

open play

I wish to play
in the domain
of your freedom —
your dance with me,
consensual and tender,
must not constrain you.

I wish to be
a desired delight,
a pleasant surprise
of joy and light.

Towards an Explanatory Framework for Learning in Artificial Environments¹

William Winn² and Mark Windschitl³

Abstract: Learning takes place in a variety of environments, ranging from the natural world to schools. Computer-generated artificial environments can emulate any of these environments and, at the same time, offer the student additional ways to learn. These derive from the ability of artificial environments to completely immerse the student in the environment, to simulate the natural environment, to make properties of the natural world, not normally accessible to the senses, available in a manner that affords interaction, and to build pedagogical strategies seamlessly into the environment itself.

Learning in artificial environments, like learning in natural environments, is most successful when students construct knowledge for themselves as they interact with the environment and observe the consequences of their actions. Research in our laboratory and elsewhere is beginning to shed light on how working in artificial environments can help students learn through a process not unlike the dynamic adaptation of any living organism to its environment. This approach suggests that an adaptive, second-order cybernetic view of learning in artificial environments might be a useful theoretical lens through which to view theory and interpret data. In this article, we suggest an explanatory framework that embodies these ideas.

Introduction

Learning occurs when people adapt to their environment so that a relatively permanent change takes place in knowledge or behavior. To understand how people learn, it is therefore necessary to know the nature and function of the environments in which learning takes place, to know how adaptation occurs and to know how it changes knowledge and behavior. For educators to guide learning, whether by helping students construct knowledge for themselves or through direct instruction, it is necessary to design and build environments that promote interaction and adaptation. Because they are designed, these environments are, to a greater or lesser extent, "artificial" (Simon, 1981). And it is in the design, building and use of artificial learning environments that technology comes into the picture. In our work, the technology is Virtual Reality.

The purpose of this article is to suggest an explanatory framework for learning in artificial environments, based on second-order cybernetics and other related theories. We first examine the kinds of learning environments to which people can

-
- [1] Some of the research reported here was supported by a grant to the first author from the National Science Foundation, REPP award no. 9873620.
 - [2] William Winn, College of Education and Human Interface Technology Laboratory, University of Washington, Seattle, WA, USA
 - [3] Mark Windschitl, College of Education, University of Washington, Seattle, WA, USA

adapt and how learning occurs in these environments. This examination leads us to conclude that a useful way to view artificial learning environments, and the students who work in them, is to think of both as dynamic systems, closely coupled in time and space through activity and feedback, in which the mechanisms for interaction are as transparent as possible, and the consequences of adaptation deep and lasting changes in understanding. We use our research to illustrate a number of these ideas.

Learning Environments

Learning occurs in just about every environment imaginable. It is useful, however, to identify four types of environment where the manner and results of learning are noticeably different. We call these: Natural environments, structured natural environments, formal environments and artificial environments.

Natural environments

A lot of learning goes on in the natural world, unconstrained and unmodified to fit educational goals. An early treatise on this kind of learning is Rousseau's *Emile* (1762/1933). Rejecting the influence of contemporary educational and social institutions, which he believed corrupted a person's true nature, Rousseau argued that the best and truest education was to be had from free interaction with nature itself. Although Rousseau's motives were as much political as they were educational, echoes of his ideas are to be found in the work of Dewey, and today among some proponents of situated learning and constructivism.

Yet our social institutions and, more generally, our culture, its norms and artifacts are as much part of our natural environment as nature was to Rousseau. Piaget's (1969) theory of cognitive and social development is clearly drawn from the study of how a child adapts to such an environment. "Equilibration" (Piaget, 1968) is a dynamic process that attempts to achieve a state of cognitive equilibrium within a child by fitting new natural and cultural experiences to existing schemata — "assimilation" — while at the same time modifying existing schemata to incorporate new experiences — "accommodation". New experiences occur when the child interacts with the environment through active exploration. The ensuing, iterative, modification of schemata ensures the development of a child's perceptual, cognitive and social abilities. Contemporary constructivist approaches to learning likewise seek to let learning occur through active exploration of the environment.

Structured natural environments

Some learning occurs in natural environments that have been modified in order to allow learning to be more formally directed towards educational goals, or in natural environments in which students' activities are organized so that they learn particular things. These ideas have been articulated by proponents of "situated learning" under the rubric of "cognitive apprenticeship" (e.g. Brown, Collins &

Duguid, 1989). The idea here is that learning out of context, such as learning Mathematics by doing pages of examples in a classroom, does not lead to a good and useful understanding of Mathematics. Learning in a context where Mathematics is applied to real-world problems, it is claimed, leads to a deeper understanding. A good example of how this works is the "Adventures of Jasper Woodbury" video project, directed by Bransford and the Cognition and Technology Group at Vanderbilt University (Cognition and Technology Group at Vanderbilt, 1997, 2000). In this series of videos, Jasper Woodbury and his friends encounter problems that require mathematical solutions as they go about their lives. The problems are authentic and often difficult to solve. Students are challenged to solve the problems using information found in the stories and with guidance from supplementary materials.

Other examples of apprenticeship, in natural environments where students are led towards educational goals, are to be found in the case studies presented by Lave & Wenger (1991) as examples of "legitimate peripheral participation". Here, apprentices perform tasks that are essential to the successful completion of a mission, but are initially peripheral to it. One such case describes how junior seamen in the United States Navy undergo apprenticeships under the guidance of senior Chiefs while they are working on a ship at sea. The seamen's learning experiences are carefully structured. They begin by learning to perform tasks that are peripheral to the running of the ship, such as operating sonar. It is only towards the end of their apprenticeship that they take the helm, an activity that is central to running the ship, and where a mistake could have very serious consequences.

Even in such structured environments, a lot of important learning occurs naturally, outside the framework of formal apprenticeship. Apprentices become socialized to a craft or profession by listening to the stories told by their more senior mentors in off-duty moments. The knowledge thus acquired has been called "stolen knowledge" (Brown & Duguid, 1996), "stolen" because those in possession of it are not conscious of imparting it (nor, more often than not, the students of receiving it).

Formal learning environments

The idea that learning is something that goes on in a place separate from where a person lives or works, and at specific times of the day and year, is the exception rather than the rule at the scale of human history. Yet that is how the formal environment we call "school" operates. What is it that students adapt to in school? Brown (1997) suggested that students adapt to the institution of schooling itself, when she asserted that the only thing students learn in school is how to be students. There is certainly some truth in this. To succeed, or just to survive, some students set a priority of discovering how "the system" works and then of complying with it. For instance, they might try to determine what a teacher expects as answers on tests or ideas in a paper, and provide them.

However there is more to say about the learning environment provided by schools than that. The most important characteristic of formal learning environments is that they largely provide *interpretations* of the natural world for students to interact with. The teacher, the textbook, the video and the computer program all mediate between the student and the natural world, providing *representations* of it. We expect students to make use of knowledge and skills, acquired thus vicariously, through *transfer* to more authentic situations outside the school.

Because formal environments only provide these kinds of "third person" experiences of the world, students must master two skills beyond those needed to learn in other environments. The first of these is skill with the symbol system in which the information to learn is represented. (A problem arises, of course, when, as is often the case, schools test students on the symbol system and not the concepts and principles it encodes.) The second skill helps students to make appropriate connections between the formal learning environment and the environment to which knowledge and skill must be transferred. Sometimes, this can be achieved by directly teaching strategies for mapping formal learning onto practical problems. Sometimes, this can be achieved by giving students internships in natural environments that run parallel to their more formal education. At other times, this can be achieved by making education the responsibility of an entire community (Abbott, 1993). In any event, like learning symbol systems, learning how to transfer knowledge and skill from a formal environment to a natural one requires effort that is not otherwise needed to learn.

Artificial environments: Simulation, Reification and Transduction

Artificial environments can be facsimiles of any of the three types of environment described so far. They can consist of complete physical mock-ups of natural environments, like flight simulators; computer-based simulations of natural environments, like virtual nuclear reactors that allow technicians to practice repairs without being exposed to radiation; or computer-based "reifications" (Winn, 1993) of concepts or processes that are not normally accessible to our senses.

In contrast to simulations, where the goal is to recreate, in hardware or software, as accurate a replica of a real environment as possible, reifications give perceptible form and behavior to objects and phenomena we cannot normally see, hear or touch. Artificial environments are created from data. For example, a virtual chair is modeled in three dimensions by graphics software and appears to the student, on a computer screen or in a head-mounted display (HMD)⁴, as a recognizable chair. Likewise, changes in ocean salinity, which cannot be experienced directly, can also be modeled from data. These data can be rendered to appear on a screen or in an HMD in any form you want. In our research, we have used animated isosurfaces for this, connecting points in the water column having the same salinity. The point is that the computer does not render data differently depending on whether they represent real objects or reifications of

abstractions. This means that, in artificial environments, *models of abstractions can be created, experienced, interacted with and controlled in exactly the same way as models of real objects.*

There are two important consequences for learning that ensue from this property of artificial environments. The first is that, although they may contain reified abstractions, artificial environments can be far more like natural environments than school environments are. Well-designed artificial environments meet three criteria (Zeltzer, 1991). They permit students to experience high levels of "presence" — the sense that you are "in" the environment, not in a classroom looking at a computer screen or wearing an HMD. They are interactive, tightly-coupled systems in which the consequences of a student's natural actions are accurately revealed in real time by how the environment behaves. They are autonomous — a measure of their behavior is independent of what the students do or do not do. When these conditions are met, and when abstractions, reified as virtual objects, are endowed with appropriate properties and behaviors, students can construct knowledge of concepts and principles directly, in the "first person", as in a natural environment, without having to learn an abstruse symbol system first (Dede, 1995).

The second important consequence of an artificial environment's ability to reify abstractions is that they can act as transducers. A transducer is a device that shifts information in a bandwidth that is beyond humans' perceptual capabilities into a bandwidth that is within them. A good example is a marine depth sounder. This device bounces a sound that we cannot hear, off the seabed that we cannot see, and converts the returned echo into a digital or analog display on a video screen that we can see and interpret. The pedagogical power of artificial environments relies to a great extent on their power to act, in this way, as transducers.

Two artificial environments in which transduction is used extensively illustrate this idea well. They are Dede's "Maxwell World" (Dede et al., 1997) and Byrne's "Atom World" (Byrne, 1996). Both of these are immersive three-dimensional simulations that meet Zeltzer's criteria for effective artificial environments. Both have served as test beds for research on virtual reality and both have been used in schools.

"Maxwell World" simulates aspects of electrostatic fields described by Maxwell's equations. Wearing an HMD that immerses the student in the environment, the student sees a three-dimensional representation of an electrostatic field, showing lines of force, charges and so on, that cannot otherwise be seen. In addition to moving around outside and inside the field to experience three-dimensional views of it from different positions, the student can perform operations on it and immediately observe the consequences. For example, the tip of the student's finger can "become" a source charge. By moving the hand, the student can move a virtual charge around in the field and observe how it affects the lines of force.

Byrne's "Atom World" first "shrinks" students to atomic size and then allows them to construct atoms, by hand, from a supply of protons, neutrons and

electrons. At this scale, subatomic particles are roughly the size of an apple, and are easy to handle. To construct each atom, the student must discover how many electrons can exist in each orbital, the spin of each electron, the energy level of electrons in each orbital, and the relationship between the nucleus and the electrons in a stable atom. The environment's behavior provides a variety of audible and visual feedback to guide the student's activity. Subatomic particles appear as one might expect, a proton as a three-dimensional cross, a neutron as a cube and an electron as a three-dimensional minus sign. The orbitals themselves are rendered as transparent, shimmering spheres, nested around the nucleus. Virtual levers, not unlike an automobile's automatic transmission shifter, control electron spin and energy level.

In spite of the demonstrated success of these and many other artificial environments, there is an important caveat. Reification and transduction are selective processes that rely on metaphors to stand for abstractions — arrows for lines of force, cubes for neutrons, and so on. Sometimes, inappropriate metaphors can induce rather than correct misconceptions. For example, in our artificial environment that simulated global warming (Jackson, Taylor & Winn, 1999), the amount of biomass available to absorb greenhouse gases (a relatively complex abstraction) was reified through the appearance and disappearance of trees in the virtual landscape. (By using this metaphor, we intended to link global warming to the destruction of the rain forests.) A number of twelve- and thirteen-year-olds came to believe that global warming was not a serious problem: all you had to do was plant more trees! The tree metaphor grossly oversimplified the environmental processes. The students forgot (or perhaps never realized in the first place) that the trees were a metaphor that stood for a more abstract, complex and general idea.

Learning in Artificial Environments

We have characterized learning as adaptation to an environment. We have seen, in particular, that in the traditional view, adaptation to formal learning environments requires mastery of the symbol systems in which information is represented and the ability to interpret descriptions of the world, rather than interacting directly with it. It develops what Scribner (1997) has called "theoretical thought" rather than "practical thought". We have also seen that learning in formal environments requires the ability to transfer knowledge and skill to the real world. The purpose of artificial environments is to overcome these difficulties using simulations, while at the same time maintaining the ability to show abstractions directly by reification and transduction using carefully-selected metaphors. This results in a first-person experience of the environment. In this section we examine how this can be done and develop our conceptual framework. We provide illustrations in the next section.

The cognitive science view of learning

The traditional view of learning from cognitive science describes an iterative process by means of which what a person comes to know is changed by new information and experiences, which in turn are interpreted in terms of what the person knows already (Neisser, 1976; Winn & Snyder, 1996). As what a person knows and the interpretation of new experiences change, the two converge until a person is satisfied that he or she understands something new, or, with the failure of convergence, gives up trying to make sense of the new experience. This convergence has been variously described as "equilibration" (Piaget, 1968), a "conversation" leading to "agreement over an understanding" (Pask, 1975, 1984), "schema tuning" (Rumelhart & Norman, 1981) or "conceptual change" (Posner et al., 1982). Moderate constructivist views of learning (for example, Spiro et al., 1992; Cognition and Technology Group at Vanderbilt, 1997) subscribe to the same basic account.

While this account of learning provides a useful conceptual framework within which to conduct experimental research on learning in the laboratory, and describes fairly well how many studying formal learning environments think students learn, it is weak when it comes to explaining how whole systems of learner and environmental variables interact as knowledge is constructed. It is also insufficient when it comes to explaining the effects of context on learning, particularly in the absence of direct instruction or even guidance from a teaching "agent". Finally, and most critically, this account assumes "turn-taking" — which Clark (1999, pp. 105-106) has evocatively called "catch and toss" — between an environment and student that are separated. On this account, the cognitive side is isolated inside the student's head, making computations with symbols caught from the environment and tossing back hypotheses to test, or eventually solutions to problems. Cognitive scientists subscribing to this account have been concerned with developing a "functional architecture" for cognition (Anderson, 1990; Pylyshyn, 1984) rather than accounts of cognition based on biology or activity. Not surprisingly, the notion of functional architecture has led to a misleading "mind = computer" analogy (Johnson-Laird, 1988; Boden, 1988), misleading because, although a computer can serve cognitive scientists well as a research tool that can simulate some aspects of cognition, it does not follow that cognition must therefore be computer-like. Searle (1992) commented wryly, "If you subscribe to functionalism, you don't need refutation, you need help."

We have seen that in formal schooling environments, students are indeed isolated from the natural environment by intervening symbol systems and by the operations needed to transfer what they learn to the real world. In artificial environments, as in natural environments, students enjoy unmediated experiences of the environment and are not isolated in this way. The reason that the traditional, computational, functionalist view of cognition fails to provide an adequate explanatory framework for learning in artificial environments therefore lies in two properties of the latter: Artificial environments provide students with first-person experiences from which to construct knowledge; they afford dynamic interactions

in which action and feedback are indistinguishable, where the time intervals between the "catch" and the "toss" are vanishingly small.

Embodiment, embeddedness, environment and Umwelt

The dynamic, first-person nature of interacting with artificial environments is seen best in an illustration. Clark (1999, pp. 171-2), elaborating on an example from Merleau-Ponty (1963), invites us to imagine trying to catch a hamster with a pair of tongs. The animal's evasive maneuvers are immediate and uninterrupted reactions to how we wield the tongs. At the same time, the movements of our hands and the tongs are driven by what the animal does. The two sets of actions are so completely coupled that they can only be thought of as the actions of a single system. There are in fact three pieces to this system: our brain, our body, and the environment containing the hamster. Since we do not consciously decide how to move the tongs, we can say that our cognition is embodied in our hands. And since our embodied actions are inseparable from the environment, we can say that we, and our actions, are embedded in the environment. Embodiment (Varela, Thompson & Rosch, 1991) and embeddedness (Clark, 1999), thus construed, feature prominently in second-order cybernetic theories of cognition — Maturana and Varela (1980, 1992) describe how organisms are structurally coupled to their environments and to other organisms in it — and provide a useful start for our account of learning in artificial environments. Roth (1999) makes the case that it is often a better strategy to consider mind, body and environment as one thing, rather than three. Similar arguments have been made by Rosch (1999) and Kelso (1999), who proposes a "Dynamical Systems" approach to interaction and adaptation.

The first-person nature of experience in artificial environments requires us to define "environment" more carefully. We are talking about a unique, personal environment brought forth by an individual, not some standard, objective world about whose nature and behavior we can all agree. Von Uexküll (1937) proposed the word "Umwelt" for such environments. (See also Roth, 1999.) Von Uexküll (1937, reproduced in Clark, 1999, pp. 26-28) shows somewhat whimsical drawings of how the world might look to various creatures, for example bees and scallops. Like these drawings, artificial environments can transduce information from Umwelten outside the experience of human beings, and bring them into an individual human's Umwelt. However, there are limitations to the extent that phenomena can be transduced in this way. Nagel (1974) reminds us that we may well be able to "see" what a bat "sees", by transcribing the sounds of echolocation into sonograms, or shifting their frequencies into the range audible to humans. But however cleverly we do this, we can never really know what it is like to be a bat.

Finally, Umwelt gives us a better way than before to conceptualize the relationship between a student and an environment in a dynamic, tightly-coupled system. Since artificial environments provide first-person experiences to a student, and since Umwelt refers to the unique personal environment that the student inhabits, it is now possible to talk about the student's actions bringing

about continuous and immediate changes in the Umwelt, just as the environment brings about changes in the student, through structural coupling (Roth, 1999). Actions and their consequences are now a genuine two-way street, which Clark (1999, p. 163) calls "continuous reciprocal causation", and which is also a good way to think of structural coupling (Maturana and Varela, 1992).

Learning as adaptation

How does learning occur in such an environment? Using the language and ideas of second-order cybernetics (Maturana and Varela, 1992), and noting their basis in biology, we can describe a process whereby a student encounters a perturbation in the environment to which she adapts by changing her structure. If the change in structure is such that it leads an observer to notice a more appropriate behavior the next time the perturbation is encountered, we can say that learning has occurred. This account rules out the possibility that information from the environment is somehow reproduced as a "mental map" in the student's head. Instead, it proposes structural coupling as a mechanism whereby a student is continuously, tightly and dynamically interacting with the Umwelt, in ways that we described in the last section. It also means that the student's Umwelt changes in equal measure to the student's structure.

We must also consider some elaborations to this basic account. First the informational closure of autopoietic entities does not inevitably lead to solipsism (see Maturana and Varela, 1992, pp. 133-137). Indeed it must not. We can communicate, through language in all its forms, because, through structural coupling, our histories of adaptations are similar to those of other people. We can also become observers of student-environment interactions and provide descriptions of what we see. Second, learning through interaction with the environment does not preclude learning from interactions among operations within the individual. Steffe (1999), for example, stresses the importance of both types of interaction in learning Math, stating that experience alone is insufficient for the construction of mathematical knowledge. Third, a lot of learning requires that students become their own observers. Thus, Keiny (2000) describes the "ecological thinking" of a group of Israeli teachers collaborating on the development of a new curriculum. In this learning community, each teacher was both an actor within the community and an outside reflector on it. Keiny's "ecological thinking" was inspired by and is similar to Schön's (1983) concept of "reflection in action", a concept by now quite familiar to the broader educational community.

How is a student's structure changed as a result of adapting to a perturbation? Reyes and Zarama (1998) suggest that knowing arises from our inventing distinctions among what we experience. As they say, we learn to distinguish an object from its background. We also distinguish objects from each other. How we make distinctions is intimately tied to our ontology, the distinctions we make are very personal, and say "more about ourselves than about the world we are describing" (Reyes & Zarama, 1998; p. 23). Once made, a distinction becomes

grounded in our history of structural change and, as we saw on several earlier occasions, embodied in our actions. We give examples of this process from our observations of students visiting artificial environments below.

For a possible account of the mechanisms of structural change through structural coupling, we look further afield. Holland (1975, 1995) allows us to maintain the biological flavor of our account so far by describing adaptation as emerging from the operation of sets of rules combined to form "genetic algorithms". To begin with, as part of his ontogeny, a student has a set of rules that guide his reactions to perturbations in the environment. It is important to note that, for Holland (1995, p. 53), rules are not facts but rather competing hypotheses whose influence and survival depend on how successfully they direct responses to perturbations. Adaptation can occur by "credit assignment", where the status and viability of a rule waxes and wanes according to its success in directing responses to the environment. Adaptation can occur by combining pieces of rules to form new ones in the same way that genes do. And adaptation can occur when one rule simply replaces another one. In successful applications of Holland's approach, it is not possible to predict the nature of the adaptation simply from knowing the rules. In other words, adaptation, and thus learning as we have construed it, is an *emergent property* of interaction with the environment and organisms that learn, like students, are *self-organizing*. Thus Ison and Russell (2000) present case studies in which phenomena are seen to emerge from observation.

We are not particularly committed to Holland's description of adaptation. After all, he is concerned with building, in software and in hardware, "intelligent agents" and robots, activities quite removed from our work and from this discussion. However, we feel his kind of approach holds promise for filling in and extending the account of learning that we are proposing. First, the emergent nature of learning and the self-organizing capability of learners stands as a safeguard against the increasing fragmentation of knowledge as our students make more and more distinctions. Both the cognitive (Merrill, 1992) and contextual (Lakoff, 1985; Rosch, 1999) approaches to learning have stressed the value of our ability to form categories into which we can classify our experiences. Such categories might emerge from adaptations to the environment in the manner that Holland describes. Second, self-organization and emergence are receiving empirical support, from a variety of scholars, as models of aspects of cognition. Kelso's (1999) experiments and subsequent conceptual analysis are compelling. Finally, our own research has shown that students test and modify rules (hypotheses) as they adapt to perturbations in artificial environments, just as Holland (1995) describes. We give examples later.

An explanatory framework for learning in artificial environments

In this section, we list eight ideas that help explain learning in artificial environments. These are derived from topics we have discussed so far, but are not in any particular order.

1. Students, like all living organisms, are informationally closed. A student's adaptation to environmental perturbations does not result in information from the environment being "mapped" in some way onto the student's mind. The "picture in your head" view of memory and, increasingly, the computational view of cognition are commonly discredited. Nonetheless, there persists even today in some circles the idea that we carry in our heads more or less realistic replicas of the world, and that these direct our behavior and determine what and how we learn. This has led to some strange suggestions from educators, for example that schools should "activate students' right hemispheres" where mental images are supposed to reside. We can now firmly discount such suggestions.

2. Learning is a biological process. This has two implications. The first is that cognitive processes are intimately tied to the mechanisms that distinguish living (autopoietic) beings from inert ones. Expressed another way, our mind is "embodied" (Varela, Thompson & Rosch, 1991) in our brain and body, not something separate and mystical. Arguments in support of this position have been put forward by Pinker (1997), Dennett (1995, mostly in pp. 335-451), Bruer (1999a), and others. The second implication is that neurobiological accounts of cognition are relevant to education. For instance, we are beginning to understand some of the neural mechanisms that support learning, (incidentally, rendering unnecessary the need for a "functional architecture" for cognition). Also, we now know that the brain is more plastic than hitherto believed (Bruer, 1999b). Synapses and neurons can be created long after the time, in the course of biological and cognitive development, when such activity was once thought to cease. This means that the ability of people to adapt to an environment, whether natural, formal or artificial, continues throughout a lifetime.

3. Students and environments are complex, interacting systems. We must therefore consider all of the factors that impinge on learning, not just a few isolated dependent and independent variables. If artificial environments are to be facsimiles of the natural environment, as in high fidelity simulations, then they must afford modes of learning to students that emulate the way people learn in the natural world. Theories of informal learning suggest that the act of constructing knowledge is not naturally constrained to a limited number of experiences or sources of information (independent variables) and changes to a student's structure (dependent variables). Indeed, Salomon (1991) has made the case that useful educational research *must* examine whole sets of variables at once. Artificial environments that simulate natural environments must describe interacting networks of systems of variables.

4. We must maintain a sensible balance between student autonomy and providing guidance to them. Artificial environments must allow students considerable freedom to choose experiences and, especially, must be tolerant of idiosyncratic adaptations as students interact with the environment. This does not mean that a "sink or swim" approach to learning is appropriate. Even in unstructured, informal learning, students need guidance (often referred to as "scaffolding" [Bell, 1997; Linn, 1995]) that, if nothing else, makes learning more

efficient. For this reason, artificial environments often build pedagogical strategies, cunningly, into the behavior of the environment itself (Winn & Bricken, 1992). The embedding of critical information into the Jasper video stories (Cognition and Technology Group at Vanderbilt, 1997), mentioned above, comes close to meeting this goal.

5. Students are embedded in their environment to form, in effect, one system. These systems are dynamic. The give-and-take that occurs with adaptation can rarely be separated out into discrete events in time. Behavior and learning are emergent properties of these systems. In artificial environments, this means that the way the environment reacts to students' actions, and acts autonomously for that matter, must be immediate, except when immediate reaction would not occur naturally. The interface that the student uses to interact with the environment must not get in the way of the interaction, either by being cumbersome or non-intuitive, or by slowing activity down. It follows that research data obtained about students learning in artificial environments is best described and analyzed using the mathematical methods of Dynamical Systems Theory (Port & Van Gelder, 1995; Kelso, 1999). The statistical techniques we are all familiar with are not up to the task.

6. Learning is a social activity (Vygotsky, 1978). The environments to which students adapt more often than not contain other people. Much current research on how technology can support learning examines the discourse that occurs in a community of learners as they work together in a learning environment. For example, Bell (1997) has created a web-based learning community of students, teachers and scientists that examines scientific controversies, such as genetically altered foods. His units of analysis are, in large part, the exchanges that take place within this community. We note, though, that immersive artificial environments, including many of ours, are designed to accommodate only one student at a time. With new, cheaper, technology, this is beginning to change and techniques like Bell's will become increasingly important to our work.

7. To teach, you must perturb the environment to induce adaptation. In natural environments, this often occurs without needing to be provoked. This is not the case in artificial environments where every event arises either from within the software or from within the student. The nature of the perturbation is critical. In certain respects, artificial environments have a lot in common with computer games (Hedden, 1998) and embody the same elements of goal attainment and challenge that one finds there. Perturbations act as surprises, as puzzles to solve. Adapting to them must somehow bring the student closer to a satisfactory conclusion, such as constructing an atom or understanding lines of force — or slaying a demon or rescuing a princess. Also, perturbations must motivate the student. They must be challenging but not excessively so, and dealing with them must be fun. Finally, they must lead students to question their assumptions, resulting from earlier adaptations. Without this open-mindedness and tolerance for ambiguity, students will not modify the hypotheses that are driving their adaptation, and no learning will occur.

8. In artificial environments, reified abstractions have equal phenomenological status with models of real objects. This allows students to enjoy first-person experiences of objects and events that they would otherwise not have. However, when we use reification and transduction as ways to perturb the environment to induce adaptation, we must be wary of poorly-chosen metaphors that lead to highly idiosyncratic changes in students' structures. The environment that is seen by each student, the *Umwelt*, arises from the student's history of prior adaptations, and will vary considerable from one to the other.

Illustrations from our research

In this section we illustrate many of the ideas we have presented so far from research that has looked at how students learn when they visit artificial environments. Our research group is studying the way in which artificial environments, with the capabilities to simulate, reify and transduce phenomena, can help or hinder students' understanding about complex natural phenomena. One environment, "Virtual Puget Sound", is an immersive simulation of tidal movement and other related underwater phenomena in Puget Sound, Washington, USA.

Virtual Puget Sound

Virtual Puget Sound (VPS) uses simulation, reification, transduction, immersion and interaction to create an artificial environment that recreates important properties of the real marine environment while allowing the student to go to places, observe phenomena, make measurements and control events in ways that are impossible in reality. VPS is a three-dimensional model of the Puget Sound region, roughly 200 kilometers North to South, seventy kilometers East to West, from the bottom of the ocean to high in the atmosphere above it. Tidal currents are simulated for one tidal cycle (just under twenty-four hours), sampled at 30-minute intervals, at twelve depths, on a grid of 600-meter squares. In normal operation, the currents are animated in a 48-frame loop, showing one 24-hour cycle in about four seconds. The bathymetry (underwater topography) is modeled accurately, but with a ten-times vertical exaggeration to make it easier to "swim" around in the model. A texture map, created from a satellite photograph of the region, lies over the land, creating a good substitute for the actual topography. The simulation, adapted by our colleague Mitsuhiro Kawase from the Princeton Oceanographic Model of water circulation (Blumberg & Mellor, 1987), provides information about water speed, direction and salinity at different depths at each grid point, and for each time step in the tidal cycle. Currents can be illustrated as vectors — arrows whose length and direction correspond to water speed and direction — or, in a most recent version, advected particles, whose movement mimics that of objects floating in the water at different depths.

So far, we have studied learning about the marine ecosystem using a problem-solving scenario. Students are told that a fish that preys on salmon,

Macrodont horribilis, has recently been discovered in Puget Sound. Should it become established, it will threaten the salmon fishery. Students are given the fish's preference for water speed and salinity in its habitat, and are told to find the most likely location and depth where it can be found. Once they put on an HMD, the students are immersed in the environment. They can visit pre-selected sites, where they may measure current speed, current direction and water salinity at any depth. They can also turn the animation of the tidal cycle on and off and, with animation off, can set the time in the cycle at which they make measurements. They take an objective pre- and posttest of their knowledge of currents, salinity, and so on. They are videotaped and they provide verbal "think-aloud" protocols while they are looking for the fish. Every action they take is logged in a file. They are interviewed after completing the task.

Learning in VPS

The learning strategy around which VPS was built was originally taken from Posner et al.'s (1982) set of conditions for bringing about conceptual change. However, this strategy is remarkably similar to the one proposed by Reyes and Zarama (1998), so for consistency we describe the strategy in their terms.

1. Declaring a break. A perturbation in the environment occurs that is counter to what the student expects and disrupts the flow of actions. In VPS, students have many opportunities to observe phenomena that are counterintuitive and that catch them by surprise. An example might be that tidal currents appear to move at different speeds in different places, which the student did not know before.

2. Drawing a distinction. From this observation, the student might make a new distinction between fast and slow currents that would change both the way the environment appears to the student — fast and slow currents are now noticed — and how the student acts in the environment. (We might also say, using Holland's [1995] ideas, that the student has modified a rule/hypothesis that will consequently alter the process of adaptation to the environment.)

3. Grounding the distinction. The new distinction must be compatible with the student's particular history of adaptations. In this case, the fact that currents in the sea can move at different speeds must be both intelligible and plausible in terms of what the student already knows. If the distinction between fast and slow currents is not intelligible, the student may simply memorize the fact and not come to understand it at all. If it is not plausible, the student may understand it but not believe it. Prior experience determines both what the student can understand and will believe about a new distinction.

4. Embodying the distinction. Simply making and grounding the distinction is not sufficient for learning to occur. It must allow new actions that, in the case of VPS, can be fruitfully employed to help solve the problem of finding the fish. Knowing that *Macrodont* likes faster-moving water, the student might then set out to visit only sites that fell on the "fast" side of the distinction between current speeds.

With VPS built, and the learning strategy embodied in it through its behavior and through scaffolding, we have conducted studies of how undergraduates, middle school students (ages 13 — 14) and dyslexic ten- to twelve-year olds fared during their visits to Virtual Puget Sound.

Observations

Complete reports of the findings from our studies have been published elsewhere (Windschitl & Winn, 2000; Winn et al., 2001). Here, we simply describe things we have observed and documented that illustrate aspects of the explanatory framework that we laid out in the previous section.

1. Students embody their ideas in actions. Immersive environments are very physical, visceral places. In VPS, students have to be physically active to learn. They must move their bodies in the environment to get to the places where they can make observations and measurements. They control this movement by gestures. They control their depth in the water column by moving a virtual lever up and down. They control the state of the tide by sliding a ball along a tide graph. They make measurements by pressing buttons. Each of these actions reveals to the observer how students are adapting to the environment and allows an assessment to be made about what knowledge the student has (see Maturana and Varela, 1992, pp. 170-172, and Reyes and Zarama, 1998, p. 21-22). More interesting are the occasions when students use actions to give descriptions of what they understand. For example, one of our younger dyslexic students had difficulty using words to explain his new discovery, that water enters Puget Sound at depth and leaves at the surface. Yet he mimicked the circular nature of this movement perfectly with his arms and hands.

2. Students adapt to perturbations by drawing distinctions. On the videotapes and in the verbal think-aloud protocols are many examples of this. A university undergraduate discovers by measuring salinity at different depths that it varies with depth and comments, "I didn't know that". Another student notices that particles released into the water do not return to their starting point at the end of the tidal cycle. Furthermore, he notices that their motion is not all that well coordinated within the particle cloud. In many of these cases, students embody their new distinctions in actions. The first student sets about exploring the manner in which salinity varies, using virtual tools, and finds that it increases with depth. The second student suggests several reasons why particles might move like they do and goes to check them out.

3. Many factors affect students' actions in VPS. VPS is a relatively complex environment and many factors influence how students act when they visit it. Some of these arise from within the environment itself, like skill in moving around, comfort with the environment and interface, the extent to which a student feels "present" in the environment, the length of time spent there and how frequently and well the virtual tools are used. Many other factors come from outside the environment. They include skill with the interface (generally, the younger students are the best at this), familiarity with the region, and gender. We find it necessary

to consider all of these factors and the complexity of VPS itself when we generate descriptions of what we see happening.

4. Students differ in how they deal with complexity. Students use a variety of strategies to deal with the complexity of VPS. Some students simplify the complexity by turning some things off, such as time. With the moment frozen, it is easier to make consistent measurements of salinity, say, and to draw more reliable inferences. You are, in effect, controlling a variable as you test a hypothesis. Other students are less systematic and make more or less random measurements until they find conditions that approximate those they are looking for. Yet others (a few) cross-check observations from several vantage points before drawing conclusions. Finally, some are just confused and never really come to understand what is happening. We attribute this variety, obviously, to differences in individual's histories of prior adaptations. But we acknowledge that VPS is weakly scaffolded. This means that happenstance allows some students to solve the problem of the fish just by happening to make the right measurement or observation at the right place and time. Others are as equally unlucky. While the element of chance also operates in similar ways in the real world, it is unfair to students for it to play a role when their time in the artificial environment is limited.

5. The student-environment system's actions exhibit continuous reciprocal causation. Wearing an HMD inevitably embeds a student in the environment. And tracking the position of the HMD so as to continuously update the visual display guarantees that even just looking around will produce an instantaneous and completely logical change in the environment. The actions students perform to accomplish tasks are likewise tightly coupled to the environment. In a relatively simple operation, like setting your depth in the water column, a movement of the depth lever causes an instantaneous change in perspective and in a numerical read-out, which immediately indicates how much more to move the lever. Watching students perform a trickier operation, such as setting time in the tidal cycle, especially if they have forgotten to turn time off, really does remind us of trying to catch a hamster with tongs! More generally, learning is greatly assisted by the instant reaction of the environment to students' actions. The think-aloud data show that rapid confirmation of hypotheses allows students' knowledge construction to proceed at the rate the student sets, unimpeded by delays on the part of the environment.

6. Poorly-chosen metaphors lead to idiosyncratic adaptation. As we said, our first version of VPS used the length and direction of arrows to reify current speed and direction. The majority of the students could interpret this metaphor. However, in some cases, the metaphor caused problems. As a case in point, students needed to find out whether tidal currents speed up or slow down as they pass through narrow passages. (They speed up.) In a post-experience interview, one student was certain the currents slow down. He explained that the greater length of the arrows in narrow passages, supposed to show greater speed, made the passages look clogged. When he drives on the highway, and the traffic gets

clogged, it slows down. Another subject simply said he did not know, but reasoned by analogy from a knowledge of aerodynamics that water would probably speed up because air does when it passes through a narrow tube. In both of these cases, we have examples of different interpretations of the same metaphor, based on different prior experiences — driving in traffic and aerodynamics. Purely by chance, one happened to be “right” and one “wrong” by our criteria. But we are reminded, yet again, of the importance of prior history on the adaptations students make to artificial environments.

Conclusion

We started out with the notion that the traditional explanation for learning, based in cognitive science, could not adequately explain what we have observed in students working with immersive artificial environments. Our reading of a wider literature, centered on second-order cybernetics but including related topics such as emergence and self-organization, biological adaptation, and dynamical systems theory, led us to a framework that we think better accounts for some aspects of our findings. It does so, we believe, because it takes into account both the unique properties of artificial environments and the full repertoire of activities that students can use to learn by adapting to them. The explanatory framework that has emerged is exactly that: explanatory, not prescriptive; a framework, not a theory. This means that it is not yet ready to guide the design of artificial environments. And it will require empirical validation, and certainly subsequent revision, before it can attain the status of a satisfactory descriptive theory. We look forward to developing both its prescriptive and theoretical power in future research.

References

- Abbott, J. (1993). The creation of effective modern learning communities: Constructivism in practice. In T.M. Duffy, J. Lowyck., & D. Jonassen (Eds.), *Designing environments for constructive learning*. New York: Springer.
- Anderson, J.R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Erlbaum.
- Bell, P. (1997). Using argument representations to make thinking visible for individuals and groups. In R. Hall, N. Miyake, & N. Enyedy (Eds.), *Proceedings of CSCL97: The Second International Conference on Computer Support for Collaborative Learning*. Toronto: University of Toronto Press.
- Blumberg, A.F., & Mellor, G.L. (1987). A description of a three-dimensional coastal ocean model. In N. Heaps (Ed.), *Three-dimensional coastal ocean models*. American Geophysical Union.
- Boden, M. (1988). *Computer models of mind*. New York: Cambridge University Press.
- Brown, A.L. (1997). Transforming schools into communities of thinking and learning about serious matters. *American Psychologist*, 52(4), 399-413.
- Brown, J.S., Collins, A., & Duguid, P. (1989). Situated cognition and the culture of learning. *Educational Researcher*, 18(1), 32-43.
- Brown, J.S., & Duguid, P. (1996). Stolen knowledge. In H. McLellan (Ed.), *Situated learning perspectives*. Englewood Cliffs NJ: Educational Technology Publications, 47-56.
- Bruer, J.T. (1999a). Neural connections: Some you use, some you lose. *Phi Delta Kappan*, 81(4), 264-277.
- Bruer, J.T. (1999b). *The myth of the first three years: A new understanding of early brain development and lifelong learning*. Riverside, NJ: Simon and Schuster.

- Byrne, C. (1996). *Water on tap: The use of virtual reality as an educational tool*. Ph.D. Dissertation, Department of Industrial Engineering, University of Washington, Seattle, WA.
- Clark, A. (1999). *Being there: Putting brain, body and world together again*. Cambridge, MA: MIT Press.
- Cognition and Technology Group at Vanderbilt (1997). *The Jasper project: Lessons in curriculum, instruction, assessment and professional development*. Mahwah, NJ: Erlbaum.
- Cognition and Technology Group at Vanderbilt (2000). Adventures in anchored instruction: Lessons from beyond the ivory tower. In R. Glaser (Ed.), *Advances in instructional psychology, educational design and cognitive science, Volume 5*. Mahwah, NJ: Erlbaum.
- Dede, C. (1995). The evolution of constructivist learning environments: Immersion in distributed, virtual worlds. *Educational Technology*, 35 (5), 46-52.
- Dede, C., Salzman, M., Loftin, R.B., & Ash, K. (1997). Using virtual reality technology to convey abstract concepts. In M. Jacobson & R. Kozma (Eds.), *Learning sciences for the 21st century: Research, design and implementing advanced technology learning environments*. Mahwah, NJ: Erlbaum.
- Dennett, D. (1995). *Darwin's dangerous idea*. New York: Simon and Schuster.
- Hedden, C. (1998). *A guided exploration model of problem-solving discovery learning*. Ph.D. Dissertation, College of Education, University of Washington.
- Holland, J. (1975). *Adaptation in natural and artificial environments*. Ann Arbor, MI: The University of Michigan Press.
- Holland, J. (1995). *Hidden order: How adaptation builds complexity*. Cambridge, MA: Perseus Books.
- Ison, R.L., & Russell, D.B., (2000). Exploring some distinctions for the design of learning systems. *Cybernetics and human knowing*, 7(4), 43-56.
- Jackson, R., Taylor, W., & Winn, W.D. (1999, April). Peer collaboration in virtual environments: An investigation of multi-participant virtual reality applied in primary science education. Paper presented at the annual meeting of the American Educational Research Association, Montreal, Canada.
- Johnson-Laird, P.N. (1988). *The computer and the mind*. Cambridge, MA: Harvard University Press.
- Keiny, S. (2000). Learning as knowledge construction within a community of learners: An attribution to Don Schön. *Cybernetics and human knowing*, 7(2-3), 31-45.
- Kelso, J.A.S. (1999). *Dynamic Patterns: The self-organization of brain and behavior*. Cambridge, MA: MIT Press.
- Lakoff, G. (1985). *Women, fire and dangerous things*. Chicago: University of Chicago Press.
- Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge: Cambridge University Press.
- 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.
- Maturana, H.R., & Varela, F.J. (1980). *Autopoiesis and cognition*. Boston: D. Reidel.
- Maturana, H.R., & Varela, F.J. (1992). *The tree of knowledge*. Boston: Shambhala.
- Merleau-Ponty, M. (1963). *The structure of behavior*. New York: Beacon.
- Merrill, M.D. (1992). Constructivism and instructional design. In T. Duffy & D. Jonassen (Eds.), *Constructivism and the technology of instruction: A conversation*. Hillsdale, NJ: Erlbaum.
- Nagel, T. (1974). What it is like to be a bat. *Philosophical Review*, 83, 435-450.
- Neisser, U. (1976). *Cognition and reality*. San Francisco: Freeman.
- Pask, G. (1975). *Conversation, cognition and learning*. Amsterdam: Elsevier.
- Pask, G. (1984). Review of conversation theory and a protologic (or protolanguage), Lp. *Educational Communication and Technology Journal*, 32, 3-40.
- Piaget, J. (1969). *The psychology of the child*. New York: Basic Books. [Translated by H. Weaver].
- Piaget, J. (1968). The role of the concept of equilibrium. In D. Elkind (Ed.), *Six psychological studies by Jean Piaget*. New York: Vintage Books.
- Pinker, S. (1997). *How the mind works*. New York: Norton.
- Port, F.P., & Van Gelder, T. (1995). It's about time: An overview of the dynamical approach to cognition. In F.P. Port & T. Van Gelder (Eds.) *Mind as motion: Explorations in the dynamics of cognition*. Cambridge, MA: MIT Press.
- Posner, G. J., Strike, K. A., Herson, P. W., & Gertzog, W. A. (1982). Accommodation of scientific conception: Toward a theory of conceptual change. *Science Education*, 66 (2), 211-227.
- Polyshyn, Z. (1984) *Computation and cognition: Toward a foundation for cognitive science*. Cambridge, MA: MIT Press.
- Reyes, A., & 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. In R. Nunez & W.J. Freeman (Eds.), *Reclaiming cognition*. Bowling Green, OH: Imprint Academic.
- Roth, W.-M. (1999). The evolution of Umwelt and communication. *Cybernetics and human knowing*, 6(4), 5-23.
- Rumelhart, D., & Norman, D. (1981). Analogical processes in learning. In J.R. Anderson (Ed.), *Cognitive skills and their acquisition*. Hillsdale, NJ: Erlbaum.
- Rousseau, J.-J. (1762/1933). *Emile*. New York: E.P. Dutton. [Translated by B. Foxley].
- Salomon, G. (1991). Transcending the qualitative-quantitative debate: The analytic and systemic approaches to educational research. *Educational Researcher*, 20(6), 10-18.
- Schön, D. (1983). *The reflective practitioner*. London: Temple Smith.
- Scribner, S. (1997). Thinking in action: Some characteristics of practical thought. In *Mind and social practice: Selected writings of Sylvia Scribner*. Cambridge: Cambridge University Press.
- Searle, J. (1992). *The rediscovery of the mind*. Cambridge, MA: MIT Press.
- Simon, H.A. (1981). *The sciences of the artificial*. Cambridge, MA: MIT Press.
- Spiro, R., Feltovitch, P.J., Jacobson, M.J., & Coulson, R.L. (1992). Cognitive flexibility, constructivism and hypertext: Random access instruction for advanced knowledge acquisition in ill-structured domains. In T. Duffy & D. Jonassen (Eds.), *Constructivism and the technology of instruction: A conversation*. Hillsdale, NJ: Erlbaum.
- Steffe, L. (1999). Individual constructive activity: An experimental analysis. *Cybernetics and human knowing*, 6(1), 17-31.
- Varela, F.J., Thompson, E., & Rosch, E. (1991). *The embodied mind*. Cambridge, MA: MIT Press.
- Von Uexküll, J. (1934). A stroll through the worlds of animals and men. In K. Lashley (Ed.), *Instinctive behavior*. London: International Universities Press.
- Vygotsky, L.S. (1978). *Mind in society*. Cambridge, MA: Harvard University Press.
- Windschitl, M., & Winn, W.D. (2000). A virtual environment designed to help students understand science. In B. Fishman & S. O'Connor-Divelbiss (Eds.), *Proceedings of the Fourth International Conference of the Learning Societies*, 290-296. Mahwah, NJ: Erlbaum.
- Winn, W.D. (1993). *A conceptual basis for educational applications of virtual reality*. HITL Technical Report R-93-9. Seattle WA: Human Interface Technology Laboratory, University of Washington.
- Winn, W.D., & Bricken, W. (1992). Designing virtual worlds for use in mathematics education: The example of experiential algebra. *Educational Technology*, 32(12), 12-19.
- Winn, W.D., & Snyder, D. (1996) Cognitive perspectives in psychology. In D. Jonassen (Ed.), *Handbook of Research for Educational Communication and Technology*. New York: Macmillan.
- Winn, W.D., Windschitl, M., Fruland, R., Hedley, N., & 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.
- Zeltzer, D. (1991). Autonomy, interaction and presence. *Presence*, 1(1), 127-132.

