Artificial Realities, Virtual Communities, and Knowbots

Chris Dede
George Mason University

ABSTRACT

As the power of information technology grows and its cost diminishes, training is shifting into complex, computer-maintained worlds. The primary reason for this transition is effectiveness; richly detailed virtual environments leverage the most important variables for educational success. These types of training applications can enhance learners' motivation to spend time on task, provide collaborative experiences to foster peer teaching, tailor material to each student's needs and background, and promote the transferability of complicated knowledge and skills into real-world settings. This report describes progress in three aspects of virtual environments that draw on ideas from artificial intelligence: artificial realities, virtual communities, and knowbots.

ARTIFICIAL REALITIES

Artificial realities are computer-maintained environments rich enough in their detail and scope to form virtual "worlds." Two technological trends are driving the evolution of artificial realities as a new medium. First, advances in information technology are giving users enormous, inexpensive power on their desktops. Very fast parallel processing is ideal for enabling sensory interaction (visual, auditory, haptic), as contrasted to the predominantly textual and symbolic representations that characterize current human-computer interfaces.

In particular, the merging of computers and telecommunications enables seamless melding of graphics and video to produce very detailed visual displays. Most of the brain is visual processing system, so opening up this channel into the mind is particularly significant, as researchers in data visualization are discovering. Algorithms that enable the rendering of sophisticated computer graphics are now to the point that, given enough computational power, a high degree of realism is achievable. However, reducing the lag between user actions and the response of the virtual environment is a difficult technical problem that limits the verisimilitude of current artificial realities; this will gradually improve during the 1990s.

Second, the barriers that have kept most people from using sophisticated media are about to collapse. Many potential users who could benefit from information tools avoid attempting to master these devices (e.g. learning to program a videocassette recorder) because they are daunted by the human-computer interface typical of present technology. Incantation-based interfaces are difficult because they demand remembering arcane commands remote from natural language or everyday experience. Graphical user interfaces build some links to common sense reasoning through icons and sensorimotor memory, but still are far from the simplicity of reaching out a finger to set the analog hands on a clock to a particular time. Even experienced users of incantation or graphical interfaces still must do a type of unconscious pretending: imagining that a two-dimensional, small screen image manipulated by some type of input device is really a three-dimensional, full field-of-view environment that can be altered by reaching through the screen into the virtual world.

However facile this pretending, cognitive resources are diverted from the task at hand to sustain the illusion. "Immersion interfaces" (head-mounted display and computerized clothing) use the brain's visual processing capabilities and gesture technology to move beyond pretending into a virtual environment in which even the most naive user can apply common sense and everyday experience to gain a mastery of its surface features. Alternatively, rooms in which multi-directional sensors detect the user's movements and images displayed on the walls respond are a similar visual interface approach that eliminates the need for computerized apparel (Krueger, 1983). Either of these immersion interface strategies opens access to sophisticated information tools to a wide spectrum of users previously discouraged by having to master counterintuitive actions to interact with these devices.

Driven by these advances in the power and the interface of information technology, artificial realities have potential applications in business, entertainment, and education. This broad spectrum of possible uses is crucial to the affordability of a complex technology. Devices prohibitively expensive for a few thousand users can be made much more cheaply for an installed base of over ten million. The gesture glove is a good example; the commercial version marketed in 1987 cost $8000, but in 1989 Mattel released the $80 PowerGlove for Nintendo.

Even though the technical challenges involved are formidable, vendors are making large front-end investments to develop inexpensive artificial reality products because the potential market is so large and spans many different types of applications. Projections of when the different enabling technologies that are converging to empower artificial realities will be widely available and affordable are given in this report's Appendix. A good overview of artificial reality approaches and applications is given in Rheingold (1991).

Leverage for Learning

Environments that mimic, but simplify the real world leverage several variables that enhance learning. If users are motivated by a virtual environment to spend more focussed time on task, that will increase their mastery of the training material. Also, learning situated in contexts similar to the real world setting in which knowledge and skills will be applied is more likely to transfer. Much recent work in cognitive science has centered on "cognitive apprenticeships," characterized by student activity in mastering authentic activities rather than in solving decontextualized problems such as those at the end of textbook chapters (Brown, Collins, & Duguid, 1989). Resnick (1987) describes ways to design learning environments so that they build on people's abilities in real world settings to master new competencies without formal instruction.

Of course, trainers have long used technology-based simulations to mimic simplified real-world environments. However, artificial realities are qualitatively different than simulations because virtual worlds incorporate both immersion and magic as powerful capabilities to leverage instruction. The immersion interface technologies described above must be experienced to be understood; a textual description cannot convey the sensation of interacting with a 3-D, full-visual-field environment. Entering and exploring a virtual world is a more intense mode of human-machine communication than "conversational" interfaces in which the user commands the computer to do something, then waits for a reply (Walker, 1990).

A weak analog to immersion interfaces that many readers will have experienced is the IMAX theater, in which a two-story by three-story screen and high resolution images can generate in the observer strong sensations of motion. Adding three-dimensionality, highly directional and realistic sound, an even wider visual field, and the ability to interact with the virtual world through natural physical actions produces a profound sensation of "being there," as opposed to watching. Because common sense responses work in artificial realities, the learner quickly develops feelings of mastery, rather than the perception of helplessness and frustration typical when first attempting to use a new information tool.

Also, these types of training environments are strongly motivating, since designers of virtual worlds can readily incorporate three powerful levers for emotional involvement: fantasy, challenge, and curiosity (Malone, & Lepper, 1985). Artificial realities can be tailored to the affective needs of individual learners, incorporating fantasy characters matched to the trainee's personality in challenging games or complex adventure-worlds. Of course, these fantasy settings must be carefully designed to suppress the learning of behaviors harmless in a virtual environment, but hazardous in the real world (e.g. walking through a buzz-saw).

In addition to providing through immersion greater mastery and motivation than simulations, artificial realities can support two types of "magic" strongly conducive to learning: visualization and virtual communities. Virtual environments that simply mimic reality have their uses, such as allowing learners to experience activities that are dangerous or expensive in the real world. However, adding the ability to magically act in ways impossible in the real world opens up new dimensions for training. Through visualization, learners can manipulate typically intangible entities such as molecules and mental models; through virtual communities, trainees can interact in rich psychosocial environments populated by simulated beings.

Visualization

Imagine a medical trainee in an artificial reality entering a room labelled "Laboratory." Inside are three types of tangible, manipulable objects. First, the learner can explore the uses of commonplace laboratory devices such as microscopes and centrifuges. Second, the learner can manipulate typically intangible physical objects such as molecules, altering size to perceive in detail their three-dimensional configurations and experimenting with maneuvering two molecules together to understand how one catalyzes a change in the other. Third, the learner can perform similar actions with typically intangible cognitive objects, such as mental models or knowledge structures, looking for geometric patterns that expose the similarities and differences of contrasting theories.

The term applied to rich artificial realities for visualizing information is "cyberspace," from a science fiction novel by William Gibson (1984). The vision of our civilization a couple decades hence that Gibson presents is both intriguing and plausible. The "nervous system" of global business is based on workers manipulating huge virtual structures of data in a shared artificial reality; teleoperation (performing activities over distance) and telepresence (mimicking face-to-face contact through video/graphics representations) dominate human activity. Some "cybernauts" are interested in having life imitate art and are building computational tools that would enable Gibson's imagined future.

Such an endeavor involves grappling with design issues such as the costs and benefits of reifying information, magic versus logic as principles underlying user actions, the presentation of the self in a virtual context for group work, the meaning of travel and action when translated from a physical to a symbolic domain, coordinate systems for (un)real estate, the form and meaning (semiotics) of data objects, and three-dimensional user interface development. Benedikt (1991) discusses the full range of these design challenges; this report will focus only on those issues related to visualization and to virtual communities in training settings.

The fundamental idea underlying visualization is that of displacing cognitive complexity into the human visual system. Human beings have very powerful pattern recognition capabilities for images. Spatial data management systems and scientific visualization have established thatQwhen symbolic meanings are mapped into visible attributes such as shape, texture, size, color, and motionQincreased insight into underlying structural patterns of information is attained (Tufte, 1990).

The virtual medical laboratory outlined above would support two types of visualization embedded in an artificial reality: "sensory transducers" that allow users' eyes, ears, and hands to access previously imperceptible phenomena (such as a molecule) and "cognitive transducers" that perform a similar function for intellectual entities. Sensory transducers provide a means of grasping reality through illusion (Brooks, 1988). Using computers to expand human perceptions (e.g. allowing a medical studentQlike SupermanQto see the human body through X-ray vision) is a powerful method for deepening learners' intuitions about physical phenomena.

Realistic, directional sound is attainable today to enable auditory transducers, and visual transducers already exist through computer graphics. However, tactile forcefeedback is a difficult and expensive hardware challenge; for the early 1990s, cost will likely restrict haptic transducers to applications, such as exploring molecular docking sites, that necessitate people's very delicate and important sense of touch (Minsky et al, 1990). Together, all these types of sensory visualization modalities can combine to intensify immersion in a magical environment of previously intangible physical entities.

A second form of visualization, cognitive transducers, make intellectual entities such as knowledge structures visible and manipulable. They are a logical extension of sensory transducers in enhancing the power of artificial realities for training. Transforming the symbolic into the geometric via data visualization is useful in situations where the amount of data is large and interacting with the data to shape its presentation can aid in interpretation.

A leading-edge illustration of generic data visualization approaches is the Information Visualizer, an experimental interface that uses color and three-dimensional, interactive animation to create information objects (Robertson, Card, & Mackinlay, 1991). Designers of educational "microworlds" (simulations in which the user can change the rules by which the virtual environment functions) also frequently incorporate cognitive transducers. For example, the Alternate Reality Kit allows the user to see and manipulate abstractions such as Newton's Law of Gravity (Smith, 1987).

As one example of cognitive transducers applied to training, the author has conducted preliminary research on the functional design of an information tool that could aid instructional developers (Dede, & Jayaram, 1990). Evaluative imaging in a cyberspace environment is naturally suited to coalescing the complex data typical of a sophisticated training process. A three-dimensional virtual space with icons, shadings, and colors can depict clustered, nested knowledge structures and their interrelationships.

The fundamental metaphor underlying this application for imaging the mental models that underlie training is traversal of virtual cognitive space. The viewer can navigate through a virtual environment populated by cognitive entities represented as physical objects. In addition, the user can transcend the metaphor of physical space by shifting among alternate contexts (e.g. informational, relational, diagnostic) that provide different perspectives on a particular cognitive entity. Special capabilities to minimize complexity (i.e. guided tours, filters, hypermedia) are also available.

However, a major challenge for this mental modeling tool, as for any cyberspace application, is management of the computer-human interface. Complex configurations of multi-dimensional data must be converted to a spatiotemporal analog that allows viewers to transfer their usual strategies for navigation and action from physical contexts to cognitive space. User disorientation has been a particularly difficult problem in cyberspace work because a nonlinear medium offers fewer clues to relative position and orientation than sequential forms of communication. While hypermedia representations offer some leverage by allowing user control over image complexity, large-scale semantic networks can present a cluttered and confusing appearance if viewed in full detail.

However, preliminary work in manipulating the visual presentation of large node/link structures indicates that interface functionalities can provide substantial aid in reducing the complexity of the image. For example, the SemNet tool developed at the Microelectronics and Computer Technology Corporation illustrates how controlled manipulations of associative webs can reveal and conceal complexity by focusing on particular types of categories and relationships. This research could generalize into a theory of semantic navigationQbased on multiple views of information and hierarchical object structuresQthat seems likely to reduce disorientation problems (Fairchild, & Wexelblat, 1989).

In addition, the hypermedia community is developing a variety of alternative strategies for facilitating user traversal through conceptual maps. Nielsen (1990) has constructed a system that uses two navigational dimensions, each with a different animation technique, to aid in orientation. The NoteCards group at Xerox's Palo Alto Research Center has studied how visual cues embedded in the contents of hypertext nodes can create a guided tour through a complex web of information (Marshall, & Irish, 1989). Brown University's Intermedia project has synthesized a variety of navigation approaches into a web-view capability for browsing (Utting, & Yankelovitch, 1989). Whether these approaches can succeed in creating information tools powerful enough to serve as a virtual associational memory is still open to question (Dede, & Palumbo, 1991), but the utility of such a cognitive transducer for trainers would be significant.

A Vignette Illustrating a Hypothetical AR for Training

During the 1990s, the sophistication of sensory and cognitive transducers will steadily evolve. Presenting a vignette of a future training situation can illustrate how artificial realities that include visualization as a type of magic could enhance learning. Many strategies for improving training center on individualization, students' active construction of knowledge, collaborative learning, situated learning, sophisticated evaluation strategies, pedagogical partnerships between instructors and intelligent tools, and the tailoring of content to learning style through the use of multiple representations for knowledge. The following scenario, modified from Dede (1991), depicts a few minutes of two learners interacting in a shared virtual environment that incorporates functionalities to support these training strategies:

Karen sat down at her educational workstation, currently configured as an electronics diagnosis/repair training device. When sign-in was complete, the workstation acknowledged her readiness to begin Lesson Twelve: Teamed Correction of Malfunctioning Communications Sensor. Her "knowbot" (machine-based agent) established a telecommunications link to Phil, her partner in the exercise, who was sitting at a similar device in his home thirty miles away.

"Why did I have the bad luck to get paired with this clown?" she thought, noting a hung-over expression on his face in the video window. "He probably spent last night partying instead of preparing for the lesson." A favorite saying of the problem solving expert to whom she was apprenticed flitted through her mind, "The effectiveness of computer-supported cooperative work can be severely limited by the team's weakest member."

"Let's begin," Karen said decisively. "I'll put on the DataArm to find and remove the faulty component. You use the CT (cognitive transducer) to locate the appropriate repair procedure." Without giving him time to reply, she put on her head-mounted display, brought up an AR (artificial reality) depicting the interior of a TransStar communications groundstation receiver, and began strapping on the DataArm. The reality-engine's meshing of computer graphics and video images presented a near-perfect simulation, although too rapid movements could cause objects to blur slightly. Slowly, she "grasped" a microwrench with her "hand" on the screen and began to loosen the first fastener on the amplifier's cover. Haptic feedback from the DataArm to her hand completed the illusion, and she winced as she realized the bolt was rusty and would require care to remove without breaking.

Meanwhile, Phil called up the CT for Electronics Repair; on the screen, a multicolored, three-dimensional network of interconnections appeared and began slowly rotating. He groaned; just looking at the knowledge web made his eyes hurt. Since the screen resolution was excellent, he suspected that last night's fourth margarita was the culprit.

Phil said slowly and distinctly, "Lesson Twelve," and a trail was highlighted in the network. He began "teleporting" among the nodes of information, simultaneously watching a small window in the upper left-hand corner of the screen that was beginning to fill with data from the diagnostic sensors on Karen's DataArm. Traversing the network at the speed with which Karen was working was difficult, given his hangover, and he made several missteps.

"Knowledge Base," Phil said slowly, "infer what the optical memory chip does to the three-dimensional quantum well superlattice." The voice of his knowbot suddenly responded,"You seem to be assuming a sensor flaw when the amplifier may be the problem." "Shut up!" Phil thought savagely, hitting the cut-off switch. He groaned when he visualized his knowbot feeding the cognitive audit trail of his actions into the workstations of his trainer and the corporation's communications repair expert; he could not terminate those incriminating records. Phil cringed when he imagined his trainer's "avatar" (a virtual representation of another person) giving him another lecture on his shortcomings.

Mentally, he began phrasing an excuse to send his instructors via email at the end of the lesson. Meanwhile, Karen was exasperatedly watching the window on her AR display in which Phil's diagnostic responses should have been appearing. "He's hopeless," she thought. Her knowbot's "consciousness sensor" (a biofeedback link that monitors user attention and mood) interrupted with a warning: "Your blood pressure is rising rapidly; this could trigger a migraine headache." "Why," Karen said sadly, "couldn't I have lived in the age when students learned from textbooksI"

Such a training vignette may seem science fiction and, in terms of pedagogical and psychosocial barriers, its occurrence by early in the next century is unlikely. However, virtual environments for training such as the vignette depicts are quite feasible technically and economically in that time frame. Research into design issues and analyses of relative cost/benefit are necessary to understand the potential utility of these environments for sophisticated instruction.

For example, as cyberspace environments become more commonly used, trainers will face new types of instructional design challenges. Just as current training systems must match a mixture of textual, auditory, visual, and psychomotor presentations to the student's learning style, so artificial realities must balance representations of virtual physical objects, sensory transducers, and cognitive transducers. Developing rhetorics for transducing typically intangible physical and intellectual entities is also a major challenge.

Visualization is one form of magic that empowers learning in artificial realities. Beyond visualization, a second type of artificial reality magic suggested by the vignette above is virtual communities. Learners can interact in rich machine-mediated psychosocial environments populated both by video-links to other people and by simulated beings. These simulated beings may be avatars (computer graphics representations of people) or knowbots (machine-based agents); each adds an important dimension to training in artificial realities.

VIRTUAL COMMUNITIES AND KNOWBOTS

Enabling people to communicate effectively across barriers of distance and time by forming virtual communities is an emerging goal for information technology. During the last decade, the field of computer-supported cooperative work (CSCW) has established that developing productive technology-mediated human interaction involves complex psychosocial issues that extend well beyond earlier models of simply transmitting data along a channel from sender to receiver (Greif, 1988). Communication depends on affective as well as cognitive interchange; researchers studying collaborative educational environments are just beginning to conceptualize the complicated emotional dynamics of peer teaching and learning.

During the 1990s, wide-area, broadband information infrastructures such as the National Research and Education Network (NREN), will increasingly link training sites. Learners' adept use of email, voicemail, videomail, and direct audio/video can improve the spectrum of expertise available in their training. However, the types of task structures and interpersonal relationships required for effective usage of these media differ from those conventionally utilized in face-to-face interactions (Sproull, & Kiesler, 1991). Training learners to use technology-mediated communication effectivelyQwhether over email or in artificial realitiesQinvolves more than simply mapping skills from face-to-face encounters into the virtual environment (Dunlop, & Kling, 1991).

As discussed earlier, situating training in virtual contexts similar to the environments in which learners' skills will be used helps their knowledge to transfer. When the material involved has psychosocial as well as intellectual dimensions, the design of authentic experiences to embed in artificial realities for training becomes more complex. In addition to physical and cognitive entities, instructional developers can include simulated beings (avatars and knowbots) in the virtual environment.

One example of such a training application involves software engineering education; students are trained in a technical process, code inspection, that is one stage of a formal methodology for software development (Stevens, 1989). Using hypermedia, Digital Video Interactive (DVI), and rule-based expert systems, the Advanced Learning Technologies Project at Carnegie Mellon University has created a virtual environment similar to a typical corporate setting. The trainee interacts with this artificial reality in the role of a just-hired software engineer still learning the profession. Through direct instruction and simulated experience, the student practices the process of formal code inspection.

The learner can access various rooms in the virtual software company, including an auditorium, library, office, training center, and conference facility. Machine-based agents (knowbots) that simulate people, such as a trainer and a librarian, facilitate the use of resources to learn about the code inspection process. Via specialized tools in the office, the student can prepare for a simulated code inspection, in which he or she can choose to play any of three roles out of the four roles possible in this formal software review process. For each inspection, a rule-based expert system utilizes DVI technology to construct knowbots that simulate the three roles not chosen by the learner. This knowledge-based system controls the topic of conversation; determines who should speak next; and models the personalities of the knowbots in the inspection meeting, altering their cognitive and affective perspective depending on what is happening.

The learner uses a menu-based natural language interface to interact with these simulated beings, who model behaviors typical in code inspection situations. The student not only can choose from a wide range of options of what to say, but can determine when to make remarks and can select the emotional inflection of his or her utterances, from a calm passive tone to an angry aggressive snarl. By mimicking the reactions likely from human participants in a real simulation, the knowbots provide the learner with a sense of the strengths and weaknesses of different intellectual/psychosocial strategies for that role in a code inspection.

Without using artificial realities and knowbots, this type of authentic experience is very difficult to simulate in classroom settings. Not only is the training environment dissimilar from the corporate context in which software development skills will be used; but also students do not know how to roleplay exemplary, typical, and problematic participants in code inspections. Through knowbots, the instructional designer can provide paradigmatic illustrations of how to handle a variety of situations, without the expense of having teams of human actors perform for each individual learner.

Distributing the intelligence incorporated into artificial realities among simulated beings raises challenging instructional design issues. The quality of what students learn is determined by the accuracy with which a knowledge-based system can simulate not only human reasoning, but also people's personalities and emotions. This is a very difficult task, necessitating knowledge acquisition about psychosocial dynamics as well as intellectual issues.

Another design issue posed by knowbots is whether trainees will profit more by interacting with a variety of partial machine-based intelligences (a collaborative learning/peer teaching perspective) or with a single, omniscient source of knowledge (an intelligent tutoring systems perspective). Students may empathize more readily with flawed simulated beings (e.g. a "good, but still learning" moderator for code inspections) than with an intelligent coach who has mastered every aspect of the training domain. Programming several machine-based agents with highly focussed skills is also easier than constructing a single, broad-based knowbot; knowledge-based systems work best for narrow, well-specified domains.

Beyond simulating virtual people in artificial realities, widespread usage of embedded machine intelligence in the form of knowbots can enhance training by making the learners' information processing activities more efficient. The scenario earlier suggested some ways in which knowbots could accomplish this purpose, such as facilitating electronic linkages among learners and instructors. As a general purpose agent, a knowbot can preprocess incoming electronic mail so that messages relating to different topics are automatically grouped into their appropriate folders. The agent can also automatically answer email messages proposing meetings by using a scheduling application that accesses every the calendar of every person on the network. Journals arriving in electronic form can be scanned by the knowbot for articles of particular interest to the user, selecting and filing that material for future reading (Lai, & Malone, 1988).

At its extreme, the knowbot could become an intelligent partner in mediating human communication. Each person in a virtual community could manage interaction with others via an intelligent device that recognizes speech, takes messages through digital voice input/output, filters phone calls, and scans databases (weather, traffic, airport) to aid human planning (Tesler, 1991). The intelligence, telepresence, and even "simulated personality" of each person's device could pose normative issues about what type of relationship to establish with a socially aware machine.

Avatars

The simulated beings in artificial realities for training need not all be machine-based agents. Military and aerospace research projects have created telerobotic technologies that allow operators to control devices across distance (Fisher, 1990). These enable the creation of training environments that provide a virtual link between remote learners and real world settings; for example, a student can manipulate industrial process technology without physically traveling to the factory (Uttal, 1989). An extension of teleoperations is telepresence, which allows people to interact across distance via avatars (computer-graphics representations), as opposed to direct video-links.

For example, SIMNET (Orlansky, & Thorp, 1991) is a training application that creates a virtual battlefield on which learners at remote sites can simultaneously operate military equipment. Complex data-objects that indicate changes in the state of each piece of equipment are exchanged via a telephone network interconnecting the training workstations ("dial-a-war"). The appearance and capabilities of graphics-based avatars representing military equipment alter second-by-second as the virtual battle evolves. Through this telepresence approach, a widely distributed group of personnel can engage in simulated real-time warfare without the necessity of gathering the participants at a single site to conduct combat.

Some of the representations in the SIMNET artificial reality are avatars controlled by human beings; others are knowbots whose actions are generated by a knowledge-based system under general human guidance. Similarly, the virtual world for code inspections described above could be redesigned to incorporate both avatars controlled by human agents and machine-based knowbots. For any given artificial reality for training, balancing the different types of beings interacting in the virtual environment (avatars, knowbots, video-links) is an interesting issue for instructional design research.

As users interact in virtual environments, intriguing interpersonal dynamics ermerge that are quite different from typical real-world encounters. People participating in artificial realities often feel as if knowbots are real human beings, an illustration of the general principle that users tend to anthropomorphize any type of machine-based agent. Joseph Weitzenbaum's Eliza program, which simulates a Rogerian therapist, is an example of this tendency; some users type in responses to Eliza's prompts for hours, attributing human understandings to a pattern-matching natural language program that has no conception of people's emotions or behaviors.

As a complement to responding to knowbots as if they were human, participants in a virtual world interacting via avatars tend to treat each other as imaginary beings. An intriguing example of this phenomenon is documented in research on Lucasfilm's Habitat (Morningstar, & Farmer, 1991). Habitat was initially designed to be an on-line entertainment medium in which people could meet in a virtual environment to play adventure games. Users, however, extended the system into a full-fledged virtual community with a unique culture; rather than playing pre-scripted fantasy games, they focussed on creating new lifestyles and utopian societies.

As an entertainment-oriented cyberspace, Habitat provided participants the opportunity to get married or divorced (without real-world repercussions), start businesses (without risking money), found religions (without real-world persecution), murder other's avatars (without moral qualms), and tailor the appearance of one's own avatar to assume a range of personal identities (e.g. movie star, dragon). Just as SIMNET enables virtual battles, Habitat and its successors empower users to create artificial societies. What people want from such societies that the real world cannot offer is magic, such as the gender-alteration machine (Change-o-matic) that was one of the most popular devices in the Habitat world.

Users learned more about their innermost needs and desires by participating in Habitat than they would have by spending an equivalent amount of time listening to psychology lectures. Similarly, social scientists are discovering more about utopias by studying Habitat's successors than they did by researching communes, which were too restricted by real-world considerations to meaningfully mirror people's visions of ideal communities. Giving users magical powers opens up learning in ways that trainers are just beginning to understand. As with any emerging medium, first traditional types of content are ported to the new channel; then alternative, unique forms of expressionQlike HabitatQare created to take advantage of expanded capabilities for communication and education.

If the medium itself is the message, as Marshall McLuhan claimed, how are artificial realities shaping their inhabitants? The cultural consequences of technology-mediated psychosocial environments are mixed. On the one hand, people have a wider range of vicarious experience and more contact with specialized human resources than they could attain through direct interaction locally; this empowers better forms of training. On the other hand, even with a low-immersion, passive medium like television, many people's perceptions of family life now come primarily from situation comedies, of crime from police shows, of sexuality from soap operas, and of physical exercise from videogames. When the emerging "reality-industrial complex" becomes a major source of human experience, then society risks retreating from reality into immersive illusion and escapismQyet simultaneously has access to a more powerful medium for learning than any previously available.

In Schrage (1990), Alan Kay analyzes the implications of new media through the question, "What does a medium ask you to become in order to use it?" Print requires a rational reader; television, a passive observer; the telephone, a conversationalist. Unlike passive, linear media such as television, artificial realities center on immersion and collaboration. Training in the future may balance learning experiences in three types of environments: contemplation and introspection via books and symbolic media, immersion and collaboration in virtual communities populated by knowbots and avatars, and traditional student/instructor interaction in conventional teaching/learning settings.

CONCLUSION

Technology evolves in waves of innovation and consolidation. The advent of "motion pictures" about a century ago ushered in civilization's fourth medium, another dimension to communication beyond spoken language, written language, and still images. Later, new technologies appeared to embellish the capabilities of moving images: broadcast and narrowcast television, videotapes, videodiscs, multimedia, hypermedia. Now all of these are merging into a synthesis so far beyond its individual components that it constitutes a new medium: artificial realities.

Part of the educational implications of this medium center around its channel, which is rich and powerful enough to mimic the meta-medium in which we live, the real world. Other instructional implications come from the content that cybernauts are embedding into this channel: sensory and cognitive transducers, virtual communities made up of people's avatars and of machine-based knowbots. Together, channel and content form the message of this new, immersive medium, which is still too indistinct to fully comprehend, but appears both fascinating and frightening.

Any powerful information technology is a double-edged sword: a source of either propaganda or education. Through advances in information technology, virtual environments can now be created that seem intensely real to participants, yet may be false to the true nature of reality in the same way that fractally-generated mountain ranges are not valid depictions of physical topography and geology. Artificial realities, virtual communities, and knowbots are emerging technologies that have enormous potential to improve training. However, trainers and learners must recognize that these instructional vehicles carry intrinsic content that can empower or subvert the goals of an educational experience. Careful research is needed to understand how to optimize the design and utilization of virtual worlds for training.

BIOGRAPHY

Chris Dede is Professor of Education and of Information Systems at George Mason University, where he directs the Center for Interactive Educational Technology and is on the core faculty of The Institute for Public Policy. He has been a Visiting Scientist at the Computer Science Lab, MIT and at NASA's Johnson Space Center. Chris also serves as a consultant on emerging technologies to a variety of corporations and organizations.

REFERENCES

Benedikt, M. (1991). Introduction. In M. Benedikt (Ed.), Cyberspace: First steps (pp. 1-
25). Cambridge, MA: MIT Press.

Brooks, Jr., F.P. (1988). Grasping reality through illusion: Interactive graphics serving science. CHI '88 Proceedings (pp. 1-11). Reading, MA: Addison-Wesley.

Brown, J.S., Collins, A., & Duguid, P. (1989). Situated cognition and the culture of learning. Educational Researcher 18(1), 32-42.

Dede, C.J. (1991). Emerging information iechnologies: Implications for distance learning. Annals of the American Academy for Political and Social Science 514, 146-158.

Dede, C.J., & Jayaram, G. (1990). Designing a training tool for imaging mental models (AFHRL Technical Report CR-90-80). Brooks AFB, TX: Air Force Human Resources Laboratory.
Dunlop, C., & Kling, R. (1991). Social relationships in electronic communities. In C. Dunlop & R. Kling (Eds.), Computerization and controversy: Value conflicts and social choices (pp. 322-330). New York: Academic Press.
Fairchild, K.F., & Wexelblat, A. (1989). Navigation through cyberspace (MCC Technical Report Number STP-161-89). Austin, TX: Microelectronics and Computer Technology Corporation.
Fisher, S. (1990). Virtual environments, personal simulation, & telepresence. Multimedia review: The journal of multimedia computing 1(2), 34-44.
Greif, I. (Ed.). (1988). Computer-supported cooperative work: A book of readings. San Mateo, CA: Morgan Kaufmann.
Krueger, M. (1983). Artificial reality. Reading, MA: Addison-Wesley.
Lai, K., & Malone,T. (1988). Object lens: A "spreadsheet" for cooperative work. Proceedings of the 1988 Conference on Computer-Supported Cooperative Work. (pp. 115-124). New York: Association for Computing Machinery.
Malone, T.W., & Lepper, M.R. (1985). Making learning fun: A taxonomy of intrinsic motivations for learning. In R.E. Snow & M.J. Farr (Eds.), Aptitude, learning, and instruction: III. Conative and affective process analysis (pp. 176-189). Hillsdale, NJ: Lawrence Erlbaum.
Marshall, C.C., & Irish, P.M. (1989). Guided tours and on-line presentations: How authors make hypertext intelligible for readers. Proceedings of Hypertext '89 (pp. 321-328). New York: Association for Computing Machinery.
Minsky, M., Ouh-Young, M., Steele, O., & Brooks, F.P. Jr. (1990). Feeling and seeing: Issues in force display. Computer graphics: Proceedings of the 1990 symposium on interactive 3-D graphics 24(2), 235-244.
Morningstar, C., & Farmer, F.R. (1991). The lessons of Lucasfilm's Habitat. In M. Benedikt (Ed.), Cyberspace: First steps (pp. 273-302). Cambridge, MA: MIT Press.
Nielsen, J. (1990). The art of navigating through hypertext. Communications of the ACM 33(3): 298-310.
Orlansky, J., & Thorp, J. (1991). SIMNET Q An engagement training system for tactical warfare. Journal of Defense Research 20(2), 774-783.
Resnick, L. (1987). Learning in school and out. Educational Researcher 16 (10), 13-20.
Rheingold, H. (1991). Virtual Reality. New York: Simon and Schuster.
Robertson, G., Card, S., & Mackinlay, J. (1991). The information visualizer. Proceedings of CHI 91 (pp. 181-188). New York: Association for Computing Machinery.
Schrage, M. (1990.) Shared Minds: The New Technologies of Collaboration.. New York: Random.
Smith, R.B. (1987). Experiences with the alternate reality kit: An example of the tension Between literalism and magic. Proceedings of CHI+GI 1987 (pp. 324-333). New York: Association for Computing Machinery.
Sproull, L., & Kiesler, S. (1991). Computers, networks, and work. Scientific American 265(3), 116-127.
Stevens, S. (1989). Intelligent interactive video simulation of a code inspection. Communications of the ACM 32(7), 832-843.
Tesler, L.G. (1991). Networked computing in the 1990s. Scientific American 265(3), 86-93.
Tufte, E.R. (1990). Envisioning information. Chesire, CN: Graphics Press.
Uttal, W.R. (1989). Teleoperators. Scientific American 261(6), 124-129.
Utting, K. & Yankelovitch, N. (1989). Context and orientation in hypermedia

networks. ACM Transactions on Information Systems 7(1): 58-84.
Walker, J. (1990). Through the looking glass. In B. Laurel (Ed.), The art of computer-human interface design (pp. 213-245). Menlo Park, CA: Addison-Wesley.

APPENDIX: THE EVOLUTION OF EDUCATIONAL TECHNOLOGY

For at least another fifteen years, the information technologies will continue to increase in power while decreasing in cost. The table below presents a forecast of when some technological capabilities useful for learning will be available on high-end personal computers and in telecommunications infrastructures.

Functionality	Uses	Time Frame
Hypermedia (nonlinear traversal of multi-media information)	Interlinking of diverse subject matter; easier conceptual exploration, training, collaboration	Now
High quality voice synthesis	Auditory natural language output	Now
Cognitive audit trails (automatic recording of user actions)	Support for finding patterns of suboptimal performance	Early 1990s
Advanced manipulatory input devices (e.g. gesture gloves with tactile feedback)	Mimetic learning which builds on real world experience	Early 1990s
Optical-disc systems with multiple read/ write and mixed-media capabilities	Support of large data and knowledge bases; very cheap secondary storage; facilitation of artificial realities	Early 1990s
High-bandwidth fiber-optic networks	Massive real time data exchange	Early 1990s
Synthesis of computers, telecommunications	Easy interconnection; realistic simulation	Early 1990s
Standardization of computer and telecommunications protocols	Easy connectivity, compatibility; lower costs	Mid 1990s
Sophisticated User Interface Management Systems	Easier development of instructional applications; reduced time for novices to master a program	Mid 1990s

Intelligent computational agents embedded in applications	Support for user-defined independent actions; knowbots	Mid 1990s
Computer-supported cooperative work (collaborative design, collective problem solving, group decision support)	Mastery of team task performance; virtual communities	Mid 1990s
User-specific, limited-vocabulary voice recognition	Restricted natural language input	Mid 1990s
Information utilities (synthesis of media, databases, and communications)	Access to integrated sources of data and tools for assimilation	Mid 1990s
High-resolution, flat-panel color monitors with 3-D graphics	Vivid simulation of reality; easy reading of text	Late 1990s
Microworlds (limited, alternate realities with user control over rules)	Experience in applying theoretical information in practical situations; sensory transducers	Late 1990s
Consciousness sensors (input of user biofeedback into computer)	Monitoring of mood, state of mind	Late 1990s
Microcomputer performance equivalent to current supercomputers	Sufficient power for simultaneous advanced functionalities	Late 1990s
Knowledge processing (contextually-linked data storage with embedded inference)	Goal-oriented, context-specific mastery of concepts and skills; cognitive transducers	Late 1990s
Intelligent tutors and coaches for restricted domains	Models of embedded expertise for greater individualization	Year 2000+
Artificial Realities (computer-maintained virtual worlds)	Intensely motivating simulation and experience	Year 2000+

The scenarios in the report illustrate how these emerging functionalities might be used in future real-world contexts.