater. You are not just reading about an event that occurred in the past; the event is happening *now*, and, unlike the action on the stage of a theater, it is happening to *you*. Once that trapdoor slams, the only navigational commands that work are the ones that lead further and deeper into the troll-filled lower world. The dungeon itself has an objective reality that is much more concrete than, for instance, the jail on the Monopoly board or a dungeon in a tabletop game of *Dungeons and Dragons*—or even a dungeon in a live-action role-playing game—because the words on the screen are as transparent as a book. That is, the player is not looking at a game board and game pieces or at a *Dungeons and Dragons* game master who is also in his or her algebra class or at a college classroom or campsite in the real world. The computer screen is displaying a story that is also a place. The slamming of a dungeon door behind you (whether the dungeon is described by words or images) is a moment of experiential drama that is only possible in a digital environment.

The dramatic power of navigation is also apparent outside the realm of the adventure game. For instance, Stephanie Tai, a student in my course on writing interactive fiction wrote a first-person poetic monologue about a sleepless night. Each screenful of text is a stanza and ends with a sentence fragment that connects syntactically with two or more stanzas, which are reached by clicking on arrows placed at the midpoint of the top, bottom, left, and right margins of the screen. Mouse-clicking through the mind of the insomniac is like a walk through a labyrinth. There are multiple end points to the maze, including one with just the single word *asleep* and another with the words *alone in this misery* in white letters on a black background. The poem is satisfying because the action of moving by arrows around a maze mimics the physical tossing and turning and the repetitive, dead-end thinking of a person unable to fall asleep. The movement through the cards makes a coherent pattern, but it is not one that could be modeled in physical space because the movement between links is not necessarily reversible. The navigational space of the computer allows us to express a sequence of thoughts as a kind of dance.

Stuart Moulthrop's ambitious hypertext novel *Victory Garden* (1992), whose title intentionally echoes the Borges story, is also in the shape of a labyrinth. Similar to a thick Victorian novel, it follows many characters with intersecting lives during the Gulf War. At the

very center of Moulthrop's web is the death of Emily Runebird, an army reserve soldier who is killed in her barracks by an enemy missile. The attack itself is represented by a striking image of shattered text, as if the enemy shell itself had landed on the previous block of writing. We reach this image by following a continuous story thread, mouse-clicking through the screens automatically as if turning the pages of a book. The shattered screen stops us dead in our tracks. The effect of moving from the intact lexia to the shattered one is like an animation of the landing of the shell. The instant of time it takes to go from one screen to the other takes on a poignancy that reflects the abruptness of the soldier's death.

These very dramatic moments mark the beginning of a process of artistic discovery. The interactor's navigation of virtual space has been shaped into a dramatic enactment of the plot. We are immobilized in the dungeon, we spiral around with the insomniac, we collide into a lexia that shatters like a bomb site.These are the opening steps in an unfolding digital dance. The challenge for the future is to invent an increasingly graceful choreography of navigation to lure the interactor through ever more expressive narrative landscapes.

Digital Environments Are Encyclopedic

The fourth characteristic of digital environments, which holds promise for the creation of narrative, is more a difference of degree than of kind. Computers are the most capacious medium ever invented, promising infinite resources. Because of the efficiency of representing words and numbers in digital form, we can store and retrieve quantities of information far beyond what was possible before. We have extended human memory with digital media from a basic unit of portable dissemination of 100,000 words (an average book, which takes up about a megabyte of space in its fully formatted version) first to 65,000,000 words (a 650-megabyte CD-ROM, the equivalent of 650 books) and now to 530,000,000 words (a 5.3 gigabyte digital videodisc, equivalent to 5,300 books), and on upward.

Once we move to the global databases of the Internet, made accessible through a worldwide web of linked computers, the resources increase exponentially.

Just as important as this huge capacity of electronic media is the encyclopedic expectation they induce. Since every form of representation is migrating to electronic form and all the world's computers are potentially accessible to one another, we can now conceive of a single comprehensive global library of paintings, films, books, newspapers, television programs, and databases, a library that would be accessible from any point on the globe. It is as if the modern version of the great library of Alexandria, which contained all the knowledge of the ancient world, is about to rematerialize in the infinite expanses of cyberspace. Of course, the reality is much more chaotic and fragmented: networked information is often incomplete or misleading, search routines are often unbearably cumbersome and frustrating, and the information we desire often seems to be tantalizingly out of reach. But when we turn on our computer and start up our Web browser, all the world's resources seem to be accessible, retrievable, immediate. It is a realm in which we easily imagine ourselves to be omniscient.

The encyclopedic capacity of the computer and the encyclopedic expectation it arouses make it a compelling medium for narrative art. The capacity to represent enormous quantities of information in digital form translates into an artist's potential to offer a wealth of detail, to represent the world with both scope and particularity. Like the daylong recitations of the bardic tradition or the three-volume Victorian novel, the limitless expanse of gigabytes presents itself to the storyteller as a vast tabula rasa crying out to be filled with all the matter of life. It offers writers the opportunity to tell stories from multiple vantage points and to offer intersecting stories that form a dense and wide-spreading web.

One early indication of the suitability of epic-scale narrative to digital environments is the active electronic fan culture surrounding popular television drama series. As an adjunct to the serial broadcasting of these series, the Internet functions as a giant bulletin board on which long-term story arcs can be plotted and episodes from different seasons juxtaposed and compared. For instance, the Web site for the intricately plotted space drama *Babylon Five* contains images of the cast and plot summaries that document the many interwoven stories portrayed over multiple seasons, allowing a newcomer to understand the large cast of characters and the richly imagined array of alien races, each with its own culture and dramatic history. But it is not only science fiction programs that attract this interest. Even viewers of the mainstream television sitcom *Wings* use Web sites and Internet newsgroups to trace plot developments that extend over several years—like Joe and Helen's on-again, off-again courtship—and that may be confusingly jumbled in syndication; they also share digitized clips of favorite moments, such as the couple's comic wedding vows. The presence of such groups is influencing these shows, holding them to greater consistency over longer periods of time. In the past this kind of attention was limited to series with cult followings like *Star Trek* or *The X Files*. But as the Internet becomes a standard adjunct of broadcast television, all program writers and producers will be aware of a more sophisticated audience, one that can keep track of the story in greater detail and over longer periods of time. Since the early 1980s, when Steven Bochco introduced multiple story arcs with *Hill Street Blues*, television series have become more complex, involving larger casts and stories that take anywhere from one episode to several years to conclude. Some stories even remain open-ended after the series is over (especially if the writers were not expecting cancellation). In some ways, television dramas seem to be outgrowing broadcast delivery altogether. To join *Babylon Five* in its second or third season or *Murder One* in midseason is to immediately want to flip back or rewind to earlier episodes. The Internet serves that purpose, making a more capacious home for serial drama than the broadcast environment affords.

Making even fuller use of the computer's properties, by combining its spatial, participatory, and procedural elements with its encyclopedic coverage, are the many on-line role-playing environments in the adventure games tradition. By the 1980s, Zork-like games had grown to accommodate simultaneous multiple players, turning them into Multi-User Dungeons or MUDs, which combine the social pleasures of interplayer communication with the standard command-driven adventures. In the MUDs of the 1990s players are no longer limited to navigating a preexisting dungeon but can use a simple programming language to build their own dungeon or adventure maze and link it up with those of other players by creating objects out of common building blocks. The MUD itself is a collective creation—at once a game, a society, and a work of fiction—that is often based on a particular encyclopedic fantasy domain, such as Tolkien's Middle Earth or *Star Trek*'s twenty-fourth century. For instance, *TrekMuse*, founded in 1990 with over two thousand players, had five hundred people enrolled in its virtual Star Fleet Academy in 1995, each of whom had made up his or her own character, based on the existing *Star Trek* races. The digital narrative environment extends the fictional universe of the television shows and films in a way that is consistent with the canonical version of the story but personalizes it for each of the players.

Some hypertext stories successfully use the encyclopedic extent of the computer to develop multithreaded stories composed of many intersecting plots. In *Victory Garden*, for instance, we can follow a radical professor and his colleagues and graduate students through the same time period as they intersect with one another in the classrooms, offices, and coffee bars, or we can follow them home to witness their tangled domestic lives; we can listen to the official coverage of the Gulf War (with CNN transcripts) or read Emily Runebird's letters. In *The Spot* and similar Web soaps, we can read through the conflicting accounts of the same love affairs and deceptions in the journals of various friends. In on-line murder mysteries like *Crime Story*,[9] we can delve through various document files, including crime scene photos, interview transcripts, and newspaper accounts. We can even leap out of the story altogether and find

ourselves in the "real" world, following a reference to the University of Mississippi right to its own Web site, or finding that the name of a witness seen in the company of the fleeing suspect belongs to a real-life software engineer whose Web page has nothing to do with the fictional crime. Not only does the weblike structure of cyberspace allow for endless expansion possibilities within the fictional world, but in the context of a worldwide web of information these intersecting stories can twine around and through the nonfictional documents of real life and make the borders of the fictional universe seem limitless.

However, the encyclopedic nature of the medium can also be a handicap. It encourages long-windedness and formlessness in storytellers, and it leaves readers/interactors wondering which of the several endpoints is *the* end and how they can know if they have seen everything there is to see. Most of what is delivered in hypertext format over the World Wide Web, both fiction and nonfiction, is merely linear writing with table-of-contents links in it. Even those documents designed explicitly for digital presentation, both fiction and nonfiction, often require too much superfluous clicking to reach a desirable destination or so much scrolling that readers forget where they are. The conventions of segmentation and navigation have not been established well enough for hypertext in general, let alone for narrative. The separation of the printed book into focused chapters was an important precondition of the modern novel; hypertext fiction is still awaiting the development of formal conventions of organization that will allow the reader/interactor to explore an encyclopedic medium without being overwhelmed.

The encyclopedic impulse and the dangers of the encyclopedic expectation are also apparent in simulation games. For instance, *SimCity* (1987) presents the player with a schematic picture of a riverside city site, and places him or her in the role of mayor. The player is free to build the city however he or she would like, by adding to the model on the screen office buildings, factories, homes, a sewer system, electric power plants, a public transportation system, highways, schools, and so on. The software calculates the effects of each change by using models very like the ones used by social scientists and policymakers to study urban systems. Truly bad decisions in *SimCity* can bring critical newspaper articles, social unrest, and even electoral defeat. Well-built cities prosper through multiple decades. Because of the importance of the role in *SimCity*, the mayor is closer in power to God than to any real-life political leader, and the player's sense of omniscient awareness of consequences and omnipotent control of resources is part of the allure of such games.

Well-designed simulations like *SimCity* allow for multiple styles of play. One young programmer friend of mine spent hours building the most prosperous skyscrapered downtown possible. When I asked him about the game, he delighted in showing me the detail in which the city's underground service grid was specified. His wife, who is also a computer professional, took a different approach. Her favorite city was a sprawling environment with tree-lined family neighborhoods whose growing population gratified her tremendously and whose children she could easily imagine happily greeting each newly built playground. When they realized how much their efforts fell along gender lines, they laughed, but they pointed out that there was a more radical difference. For the husband, the program was a satisfyingly complex engineering problem, reinforcing his habitual sense of competence. For the wife, it was a narrative, in which the little parades and cheers of her contented townsfolk were the most memorable dramatic events. And, in fact, later versions of the game have been expanding this narrative quality by allowing the player to live inside a more detailed three-dimensional city rather than only manipulate it from on high.

Both the narrative possibilities and the godlike pleasures of simulation format are further developed in *Sid Meier's Civilization*, a game that puts the player in the role of leader of a civilization over the course of many centuries, while the computer plays the role of adversary civilizations that compete with the player for global resources and technical advancement. Like *SimCity*, *Civilization* allows multiple strategies of play and can accommodate the idealistic seeker of social harmony as well as the warrior player. The narrative interest of the game consists of creating multiple possible versions of an Earth-like history. For instance, it is possible to invent the railroad in B.C. times or to become an undefeated Napoleon. Winning the game is defined as either conquering all the other civilizations (in which case you are rewarded with pictures of the other leaders frowning at you) or sending twenty thousand people into space (in which case you see the spaceport).

Simulations like these take advantage of the authority bestowed by the computer environment to seem more encyclopedically inclusive than they really are. As its critics have pointed out, the political assumptions behind *SimCity* are hidden from the player.[10] This is less true in *Sid Meier's Civilization*, whose title alerts us to the fact that we are receiving a particular person's interpretation of human history rather than a scientific formula. The game also explicitly informs us that the behavior of each of the leaders is the result of three variables: their degree of aggression/friendliness, of expansionism/perfectionism, and of militarism/civilization. Since these are assumptions that players are aware of, they are free to accept or reject them as a reflection of the real world. Nevertheless, the basic competitive premise of the game is not emphasized as an interpretive choice. Why should global domination rather than, say, universal housing and education define the civilization that wins the game? Why not make an end to world hunger the winning condition? Why is the object of the game to compete with other leaders instead of to cooperate for the benefit of all the civilizations without jeopardizing any one country's security?

In an interactive medium the interpretive framework is embedded in the rules by which the system works and in the way in which participation is shaped. But the encyclopedic capacity of the computer can distract us from asking why things work the way they do and why we are being asked to play one role rather than another. As these systems take on more narrative content, the interpretive nature of these structures will be more and more important. We do not yet have much practice in identifying the underlying values of a multiform story. We will have to learn to notice the patterns displayed over multiple plays of a simulation in the same way that we now notice the worldview behind a single-plot story. Just as we now know how to think about what made Tolstoy propel Anna Karenina in front of that train or what made the producers of Murphy Brown offer her happiness as a single mother, we need to learn to pay attention to the range of possibilities offered us as interactors in the seemingly limitless worlds of digital narrative.

Digital Structures of Complexity

Like every human medium of communication, digital media have been developed to perform tasks that were too difficult to do without them. Hypertext and simulations, the two most promising formats for digital narrative, were both invented after World War II as a way of mastering the complexity of an expanding knowledge base. The mathematician Vannevar Bush put it this way in his landmark 1945 magazine article, "As We May Think": "The summation of human experience is being expanded at a prodigious rate, and the means we use for threading through the consequent maze to the momentarily important item is the same as that used in the days of square-rigged ships" (p. 102).

Bush's solution was "associational indexing" in a kind of magical desk based on microfilm files, a solution he called a "memex" and described as follows:

> The owner of the memex, let us say, is interested in the origin and properties of the bow and arrow. Specifically he is studying why the short Turkish bow was apparently superior to the English long bow in the skirmishes of the crusades. He has dozens of pertinent books and articles in his memex. First he runs through an encyclopedia, finds an interesting but sketchy article, leaves it projected. Next, in a history, he finds another pertinent item, and ties the two together. Thus he goes building a trail of many items. Occasionally he inserts a comment of his own, either linking it into the main trail or joining it by a side trail to a particular item. . . . Thus he builds a trail of his interest through the maze of materials available to him.
>
> And the trails do not fade. (P. 107)

This earliest vision of hypertext reflects the classic American quest—a charting of the wilderness, an imposition of order over chaos, and the mastery of vast resources for concrete, practical purposes. In Bush's view, the infinite web of human knowledge is a solvable maze, open to rational organization.

By contrast, Ted Nelson, who coined the term *hypertext* in the 1960s and called for the transformation of computers into "literary machines" to link together all of human writing, has been more in love with the unsolvable labyrinth. He sees associational organization as a model of his own creative and distractible consciousness, which he describes as a form of "hummingbird mind."[11] Nelson has spent most of his professional life in the effort to create the perfect hypertext system, which he has appropriately named *Xanadu*. He describes this pursuit as a quixotic quest, "a caper story—a beckoning dream at the far edge of possibility that has been too good to let go of, and just too far away to reach, for half my life."[12] Nelson's vision of hypertext is akin to William Faulkner's description of novel writing as a futile but noble effort to get the entire world into one sentence. Those like Nelson who take delight in the intricacies of hypertext, the twisting web rather than the clear-cut trail, are perhaps seeing it as an emblem of the inexhaustibility of the human mind: an endless proliferation of thought looping through vast humming networks whether of neurons or electrons.

The allure of computer simulations comes from a similar attempt to represent complexity. Three years after Bush's suggestion of the memex machine, Norbert Wiener founded the discipline of system dynamics with his book *Cybernetics*. Wiener observed that all systems, whether biological or engineered, have certain characteristics in common, such as the intertwining of multiple cause-and-effect relationships and the creation of feedback loops for self-regulation. Wiener called attention to parallels, for instance, between the way the body keeps a constant internal temperature by instituting changes (like sweating) and monitoring their effects (like feedback on skin temperature) and the way a home thermostat maintains a set temperature. Over the past fifty years, systems thinking has been applied to everything from family structure to frog ponds. It is now commonplace for us to think of the earth itself as a giant ecosystem, in both biological and political terms.

The computer has developed during this time into a versatile tool for modeling systems that reflect our ideas about how the world is organized. Early uses of computer simulations involved putting different values into a constant model and running the system through several "time steps" to see, for example, what would happen to crime statistics five, ten, and fifteen years down the line if police presence went up and cocaine prices went down. These systems were run in batch jobs, which spit out big chunks of numerical data. Other more responsive systems modeled a dynamically changing world open to real-time interaction, like the cockpit simulations used for training airplane pilots. In recent years, computer scientists have designed networked systems that are like a society full of autonomous individuals who talk and work with one another but have no single leader or controller.

In the late 1970s computer system design reached an intriguing milestone with a simple but elegantly conceived program that seemed to simulate life itself. The system is based on a checkerboard grid with markers that are white on one side and black on the other. The markers begin in a random arrangement and are then turned over according to a set of rules that makes decisions based on the colors of a marker's neighbors. Each round of turning causes more turning on the next round, eventually causing remarkable patterns to emerge and move across the board. The Game of Life system does not require a computer, but the patterns look particularly striking on the computer screen, which can run through multiple turns very quickly.[13]

Although no one would claim that such a system is alive in the same way as an animal or plant, it does capture one of the chief attributes of life—the creation of large patterns as a result of many smaller effects. Computer simulations like this are tools for thinking about the larger puzzles of our existence, such as how anything as soulless as a protein can give rise to something as complex as consciousness.

T. S. Eliot used the term *objective correlative* to describe the way in which clusters of events in literary works can capture emotional experience.[14] The computer allows us to create objective correlatives for thinking about the many systems we participate in, observe, and imagine. The rules for artificial life forms can be described as a kind of a game, but the knowledge about the world that the model offers us is not gamelike. It is a behavioral artifact that speaks to one of the most profoundly important aspects of our lives. The more we see life in terms of systems, the more we need a system-modeling medium to represent it—and the less we can dismiss such organized rule systems as mere games.

Current narrative applications overexploit the digressive possibilities of hypertext and the gamelike features of simulation, but that is not surprising in an incunabular medium. As digital narrative develops into maturity, the associational wildernesses will acquire more coherence and the combat games will give way to the portrayal of more complex processes. Participating viewers will assume clearer roles; they will learn how to become orienteers in the complex labyrinths and to see the interpretive shaping in simulated worlds. At the same time as these formal qualities improve, writers will be developing a better feel for which patterns of human experience can best be captured in digital media. In this way a new narrative art will come into its own expressive form.

The process by which this new art form will emerge is already under way and is itself interactive. Each time developers create new genres of digital stories or more immersive games, interactors try them out and grow frustrated or enchanted. Most often these incunabular products arouse expectations they cannot yet fulfill—for more encyclopedic coverage, for greater freedom of navigation, for more direct manipulation of the elements of the story. Every expressive medium has its own unique patterns of desire; its own way of giving pleasure, of creating beauty, of capturing what we feel to be true about life; its own aesthetic. One of the functions of early artifacts is to awaken the public to these new desires, to create the demand for an intensification of the particular pleasures the medium has to offer. Therefore, the next step in understanding what delights or dangers digital narrative will bring to us is to look more closely at its characteristic pleasures, to judge in what ways they are continuous with older narrative traditions and in what ways they offer access to new beauty and new truths about ourselves and the world we move through.

PART II

The Aesthetics of the Medium