

Learning in Artificial Environments: Embodiment, Embeddedness and Dynamic Adaptation

WILLIAM WINN*

412 Miller, Box 353600, University of Washington, Seattle, WA, 98195, USA

Reacting against the inadequacy of traditional cognitive theory to explain how learning occurs, many educational researchers have turned to a socio-cultural, situated model of learning within which to conduct their research. However, this model has, in its turn, failed to account for some of what is observed when students work with complex, computer-supported simulations of natural environments, referred to as “artificial environments.” What is more, traditional cognitive theory has continued to evolve and, considered together with systems theory and cognitive neuroscience, is now in a better position to provide an adequate account of learning. This article brings together three ideas to form a conceptual framework for studying learning in artificial environments. These are the ideas that cognition is embodied in physical activity, that this activity is embedded in a learning environment, and that learning is the result of adaptation of the learner to the environment and the environment to the learner. The conceptual framework assumes that embodiment, embeddedness and adaptation are completely interdependent. These ideas are illustrated from research on artificial environments, particularly those that use virtual reality technologies.

Keywords: Learning, artificial environments, system theory, cognition.

INTRODUCTION

Over the last two decades, educational researchers, who rely on cognitive science to guide their work, have felt frustration at the intractability of many

*e-mail: billwinn@u.washington.edu

of the problems cognitive scientists were once expected to solve quickly and easily. The most cited example is the failure of research and development in artificial intelligence to produce effective and cost-effective systems that can reason as people do (Dreyfus, 1979; Dreyfus & Dreyfus, 1985). The perceived failing of computational, AI-like, views of cognition is mostly attributed to their being based on a model of thinking that assumes: 1) Knowledge is represented in the mind by symbols that have a one-to-one mapping with phenomena in the environment, 2) Cognition consists of operating on these symbols by applying computer-like algorithms, 3) Cognitive activity is therefore separated from the learner's context; it goes on in the mind, not in the environment, 4) Computed solutions to problems appear as behavior that reduces discrepancies between a learner's knowledge and actions, and "correct" ways to understand and act in the world.

Researchers have claimed, with justification, that human thinking does not work like this (Streibel, 1991; Winograd & Flores, 1986). For example, some argue that the cognitive system is closed to information and does not allow direct mapping of external events to internal symbol systems (Bickhard, 2000; Maturana & Varela, 1987). Others make the case that cognitive activity is contextually situated, which is to say that its workings and outcomes are determined by the external environment in which cognition takes place (Brown, Collins & Duguid, 1989; Lave & Wenger, 1991) and that thinking is uniquely influenced by a person's activity in the environment in which thinking occurs (Clark, 1997; Varela, Thompson & Rosch, 1991). Still others claim that knowledge is constructed by the learner, not received from a teacher, and that it is consequently not possible, nor even worthwhile, to predict, and therefore prescribe, what and how a student will learn (Bednar et al., 1992; Cunningham, 1992).

When educators turned away from a computational view of cognition as their framework for research, there were a number of paths they could have taken. The path chosen by the majority led away from the direct study of cognitive processes towards examination of the context and culture within which cognition takes place. This approach falls, somewhat loosely, under the rubrics of "situated cognition" and "constructivism." Another path, not taken by the majority, was one that leads to a more rigorous examination of learning processes. This path requires the study of interaction and mutual adaptation of students and environments, considered to act like complex systems, and to explanations of cognition emerging from the neurosciences. The assumptions here are different, namely that: 1) Students and the environments in which they learn are much more complicated than we assumed

when we developed the constrained, AI-like view of cognition. 2) To understand learning requires that this complexity be faced head on, challenging us to analyze and directly describe learning processes, regardless of their complexity. 3) This analysis is potentially feasible given recent developments in system theory, which is understood to include dynamics, robotics, models of adaptation, and theories of self-organization and emergence. 4) Explanations of learning and cognition can be reduced no further than those emerging from the cognitive neurosciences which, after “the decade of the brain,” are coming closer to accounting for some of the behaviors we observe in learners.

This article makes the case that it is necessary and feasible to take this more challenging path, and to return to the direct study of learning with the help of system theory and neuroscience. It begins with an examination of the assumptions of constructivism as educators have understood it. Then, it sketches an alternative framework based on the assumptions that learning occurs when people adapt to their environment. To understand adaptation, we must think of the learner as embedded in the learning environment and physically active in it, so that cognition can be thought of as an embodied as well as a cerebral activity. These arguments are illustrated from research on learning in computer-supported artificial environments, particularly from research by our research group (Windschitl & Winn, 2000; Winn et al., 2002) on a simulation of the salinity and tidal currents in Puget Sound, Washington, called “Virtual Puget Sound.”

CONSTRUCTIVISM

“Constructivism” is the name commonly given to a cluster of ideologies and practices that have largely replaced the computational view of cognition and positivist methodology in educational research. Constructivism has its own set of assumptions about thinking and learning that are quite different from those that underlie the computational view. For educational technologists, these are laid out by Duffy and Jonassen (1992). Briefly stated, understanding is constructed by students, not received in messages simply to be encoded, remembered and recalled. How knowledge is constructed and with what results depends far more on a student’s history of adaptations to the environment (Maturana & Varela, 1987) than on particular environmental events. Therefore, learning is best explained in terms of the student’s evolving, contextualized understanding and is valued on that criterion, rather than on the basis of traditional objective assessments. The notion that what stu-

dents learn may be shaped in any significant way by preplanned instructional interventions is played down if not rejected out of hand. Finally, constructivists stress the social nature of learning. Knowledge is not constructed in a vacuum, but through the negotiation of meaning within groups of people.

There is nothing inherently wrong with any of these ideas. In fact, the emphasis on the context in which learning takes place and on learning's social and constructed nature have drawn the attention of educational researchers to topics that must be studied if we are to understand, more broadly, what happens to students as they go through the education system. These include students' socio-economic status, their ethnicity, the extent of family support for learning, the quality and preparation of their teachers, the fabric of their schools, and so on.

Yet the intense focus of recent educational research on these topics has led to the neglect of more basic research that extends our understanding of how students learn. What is more, much of the same cognitive research that followed the computational model, against which constructivists reacted, has evolved to a point where it can explain more complex aspects of learning, while retaining scientific rigor, and while remaining centrally focused on basic cognition and learning. Here are three examples:

1. The closure of the central nervous system to information has been the basis for arguments against the idea that the environment is directly mapped onto symbolic representations in the mind, and against the idea that cognition consists of operating on these symbols algorithmically. This has led some to question the nature of mental representation, even to doubt that representation, as traditionally construed, occurs at all (Bickhard, 2000, pp. 38-42; Rosch, 1999; Skarda & Freeman, 1987; Thelen & Smith, 1994). Yet people have persuasively argued for the importance, indeed the necessity, of representation to cognition. Without mental representation, how can we reason in the absence of the objects and events that created memories for them in the first place (Haugeland, 1991)? Without mental representation, how can we develop and reason with abstractions that are essential for developing general skill in many disciplines, like mathematics (Steffe, 2000)? Our idea of mental representation needs to be redefined, not rejected out of hand. Cognitive neuroscience offers one alternative. Since learning affects the ways in which neurons are connected and the mechanisms that create those connections (Markowitsch, 2000), we might consider mental representation to be networks of associations among neurons.

2. Cognition *is* computational (Clark, 1997, pp. 153-160), though in different ways than those proposed by proponents of the “AI” view of cognition. Dietrich and Markman (2000) describe this newer view of computation.

Computational accounts can explain the capacities of a system to exhibit certain behaviors. They do not describe causal laws; rather they describe what the system—student—is capable of doing when certain conditions arise. The behavior that the system exhibits is analogous to the result of computing a function whose parameters are inputs to the system. Marr’s (1982) account of vision is a good example of how the behavior of a complex system—human vision—can be explained through mathematical functions, without needing to claim that the functions correspond to causal mechanisms in the central nervous system. (It was the failure to understand this distinction between functional and causal models that got us into trouble in the first place. Boden’s [1988] “Computer models of mind” is a case in point.)

Computational accounts can describe how entire systems behave. They do this by positing collections of interacting sub-functions that, together, explain the system’s behavior. What is more, the system’s behavior can be computed at different levels of granularity, to provide descriptions of single processes or of their collective effect. The ability of computation, thus conceived, to explain behavior systemically appears as early as in Von Bertalanffy’s (1968) equations describing metabolism, which lie at the roots of General System Theory, and continue today in current system-theoretic accounts of cognition (Port & Van Gelder, 1995).

Computational accounts can interpret what a system is doing as it changes from one state to another. They are grounded in the system’s behavior, not in a mathematical model established *a priori*. The role of research is to find or develop functions that describe changes in the system, not to find systems that behave according to given functions. It is therefore important to draw a distinction between “scientific AI” and “engineering AI” (Dietrich & Markman, 2000). The former is used to model, descriptively, and to interpret how systems behave. The latter attempts to use established AI techniques, prescriptively, to solve problems.

3. The idea that cognitive activity depends on the context in which it takes place has been used as an argument for the ineffectiveness of instructional strategies that are employed uniformly with different kinds of students and in different contexts. Yet organisms and agents that are, in many

respects, heterogeneous, change in predictable, though complex ways as a result of interactions with their surroundings. Accepting the idea that cognitive activity occurs in an environment is not grounds, therefore, for rejecting causal models of adaptation as the basis for cognition. The required accounts are certainly more complex than previous accounts. However, cognition can be described using models of biological adaptation (Holland, 1992, 1995) and dynamical system theory (Beer, 1995; Kulikowich & Young, 2001; Port & Van Gelder, 1995) that are capable of handling the complexity of the interactions between entire environments and complete cognitive systems.

There appear, then, to be good reasons to reconsider the traditional views of representation, computation and predictability as alternatives to constructivist views of learning. Mental representations of the world are real and necessary for cognition. Rather than “pictures in the head,” these might now be thought of in terms of associative networks, which have a neurological basis and which are activated by sensory inputs. The experiences that created them in the first place are re-experienced, and restructured, as a consequence of this activation, as Farrah (2000), for example, has demonstrated for mental images. Also, cognition does operate through computation. This does not mean that the brain works like a digital computer. But cognition works by taking the state of the learner and the environment as “input,” performing operations on it, and “outputting” the result as behavior. Finally, situating learning in a context does not mean that there is no regularity in the processes or outcomes of learning. Techniques exist that allow us to describe the dynamic interactions of two complex systems - student and environment - as each changes the other.

The next section examines ideas that have, at one time or another, been used to challenge the centrality of cognitive science to educational research and practice. The notion that cognition is “embodied” and “embedded” has, rightly, been used to challenge the traditional separation among brain, body and world. Yet careful examination of these ideas does not force us to conclude that representational and computational theories of cognition should be rejected in favor of predominantly socio-cultural ones. Acceptance of embodiment and embeddedness leads, in turn, to the idea that learning is adaptation to an environment, where “adaptation” retains a large measure of its biological sense, as developed, for example, in the theories of Darwin. Taken together, the three concepts, embodiment, embeddedness and adaptation, form a viable integrated theoretical framework within which to study learning in artificial environments.

AN ALTERNATIVE FRAMEWORK FOR DESCRIBING LEARNING IN ARTIFICIAL ENVIRONMENTS

Embodied and Embedded Cognition

Once we start to think of cognition as the interaction between a person and their environment, it is necessary to consider how that interaction occurs. This, in turn, requires the consideration of how our physical bodies serve to externalize the activities of our physical brains in order to connect cognitive activity to the environment. This physical dimension of cognition is referred to as “embodiment.” Once this direct connection between cognitive action and the environment is established, we must acknowledge that cognitive activity is far more closely coupled to the environment than many have hitherto acknowledged. This interdependence of cognition and environment is referred to as “embeddedness.” It is clear that the embodiment of cognition in physical action and the embeddedness of cognition in the environment are closely connected. In this section we examine what this implies for learning in artificial environments.

The role of our bodies in learning

To say that cognition is embodied is to say that it involves our entire bodies, not just our brains. We can think of cognition as being embodied in three ways. First, somewhat trivially, the brain is an organ like any other in our body. Though complex, it does not possess magical or mystical powers (Pinker, 1997, p. 64). We understand something about how it develops and works. We do not yet understand enough to paint a complete picture of cognition using its palette. However, we do know enough to assert that cognitive neuroscience is important for educators to study, both as a source of explanations about what happens to students as they learn, and as a source of prescriptions for instruction (Berninger & Richards, 2002), although the latter are, for now, commonly confined to clinical cases.

Second, we must consider how cognition operates within constraints imposed by our physiology. The bandwidth of the data we can detect in the environment is limited. We cannot see beyond infrared and ultraviolet wavelengths in the electromagnetic spectrum. We cannot hear sounds below 20 hz. or above 20,000 hz. Our natural view of the world is therefore considerably limited. Also, we experience the world at particular scales in time and space. Sparrows are small, elephants are large. Meteorites move quickly, glaciers slowly. We can use technology to reduce the limits imposed by our sensory bandwidth. We can show a bat’s echolocation as a sonogram. Or we




FIGURE 1 Virtual Puget Sound, showing the “arrow” metaphor for tidal currents. The view is to the north west, looking up the main channel towards Admiralty Inlet, Port Townsend and the Straits of Juan de Fuca. See Color Plate I at back of issue.

can use animation to show the movement of a glacier by compressing time to, say, 50 years per second. However, we can never really know the world as a bat does (Nagel, 1974). Nor can we understand the advance and retreat of ice sheets by watching them in real time. Artificial environments can use computer technology to create metaphorical representations in order to bring to students concepts and principles that normally lie outside the reach of direct experience.

Research on artificial environments has had to confront this issue head on. In one sense, the limitations imposed by our physiology offer an advantage. They force students to deal with the world in ways “real” scientists must—by making inferences from indirect, instrumented observations of phenomena (Winn & Windschitl, 2000). Computer-created environments can act as transducers of data that lie beyond direct sensory detection. Students can use virtual instruments to measure environmental phenomena. They can also experience phenomena through metaphorical visualizations. But, using metaphors, one runs the risk of misinterpretation. Incomplete understanding of readouts from virtual instruments, or of the meaning of metaphors generally, can lead to misconceptions. For example, Winn et al. (2001) used vectors to show the speed and direction of currents in a simulation of selected aspects of physical oceanography. Longer arrows showed that the current was faster. (Figure 1 shows this environment, “Virtual Puget Sound.”) To solve a problem, students had to learn that water speeds up when it moves through narrow channels. However, one student concluded that currents were slower in narrow passages. For him, the longer arrows made the passages look more clogged. A clogged highway slows traffic down. So water slows down in narrow passages. The metaphor meant the

opposite of what it was intended to describe for that student.

Another case shows a problem with metaphorically changing temporal scale. Jackson (2000) described a study in which an artificial environment simulated global warming over a period of 2000 years. Middle school students could control the amount of greenhouse gases in the atmosphere by reducing emissions from vehicles and factories, and by increasing the amount of green plant material available to absorb CO₂. In an attempt to connect global warming to destruction of the rain forest, Jackson used the number of trees as a metaphor for this process. As students changed the amount of green plant material, trees appeared and disappeared in the environment. Students made measurements, traveled into the future and made more measurements, returned to the present to make changes to the environment, including planting or harvesting trees, and visited the future again to see what effects their changes had made. Several students concluded that global warming was not a problem: All we have to do is plant more trees! Time travel was almost instantaneous and the time scale therefore much distorted. As a result, planting trees had an immediate effect on global warming. The tree metaphor failed to convey the fact that restoring the environment takes a long time, and this simplification led to a misconception.

It is clear that having to invent ways of showing phenomena that have no natural appearance for humans can be problematic. For example, there are no hard-and-fast rules for deciding what salinity looks like, nor, therefore, for interpreting symbols chosen to represent salinity. Where visual metaphors do exist, they tend to be conventions arising from within the scientific community, as in the case of arrows showing current vectors, which may not be intuitive to non-experts. More generally, there are conventions for creating and interpreting graphics that are learned as a person becomes visually literate in their culture (Tversky, 2001; Tversky, Kugelmas & Winter, 1991), such as reading flow diagrams from left to right across the page. However, these are mostly not specific enough to permit accurate interpretation of metaphors intended to convey information about particular phenomena. In most cases, this means that, to be safe, metaphors can be explicitly taught to students before entering learning environments where they are used.

On the whole, though, good metaphors selected to overcome the limits of direct sensory experience can work very well in artificial environments, as in Dede et al.'s (1997) environment that describes Newton's laws, and in the "Virtual Puget Sound" simulation of physical oceanography (Fruland, 2002; Winn et al., 2001; Winn & Windschitl, 2002). Most of the students working




FIGURE 2 Drawings showing vertical circulation of water in Puget Sound made by a 14-year-old, before and after visiting Virtual Puget Sound. (Note the vertical exaggeration in the diagram provided for the student to draw on.) The top two drawings show the student's initial conception of water movement on a rising and a falling tide. The bottom drawing, made after visiting Virtual Puget Sound, clearly shows the vertical circulation pattern.

with Virtual Puget Sound learn how water moves in the ocean, from interacting with metaphors showing the behavior of tides and currents. To illustrate, an important characteristic of Puget Sound is that more saline, denser water enters from the north. As it moves south, it is diluted with fresh water entering the Sound from rivers and, becoming less dense, rises. Water therefore leaves Puget Sound at the surface, creating a vertical circulation pattern in addition to the more familiar horizontal one. Figure 2 shows drawings made by a 14-year-old before and after visiting Virtual Puget Sound. It is

clear that the student, like most others, understood what the metaphors, showing this circulation pattern, meant, and that he learned how the circulation process works.

The third way in which cognition can be thought of as embodied is when we use our bodies to solve problems. Our physical behavior often externalizes our thinking and extends cognition beyond our brain. Counting on our fingers is an obvious example, as is using gestures as we tell stories (Kita, Danziger & Stolz, 2001). In an unpublished study, a dyslexic sixth-grader, who visited Virtual Puget Sound, found it difficult to explain in words the vertical circulation pattern he had just learned about. He described it perfectly using his hands and arms in a circular gesture. (Roth [2001] has reviewed research on how students use gestures as they learn.) Going further, Varela et al. (1991) make the case that all of cognition is “enactive.” They argue that the way we organize ideas directly reflects how we act in the world. From there, Varela et al. construct a view of cognition that is based, not on the idea that the mind is a mirror of the environment, but that cognition consists of the constant, reciprocal, interaction between the mind and the environment.

Bodily activity is often essential to understanding what is going on in an artificial environment. The ability to move about makes it easier to remember three-dimensional spatial layouts (Arthur, Hancock & Telke, 1996). Students who frequently change their positions and points of view learned more about how water moves in Virtual Puget Sound than students who do not (Winn & Windschitl, 2002). A most interesting case is Gabert’s (2001) study of high school Chemistry students learning about changes of the states of matter. Her immersive environment consisted of three intersecting surfaces—a three-dimensional graph of temperature, volume and pressure. At various places, the students could observe processes at the molecular level, for example where solids changed to liquids and liquids to gases. The scale was such that the graph was very large relative to the students, towering above them and reaching far into the distance. As a consequence, the students had to “fly” around the environment in order to visit these locations and to observe changes of state at the molecular level. The actual position of the student’s body in three-dimensional space was therefore meaningful. To observe what happened when temperature increased, the student had to fly up to a higher point in the environment, rising with the temperature. To observe the result of an increase in pressure, the student had to move to the right. It was as if the student’s body became a data point on the graph—a cursor moving around in three dimensions, marking points where important processes occurred.

Connecting brain and body to the environment

It is largely through physical action that cognitive activity is connected to the environment. Yet, while all activity occurs in a context, to say cognition is “embedded” in an environment is to say far more than that it is contextualized. In what follows next, a distinction is made between the environment as a whole and the environment that the student knows about. The latter exists for the student either from direct perception (limited, as we have seen, by the constraints of the human sensory apparatus), or through the acquisition of knowledge about it in ways other than direct experience. We then look at the means by which students become embedded in an environment. These involve either natural proclivities, or deliberate affective and cognitive strategies that are built into the environment.

Environment and “Umwelt”

A central premise of the constructivist position is that all knowledge is constructed by the student and that every student’s understanding of the environment is idiosyncratic. It follows, the argument goes, that there can be no objective, fixed standard against which to assess what a person knows. The premise is not problematic. The conclusion is. It leans dangerously towards solipsism, as Maturana and Varela (1987) point out. Of course, no one’s knowledge of the world can be complete, and therefore everyone knows the world in a somewhat different way. But these differences in knowledge arise because everyone has a different set of experiences, not because there is no objective reality.

Following Roth (1999) and others, the word “Umwelt” is used to refer to the environment as seen and understood, idiosyncratically, by different individuals. The first use of the word “Umwelt” (German for “environment”) in this way is attributed to the biologist Von Uexküll (1934), who described what the world might look like if you were a scallop, or a bee, or some other creature. His somewhat whimsical drawings, reproduced in Clark (1997, pp. 26-28), show that each creature’s Umwelt is quite different from every other creature’s. A student’s Umwelt is the environment as the student sees and knows it—a limited view of the real world, ever changing as the student explores it and comes to understand it.

There are three further points to make. First, the uniqueness and variability of Umwelt are not the result of limited sensory capacity, which we saw above is a physiological constraint. Rather, they arise from differences in each individual’s experience of the environment. Put another way, the embodiment of sensory perception varies across species; normal variability

in Umwelten is a difference among individuals.

Second, differences among Umwelten do not just reside in knowledge and internal representations. Individuals *see* their Umwelt differently. The proverbial Inuit can actually see differences in types of snow. Goldstone et al. (2000) cite evidence for this kind of heightened perceptual discrimination in expert radiologists, beer tasters and chick sexers. There is also evidence that the exposure to environmental stimuli that leads to heightened sensory sensitivity brings about measurable changes in the auditory (Weinberger, 1993) and visual (Logothetis, Pauls & Poggio, 1995) cortex.

Third, an argument against prescriptive approaches to learning and teaching has been that we cannot predict the behavior of natural environments, because they are too complex. Nor can we predict the behavior of a student's Umwelt, because it is too idiosyncratic. In contrast, every function and feature of an artificial environment is known. An artificial environment is completely predictable, because we have made it. How it will respond to every kind of input that it is programmed to accept is predictable. Also, we can control precisely what a learning environment reveals about itself, as when we turn off gravity to illustrate aspects of Newton's laws (Dede et al., 1997), or turn on ocean currents to show water movement in Virtual Puget Sound. This makes it easier to predict, though not infallibly so, how students will adapt to artificial environments as they struggle to understand the concepts and principles they embody.

When we say cognitive activity is embedded, we therefore mean that the student interacts with an Umwelt, which changes with experience, and that the Umwelt provides an idiosyncratic and incomplete view of an environment, which may or may reveal or conceal information about itself. Embeddedness therefore depends on the nature of the interaction of the student with the Umwelt and how well the Umwelt reflects properties of the environment. We turn to these matters next.

Coupling students to the environment

Some recent thinking suggests that it is better to consider students to be tightly *coupled* to the environment rather than *embedded in* it. Being embedded suggests the student is passive, carried along as the environment changes. Successful students are anything but passive. The idea of "coupling" (Maturana & Varela, 1987; Reyes & Zarama, 1998; Roth, 1999) describes mutually influencing dynamic interaction between learners and the environment. Clark's "continuous reciprocal causation" (1997, p. 163) likewise describes an unbroken process where the actions of learners and of

the environment, and the consequences of these actions, are a two-way street. Clark (1997, p. 171-2) gives as an example trying to catch a hamster with a pair of tongs. Each move the animal makes instantly determines how you wield the tongs. And how you wield the tongs instantly determines how the animal will behave. Cause and effect are not clearly distinguishable. The two systems are tightly coupled. Winn and Windschitl (2001) have argued that when interactions like these occur, students can learn in artificial environments in the same ways that they learn in the natural world – intuitively, constructively and actively.

Students can become tightly coupled with artificial environments as a result of affective and cognitive factors. Affect is exploited by creators of computer games to keep students engaged with their products for extended periods. Cognitive strategies for engaging students in learning environments are aimed at bringing about conceptual change, usually through problem-solving or discovery learning. We examine each of these in turn.

Affective strategies

The experience of being coupled to an artificial environment is called “presence” (Zeltzer, 1991). Presence is the belief that you are “in” the artificial environment, not in the laboratory or classroom interacting with a computer. Typically, during a visit to an artificial environment, attention is divided between the environment created at the computer interface, be it a computer screen or virtual reality helmet, and the environment outside, which might be noisy, or contain someone giving you instructions about what to do, or be distracting in other ways.

Presence varies with the extent to which attention is divided between the artificial and the real environment (Witmer & Singer, 1998). A high level of presence requires complete attention to the artificial environment, produces an almost complete immunity from distraction and allows total engagement with the artificial environment. Witmer and Singer also say that presence can be improved by a high level of enjoyment and by being immersed in the environment. This last condition is met best when a student wears a virtual reality helmet with a wide field of view and when head movements are tracked in real time, allowing the student to look around in the artificial environment in the same way as in the natural world. Lin et al. (2002) have shown that increasing the field of view during immersion in an artificial environment improves presence and enjoyment—but also simulator sickness. Presence can heighten both pleasant and unpleasant experiences.

Hedden’s (1998) study of why people can spend uninterrupted hours

totally engrossed in a computer game offers some insights into affective coupling with artificial environments, by heightening presence through the direction of attention. Hedden draws on Lepper and Malone's theory of motivation (Malone, 1980; Malone & Lepper, 1987; Lepper & Chabay, 1985) and proposes that challenge, curiosity and fantasy act to focus attention on games, thereby increasing presence. Challenge is greatest when the goal is clear, initial uncertainty is high, and the activities necessary to attain the goal are of intermediate difficulty. Curiosity is aroused when "environments are neither too complex nor too simple with respect to the person's existing knowledge" (Malone, 1980, cited in Hedden, 1998, p. 36). Environments therefore hold attention if they make students curious to complete understanding they have only partially acquired. Fantasy arises when the student can imagine a number of possible outcomes from the activity. As the activity progresses, the possibilities are eliminated one by one until just one remains—the solution to the problem.

When the conditions required to optimize challenge, curiosity and fantasy are met in an environment, Hedden (1998) proposes that students enter what Csikszentmihalyi (1988, 1990) has called a state of "flow." Flow is characterized by a high level of enjoyment, total engagement in the task, a loss of awareness of the passage of time, a resistance to distraction, and a pleasant sense of fatigue when finished. Obviously, these conditions correspond to characteristics of presence identified earlier—engagement, immersion and enjoyment.

Building artificial environments that draw attention exclusively to themselves, and even induce flow, can have extremely powerful effects. Hoffman et al. (2000) placed children into an immersive game while they were undergoing wound care for severe burns. The children were completely drawn into the artificial environment, reporting high levels of presence. They also paid so little attention to what was happening to them in the real world that they experienced remarkably less pain. Carlin, Hoffman and Weghorst (1997) described an effective treatment for arachnophobia using an artificial environment. The patient's presence in the environment was so high, and the virtual spider so believable, that a variation of aversion therapy was completely successful.

Work with students of all ages has shown that presence consistently predicts the amount students learn and that reduced presence, caused by distraction or discomfort, impedes learning (Winn et al, 2001; Winn et al., 2002). These studies have also consistently shown positive correlations between presence and enjoyment. Hedden's predictions seem therefore to

apply to immersive artificial environments as well as to games. This is born out by observation of students in these studies, most often younger ones, who become completely engaged in an artificial environment. They become isolated from all external distraction. They perform physical actions that are appropriate to the artificial rather than the real world, to the extent that care has to be taken that they do not hurt themselves. They are bewildered when objects in the artificial environment do not act as they would in the real world—when they do not fall when dropped, for example.

Cognitive strategies

The cognitive strategies that couple students to artificial environments are best illustrated from studies of conceptual change. Students bring a whole host of assumptions about how the world works to their activities in artificial environments. Students visiting Virtual Puget Sound, for example, have started out believing that there are no tides in Puget Sound because there are none in Lake Washington—a completely landlocked body of water; that the water in Puget Sound is saltier than the open ocean, because evaporation in a relatively small body of water will concentrate the salt more quickly; that objects released to float in the water will eventually return to where they started from. These are common misconceptions. The purpose of artificial environments is to persuade the students to reject such misconceptions and accept scientifically accurate conceptions in their place.

Long-held misconceptions are notoriously difficult to change (Chinn & Brewer, 1993). An effective strategy for doing so, which draws on a long empirical history (Posner et al., 1982), is proposed by Windschitl and André (1998). The basic idea is to have students find and confront compelling evidence that can neither be predicted from nor explained by their current conceptions. To be accepted in place of their current conceptions, the evidence for the new conception must meet a number of criteria. The new evidence must be understandable, otherwise it may simply be learned by rote. It must fit within the student's epistemological stance towards the phenomenon in question—as well as being understood, it must be believed. It must be fruitful and allow the student to solve this and other problems. It must be acquired through interaction with the environment—the student must actively discover it, test it and apply it without too much direction from someone else.

It is clear that these strategies for encouraging conceptual change are somewhat similar to those used by game designers. The scenarios within which students visiting Virtual Puget Sound solve problems are designed to be optimally challenging and to arouse curiosity—though not to encourage

fantasy. In one, students are told that a non-native species of fish has been caught in Puget Sound. Should it become established, it will destroy the salmon fishery. Their job is to find its lair, given that it prefers certain levels of salinity and certain current speed. In another, students have to recommend a site for a new sewage treatment plant that will discharge treated sewage into the Sound. They must determine where to put the end of the discharge pipe, select its location and its depth, so that the sewage will be quickly dispersed and will leave the Sound as rapidly as possible.

Studies using Virtual Puget Sound have shown that the extent of conceptual change is greater when engagement of the student in the environment is high. Reported presence predicts conceptual change, measured by gain scores (Winn et al., 2001), and when posttest scores are regressed onto presence scores (Winn et al., 2002). This second study also showed that students who were immersed in the environment, using virtual reality technology, reported more presence than those who worked with an equivalent environment on the desktop, and that immersed students also learned more than non-immersed students about water movement, a dynamical three dimensional process. These findings suggest that immersion and presence, which couple students more tightly to the learning environment, support cognitive processes that lead to conceptual change.

LEARNING AS DYNAMIC ADAPTATION

The changes that occur to a student, and to the environment to which the student is closely coupled through reciprocal interaction, can best be explained by mechanisms of adaptation. Adaptation of a species to an environment, over tens of thousands of years, is explained by theories of evolution. Adaptation of an individual to an environment, over a lifetime, is accounted for by theories of physical, cognitive and social development. Adaptation of an individual to an environment in the short term—in a course, a lesson or a visit to an artificial environment—can be explained, in part, by theories of conceptual change that we looked at earlier, and, more completely, by accounts that allow us to capture something of the complexity of student-environment interactions. In this section, we look more closely at how this last kind of adaptation to learning environments might occur. We examine both the nature of the changes that adaptation leads to and possible mechanisms for those changes.

What changes with adaptation?

Earlier, we saw that exposure to an environment can lead to physical changes in the brain, resulting in heightened perceptual sensitivity, which leads a person to actually see things differently in the environment. The Inuit really sees different types of snow. The professional beer taster really tastes differences between kinds of beer. It seems reasonable, then, that a basic change that occurs with adaptation is a greater clarity with which students make distinctions among objects and phenomena in the environment. This is not simply to say that a student distinguishes among more things, although this certainly occurs. Rather, the student makes distinctions with more certainty.

Reyes & Zarama (1998) put forward an empirically-supported theory of learning, based on system-theoretic principles, in which they describe learning as a process for “embodying distinctions.” The theory describes four steps. 1) “Declaring a break.” An unexpected perturbation in the environment disrupts the flow of reciprocal action that couples the student and the environment. The student might notice something untoward or new that cannot be accounted for by current understanding, which interrupts the activity. Observing visitors to Virtual Puget Sound has provided many examples of this, as, for instance, when a student notices for the first time a current below the surface moving in the opposite direction to the surface current. 2) “Drawing a distinction.” The student distinguishes the new phenomenon from those that are familiar. Rather than just currents, the student now understands that there are currents that move in one direction and others, at different depths, that move in other directions. The concept “current” has been divided into distinct categories. 3) “Grounding the distinction.” The new distinction must be compatible with other concepts and principles that the student knows. The student must believe it to be reasonable that currents can move in different directions at different depths. If, for some compelling reason, the student has a deeply-rooted belief that water moves homogeneously, regardless of its depth, then the new distinction will simply be memorized (for the test, maybe), but not believed or understood. 4) “Embodying the distinction.” For learning to occur, the new distinction must be used fruitfully. Thus, if the student can use the distinction between currents moving in one direction and those moving in another direction to solve a problem, then learning is more likely to take place. Winn and Windschitl (2002) have shown that students who discover and apply new distinctions while visiting Virtual Puget Sound (here called “principles” rather than “distinctions”) will be more successful at solving the problem of where to find

the lair of a predatory fish.

The explanation of learning that Reyes and Zarama propose obviously contains features found in other learning theories. Making distinctions is not unlike Gagné et al's (1988) long-standing "discrimination learning", a process of learning to tell one thing from another that is a prerequisite for learning concepts and for solving problems. Nor is their four-step process that unlike Posner et al's (1982) or Windschitl and André's (1998) accounts of conceptual change, which require that a new conception be understandable, compatible with currently-held beliefs, and fruitful in application. However, their particular view of learning has two advantages. First, it allows, indeed encourages, consideration of how an entire student-environment system co-adapts, rather than focusing on a narrow cluster of phenomena and concepts isolated in the student's brain. In proposing distinction-drawing as the basic mechanism for learning, the scope of what influences learning is not constrained. Students may make distinctions between a new conception and any conception they already hold. In practice, the distinctions are usually among ideas that are conceptually close. But there is nothing to stop a student distinguishing among the new concept and any ideas in the cognitive system, which can explain the idiosyncratic interpretations of environments and solutions to problems that are frequently observed in students visiting artificial environments. Often, indeed, the distinctions we make, "tell us more about ourselves than about the world we are describing," (Reyes & Zarama, 1998, p. 23).

Second, Reyes' and Zarama's view of learning parallels neurological accounts. If thoughts are represented in the brain as networks of neurons, defined principally by the pattern and strength of connections among them (Markowitsch, 2000), then the brain is wired to make distinctions. The differences among neural networks can be computed from the differences among the strengths of individual connections. As the strengths change, the differences become more or less marked and the clarity of distinctions evolves.

Mechanisms for adaptation

To view learning as adaptation to an environment is to bring to it the biological flavor of evolutionary and developmental theories. From a biological perspective, adaptation results from the interaction of two kinds of influence: Environmental pressure, and genetic predisposition to change. Recent research has shown that people are not born genetically "wired" so that their development runs its course like a computer program. Rather, their genetic

make-up predisposes them to develop in particular ways should contact with the environment trigger particular genetic programs (Neville & Bavelier, 2000; Johnson, 2000). So far, we have looked mostly at the kinds of interactions with the environment that can bring about these changes. We now look to conditions within the learner (the “genetic” part) for mechanisms that drive adaptation.

We have seen that tight coupling between a student and a learning environment leads to changes in both the student and the environment. Adaptation is mutual. Since our focus is on computer-created learning environments, we need to find mechanisms that apply equally to biological and machine adaptation. Such mechanisms are proposed by Holland (1992, 1995).

Holland has described “genetic algorithms” for adaptation to environments by humans and by machines. Genetic algorithms are prescriptions for procedures that execute when a certain event ensues from interaction with the environment. As such, they resemble “if-then rules,” as rules are defined in other theories of learning (see, for example, Scandura, 1983). Yet genetic algorithms have two properties that distinguish them from rules traditionally construed. First, they are more like competing hypotheses than rules, meaning that they thrive or fail as a result of their success in guiding fruitful interactions with the environment. Second, their success or failure, like that of living organisms, depends on their own ability to adapt. So not only does a student’s knowledge change through adaptation to the environment. The rules, or procedures, that specify how the student interacts with the environment in the first place also change through adaptation, based on their success at producing fruitful behavior.

Holland (1995, 53-80) describes a number of processes by means of which these rules are created, get abandoned and change over time. “Credit assignment” changes the perceived value of rules depending on the consistency of their success at predicting how the environment behaves. Rules that predict consistently well survive. Those that do not die out. Winn and Windschitl (2002) we have documented cases of visitors to Virtual Puget Sound explicitly rejecting rules that fail to predict observations consistently. One student started out believing that water in Puget Sound was more saline than that in the open ocean, because evaporation in a relatively small body of water would concentrate the salt more. (In fact, water in Puget Sound is less saline because more fresh water flows into it from rivers than is lost through evaporation.) When application of this rule failed to solve a problem that required prediction of salinity, the student stated, “I’ll have to

change my theory here,” discovered, by experimenting, the correct principle, and proceeded to solve the problem.

“Rule discovery” creates new rules *ex nihilo*. More correctly, the rules are new to the student and govern only the Umwelt. They are already programmed into the simulation that drives the environment, remaining for the student to discover. Many students do not know that the water in Puget Sound gets saltier as it gets deeper. Yet this rule cannot be ignored if you use the virtual salinity meter at different depths—something the student must do in order to solve a number of problems. Typically (Fruiland, 2002; Winn et al., 2001), students comment, often with surprise, on this discovery and use it from then on to help solve problems.

Finally, “crossover,” in genetics, is a process through which pieces of two sets of genetic material are combined to produce offspring with new properties. In Holland’s view of adaptation, crossover likewise combines parts of two rules to produce a new one. The pieces to combine may be chosen randomly, as in the BEAGLE expert system (Forsyth, 1984, pp. 162-165), in which case the new rule thrives or dies depending solely on how well it predicts events in the environment, through credit assignment, described above. Or the pieces may be chosen in some principled way (Holland, 1995, pp. 65-69), increasing the new rule’s chances of survival. In either case, the rules evolve. One student working in Virtual Puget Sound (Winn & Windschitl, 2002) knew that tides are cyclical (one rule) and that the speed of currents varies over the tidal cycle (another rule). However, she did not know when in the tidal cycle the currents were the fastest. From her think-aloud protocol, it was determined that she reasoned the current would be slowest at high and low tide, because at those points the current would slow down and change direction. Therefore, she reasoned, current would be fastest between high and low tide. She then measured current speed midway between high and low tide, and then repeated the measurement one hour earlier and one hour later in the cycle. The second two measurements showed slower current. She then combined the two rules into one, and from then on proceeded to look for the fastest water only at times half way between high and low tide. In this case, the combination of rules was principled—the student reasoned through the alternatives to arrive at a conclusion. In other cases, students stumbled onto the same principle by chance, in a sense randomly combining the two rules, yet nonetheless proceeding to use the new rule effectively.

Holland’s mechanisms for adaptation are not the only ones that can explain learning by interacting with an environment. However, they suit our

purpose. They apply equally well to adaptation by people and by machines, a necessary feature for our work with artificial environments. Also, genetic algorithms allow adaptation to be described quantitatively. This permits data to be described in ways that can be submitted to numerical analysis, can be tested in simulations, and can be displayed, for interpretation, in a number of useful ways. In short, they meet several of the criteria for the scientific rigor that needs to exist in educational research.

SUMMARY AND CONCLUSION

Putting all of this together, we arrive at a description of learning that is quite different from accounts given by traditional cognitive psychology and by constructivism. The new account is grounded on a framework that integrates three concepts, embodiment, embeddedness and adaptation. The framework brings together recent research and theory that extend the purview of cognitive activity from the brain, through the body, to the environment itself. Learning is considered to arise from the reciprocal interaction between external, embodied, activity and internal, cerebral, activity, the whole being embedded in the environment in which it occurs. Learning is no longer confined to what goes on in the brain. Indeed, we may ask, like Clark (1997, p. 213), "Where does the mind stop and the rest of the world begin?" One answer to this question is to assume, like Beer (1995), that sometimes the coupling between a person and the environment is so tight that it is more convenient to think of person and environment as one evolving system rather than two interacting ones. In this view, learning can be thought of as self-organization by the system, and new knowledge as an emergent property of that self-organizing activity. In any event, embodiment, embeddedness and adaptation can no longer be considered to be independent of each other.

The framework has the added advantage of accommodating some recent research in cognitive neuroscience. While we are still a long way from being able to explain cognition in terms of brain structures and processes, it is nonetheless important to note the convergence of evidence from the neurosciences and the learning sciences toward explanations of how some learning takes place. We have also noted commonalities among system-theoretic and neuroscientific views of cognition. For example, both include the possibility that complexity can arise from the application of relatively simple rules, that cognition can be explained in terms of the collaborative activity of simple units, as when, for example, a series of new distinctions can lead

a student to understand whole environments, such as Virtual Puget Sound.

The framework also allows us to rethink some of the details that have been controversial in recent educational literature. For example, mental representation and computation are restored to central roles in learning and cognition. However, their natures are seen somewhat differently than before. Representation is no longer considered to consist of mapping objects and phenomena in the environment onto mental symbols. Rather, representation becomes a dynamic, environmentally-triggered activity, in which mental structures are activated by external or internal events, recreating some of the experience that was felt when the event was actively perceived for the first time. Computation is now viewed in terms of the brain's ability, first, to draw distinctions, and then to extend these distinctions in ways that permit higher cognitive operations.

The coupling of student and environment can be tightened affectively by increasing the sense of presence the student feels in the environment. Presence emerges naturally, in most cases, if the student is immersed in the environment using virtual reality technology. However, research on motivation to play computer games suggests many other strategies expressly designed to keep students' attention directed towards the learning environment and away from the world outside. Also, cognitive strategies for coupling student and environment require students to be actively engaged in changing old conceptions into new ones as their current conceptions are challenged by unexpected events in the environment.

The environment should be thought of in two ways. The *Umwelt* is the environment that the student knows, constrained by sensory limitations. It is changed both through direct experience and through interaction with metaphorical representations of phenomena not accessible to the senses, which lead to new distinctions among concepts and principles. Beyond that lies an environment that has properties, separate from knowledge constructed about it, that obey the laws of nature in predictable, objectively verifiable ways. This means that adaptation to the environment can lead to conceptions that can be considered right or wrong. In the case of artificial environments, whose behaviors are programmed by us and are therefore completely known, this also means that a considerable measure of control can be placed on how the student-environment interaction proceeds. Beyond scaffolding (Linn, 1995), we can now embed pedagogical strategies into the very fabric of the environment. Since learning arises from adaptation to the environment, it can be guided by the behavior of the environment itself.

How adaptation occurs can be considered from several perspectives,

from traditional explanations of conceptual change, to descriptions based on changes to connections among neurons. Holland's account of adaptation through the action of genetic algorithms is useful, in that it applies both to people and to machine-constructed agents and environments, and in that it can account for the creation and evolution of rules we observe students discovering and using.

Finally, this system-theoretic view of learning opens the door to quantitative descriptions of adaptation and to building mathematical models of learning that can be simulated and tested. Connectionist models (Rumelhart & McClelland, 1986) and dynamical system theory (Abraham & Shaw, 1992; Port & Van Gelder, 1995) offer techniques for describing and simulating adaptation to learning environments in all their complexity and, dynamically, in real time.

In conclusion, it is important for educational researchers to take these views of learning and cognition seriously, even though they are, as yet, not fully formed. We have spent too many years away from studying learning directly. We have traded rigorous, quantitative science for methodologies, poorly adapted from anthropology and sociology (and even from literary criticism), which are appropriate only for the study of things that are peripherally relevant to learning. Our educational techniques and technologies are developing at an ever-increasing pace. Our research and its application must keep up.

REFERENCES

- Abraham, R. and Shaw, C. (1992). *Dynamics: The geometry of behavior*. New York: Addison-Wesley.
- Arthur, E., Hancock, P. and Telke, S. (1996). Navigation in virtual environments. *Proceedings of the SPIE – International Society for Optical Engineering*, 2704, 77-85.
- Bednar, A., Cunningham, D., Duffy, T. and Perry, J.D. (1992). Theory into practice: How do we link? In T. Duffy and D. Jonassen (Eds.), *Constructivism and the technology of instruction: A conversation*. Hillsdale, NJ: Erlbaum.
- Beer, R.D. (1995). Computation and dynamical languages for autonomous agents. In R.F. Port and T. Van Gelder, *Mind as motion: Explorations in the dynamics of cognition*. Cambridge, MA: MIT Press.
- Berninger, V. and Richards, T. (2002). *Brain literacy for psychologists and educators*. New York: Academic Press.
- Bickhard, M.M. (2000). Dynamic representing and representational dynamics. In E. Dietrich and A.B. Markman (Eds.), *Cognitive dynamics: Conceptual and representational change in humans and machines*. Mahwah NJ: Erlbaum.
- Boden, M. (1988). *Computer models of mind*. New York: Cambridge University Press.

- Brown, J.S., Collins, A. and Duguid, P. (1989). Situated cognition and the culture of learning. *Educational Researcher*, 18 (1), 32-43.
- Carlin, A.S., Hoffman, H.G. and Weghorst, S., (1997). Virtual reality and tactile augmentation in the treatment of spider phobia: A case study. *Behaviour Research and Therapy*, 35, 153-158.
- Chinn, C.A. and Brewer, W.F. (1993). The role of anomalous data in knowledge acquisition: A theoretical framework and implications for science instruction. *Review of Educational Research*, 63, 1-49.
- Clark, A. (1997). *Being there: Putting brain, body and world together again*. Cambridge, MA: MIT Press.
- Csikszentmihalyi, M. (1988). The flow experience and human psychology, In M. Csikszentmihalyi and I.S. Csikszentmihalyi (Eds.), *Optimal experience: Psychological studies of flow in consciousness*. New York: Cambridge University Press.
- Csikszentmihalyi, M. (1990). *Flow: The psychology of optimal performance*. New York: Harper and Row.
- Cunningham, D. (1992). Assessing constructions and constructing assessments. In T Duffy and D Jonassen (Eds.), *Constructivism and the technology of instruction: A conversation*. Hillsdale, NJ: Erlbaum.
- Dede, C., Salzman, M., Loftin, R. B. and Ash, K. (1997). Using virtual reality technology to convey abstract scientific concepts. In M. J. Jacobson and R. B. Kozma (Eds.) *Learning the sciences of the 21st century: research, design and implementing advanced technology learning environments*. Mahwah, NJ: Erlbaum.
- Dietrich, E. and Markman, A.B. (2000). Cognitive dynamics: Computation and representation regained. In E. Dietrich and A.B. Markman (Eds.), *Cognitive dynamics: Conceptual and representational change in humans and machines*. Mahwah NJ: Erlbaum.
- Dreyfus, (1979). *What computers can't do*. New York: Harper and Row.
- Dreyfus, H.L. and Dreyfus, S.E. (1986). *Mind over machine*. New York: The Free Press.
- Duffy, T. and Jonassen, D. (1992). Constructivism: New implications for instructional technology. In T. Duffy and D. Jonassen (Eds.), *Constructivism and the technology of instruction: A conversation*. Hillsdale, NJ: Erlbaum.
- Farrah, M. (2000). The neural bases of mental imagery. In M. Gazzaniga (Ed.), *The new cognitive neurosciences, second edition*. Cambridge MA: MIT Press.
- Forsyth, R. (1984). Machine learning strategies. In R. Forsyth (Ed.), *Expert systems: Principles and case studies*. London: Chapman and Hall.
- Fruland, R.M. (2002). *Using immersive scientific visualizations for science inquiry: Co-construction of knowledge by middle and high school students*. Paper presented at the annual meeting of the American Educational Research Association, New Orleans, LA.
- Gabert, S.L. (2001). *Phase world of water: A case study of a virtual reality world developed to investigate the relative efficiency and efficacy of a bird's eye view exploration and a head-up-display exploration*. Ph.D. Dissertation, College of Education, University of Washington, Seattle, WA.
- Gagné, R.M., Briggs, L.J. and Wager, W.W. (1988). *Principles of instructional design: Third edition*. New York: Holt Rinehart and Winston.
- Goldstone, R.L., Steyvers, M., Spencer-Smith, J. and Kersten, A. (2000). Interactions between perceptual and conceptual learning. In E. Dietrich and A.B. Markman (Eds.), *Cognitive dynamics: Conceptual and representational change in humans and machines*. Mahwah NJ: Erlbaum.
- Haugeland, J. (1991). Representational genera. In W. Ramsey (Ed.), *Philosophy and connectionist theory*. Hillsdale NJ: Erlbaum.

- Hedden, C. (1998). *A guided exploration model of problem-solving discovery learning*. Ph.D. Dissertation, College of education, University of Washington, Seattle, WA.
- Hoffman H.G., Doctor J.N., Patterson D.R., Carragher G.J. and Furness T.A., (2000). Use of virtual reality for adjunctive treatment of adolescent burn pain during wound care: A case report. *Pain*, 85, 305-309.
- Holland, J. (1992). *Adaptation in natural and artificial environments: An introductory analysis with applications in biology, control and artificial intelligence*, Second edition. Cambridge, MA: MIT Press.
- Holland, J. (1995). *Hidden order: How adaptation builds complexity*. Cambridge, MA: Perseus Books.
- Jackson, R. (2000). *Collaboration and learning within tele-immersive virtual environments*. Ph.D. Dissertation, College of Education, University of Washington, Seattle, WA.
- Johnson, M.H. (2000). Cortical mechanisms of cognitive development. In M. Gazzaniga (Ed.), *Cognitive neuroscience: A reader*. Malden MA: Blackwell.
- Kita, S., Danziger, E. and Stolz, C. (2001). Cultural specificity of spatial schemas as manifested in spontaneous gestures. In M. Gattis (Ed.), *Spatial schemas and abstract thought*. Cambridge MA: MIT Press.
- Kulikowich, J.M. and Young, M.F. (2001). Locating an ecological psychology method of situated action. *Journal of the Learning Sciences*, 10(1-2), 165-202.
- Lave, J. and Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge: Cambridge University Press.
- Lepper, M.R. and Chabay, R.W. (1985). Intrinsic motivation and instruction: Conflicting views on the role of motivational processes in computer-based education. *Educational Psychologist*, 20, (4), 417-430.
- Lin, J., Duh, H.B.L., Parker, D.E., Abi-Rached, H. and Furness, T.A. (2002, March). Effects of Field of View on Presence, Enjoyment, Memory, and Simulator Sickness in a Virtual Environment. *Proceedings of the meeting IEEE VR*, Orlando, FL.
- Linn, M. (1995). Designing computer learning environments for engineering and computer science: The scaffolded knowledge integration framework. *Journal of Science Education and Technology*, 4 (2), 103-126.
- Logothetis, N.K., Pauls, J. and Poggio, T. (1995). Shape representation in the inferior temporal cortex of monkeys. *Current Biology*, 5, 552-563.
- Malone, T.W. (1980). *What makes things fun to learn?* Ph.D. Dissertation, Stanford University, Palo Alto, CA.
- Malone, T.W. and Lepper, M.R. (1987). Making learning fun: A taxonomy of intrinsic motivations for learning. In R.E. Snow and M.J. Farr (Eds.), *Aptitude, Learning and Instruction Volume 3: Conative and Affective process analysis*. Englewood Cliffs, NJ: Erlbaum.
- Markowitsch, H.J. (2000). The anatomical bases of memory. In M. Gazzaniga (Ed.), *The new Cognitive Neurosciences, Second edition*. Cambridge, MA: MIT Press.
- Marr, D. (1982). *Vision*. New York: Freeman.
- Maturana, H. and Varela, F. (1987). *The tree of knowledge*. Boston: New Science Library.
- Nagel, T., (1974). What it is like to be a bat. *Philosophical Review*, 83, 435-450.
- Neville, H.J. and Bavelier, D. (2000). Specificity and plasticity in neurocognitive development in humans. In M. Gazzaniga (Ed.), *The new Cognitive Neurosciences, Second edition*. Cambridge, MA: MIT Press.
- Pinker, S. (1997). *How the mind works*. New York: Norton.
- Port, R.F. and Van Gelder, T. (1995). It's about time. In R.F. Port and T. Van Gelder, *Mind as*

- motion: Explorations in the dynamics of cognition.* Cambridge, MA: MIT Press.
- Posner, G.J., Strike, K.A., Hewson, P.W. and Gertzog, W.A.. (1982). Accommodation of scientific conception: Toward a theory of conceptual change. *Science Education*, 66, 211-227.
- Reyes, A. and Zarama, R. (1998). The process of embodying distinctions: A reconstruction of the process of learning. *Cybernetics and Human Knowing*, 5 (3), 19-33.
- Rosch, E. (1999). Reclaiming concepts. *Journal of consciousness studies*, 6 (11), 61-77.
- Roth, W.M. (1999). The evolution of Umwelt and communication. *Cybernetics and Human Knowing*, 6(4), 5-23.
- Roth, W.M. (2001). Gestures: Their role in teaching and learning. *Review of Educational Research*, 71, 365-392.
- Rumelhart, D.E., McClelland, J.L. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition.* Cambridge MA: MIT Press.
- Scandura, J.M. (1983). Instructional strategies based on the structural learning theory. In C.M. Reigeluth (Ed.), *Instructional design theories and models.* Hillsdale, NJ: Lawrence Erlbaum Associates.
- Skarda, C. and Freeman, W. (1987). How brains make chaos in order to make sense of the world. *Behavioral and Brain Sciences*, 10, 161-195.
- Steffe, L. (1999). Individual constructive activity: An experimental analysis. *Cybernetics and Human Knowing*, 6 (1), 17-31.
- Streibel, M.J. (1991). Instructional plans and situated learning: The challenge of Suchman's theory of situated action for instructional designers and instructional systems. In G.J. Anglin (Ed.), *Instructional technology past, present and future.* Englewood, CO: Libraries Unlimited.
- Thelen, E. and Smith, L. (1994). *A dynamic systems approach to the development of cognition and action.* Cambridge MA: MIT Press.
- Tversky, B. (2001). Spatial schemas in depictions. In M. Gattis (Ed.), *Spatial schemas and abstract thought.* Cambridge MA: MIT Press.
- Tversky, B., Kugelmass, S. and Winter, A. (1991). Cross-cultural and developmental trends in graphic productions. *Cognitive Psychology*, 3, 515-557.
- Varela, F.J., Thompson, E. and Rosch, E. (1991). *The embodied mind.* Cambridge, MA: MIT Press.
- Von Bertalanffy, L. (1968). *General System Theory.* New York: Braziller.
- Von Uexküll, J. (1934). A stroll through the worlds of animals and men. In K. Lashley (Ed.), *Instinctive behavior.* New York: International Universities Press.
- Weinberger, N.M. (1993). Learning-induced changes of auditory receptive fields. *Current opinion in neurobiology*, 3, 570-577.
- Windschitl, M. and André, T. (1998). Using computer simulations to enhance conceptual change: The roles of constructivist instruction and student epistemological beliefs. *Journal of Research in Science Teaching*, 35 (2), 145-160.
- Windschitl, M. and Winn, W. (2000). Learning from virtual worlds. In B. Fishman and S. O'Connor-Divelbliss (Eds.), *Proceedings of the International Conference of the Learning Sciences.* Mahwah, NJ: Erlbaum.
- Winn, W.D. and Windschitl, M. (2000). Learning science in virtual environments: The interplay of theory and experience. *Themes in Education*, 1, 373-389.
- Winn, W.D. and Windschitl, M. (2001). Learning in artificial environments. *Cybernetics and Human Knowing*, 8 (3), 5-23.
- Winn, W.D. and Windschitl, M. (2002, April). *Strategies used by university students to learn*

- aspects of physical oceanography in a virtual environment.* Paper presented at the annual meeting of the American Educational Research Association, New Orleans, LA.
- Winn, W.D., Windschitl, M., Fruland, R., Hedley, N. and Postner, L. (2001, April). *Learning science in an immersive virtual environment.* Paper presented at the annual meeting of the American Educational Research Association, Seattle, WA.
- Winn, W.D., Windschitl, M., Fruland, R. and Lee, Y. (2002). When does immersion in a virtual environment help students construct understanding? In P. Bell and R. Stevens (Eds.), *Proceedings of the International Conference of the Learning Societies, ICLS 2002.* Mahwah, NJ: Erlbaum.
- Winograd, T. and Flores, F. (1986). *Understanding computers and cognition: A new foundation for design.* New York: Addison Wesley.
- Witmer, B.G. and Singer, M.J. (1998). Measuring presence in virtual environments: A presence questionnaire, *Presence: Teleoperators and Virtual Environments*, 7, 225-240.
- Zeltzer, D. (1991). Autonomy, Interaction and Presence. *Presence: Teleoperators and Virtual Environments*, 1, 127-132.