

Bereiter, C. (1997). *Situated cognition and how to overcome it*. In D. Kirshner & J. A. Whitson (Eds.), *Situated cognition: Social, semiotic, and psychological perspectives* (pp. 281-300). Hillsdale, NJ: Erlbaum.

Situated Cognition and How to Overcome It

Carl Bereiter

Centre for Applied Cognitive Science

The Ontario Institute for Studies in Education

The first half-century of American psychology is often represented nowadays as a dark age of rat-running and rote verbal learning experiments, which ended only when it was overwhelmed by the rise of the new cognitive psychology that started in the 1950s (cf. Gardner, 1985). But throughout that early period there was a strain of mainstream experimental psychology that dealt with what we would now call situated learning. It was rat psychology, to be sure, but it was a psychology of situated rat behavior. The central tenet was that animals do not simply learn responses, they learn their environments. Run rats in a maze under typical tightly controlled conditions and they learn a fixed route to the goal. Change things a little and they are lost. But let them run around on their own and they quickly learn the whole maze, so that they can get from wherever you drop them to wherever they want to go by an efficient route. Edward C. Tolman was responsible for the seminal experiments and also for the proposal that ‘place learning,’ as he called it, is mainly perceptual and not behavioral (Tolman, 1949). It was Robert S. Woodworth, however, for some time the dean of American experimental psychologists, who made the most forceful effort to shift the mainstream of experimental psychology toward a consideration of organisms interacting with their environments. “If behavior consists in an active dealing with the environment,” he said (Woodworth, 1958, p. 300), “and if it depends very largely on learning, it must depend on what we can properly call ‘learning the environment.’” He then went on to elaborate:

If we said, “learning about the environment,” we should imply the sort of information that can be picked in conversation and reading. What we mean is a direct acquaintance with persons, places, and things—the kind of learning that occurs when we see a person, hear him talk, watch him act, and participate with him in social activity.... (Woodworth, 1958, p. 300)

I would like to credit Woodworth with having arrived at the basic idea of situated cognition by the unlikely route of animal experimental psychology. Despite his eminence, his 1958 book seems to have been largely ignored. Interesting observations, but they don't lead anywhere theoretically—that was the opinion expressed by professors of mine at the time. Woodworth was a couple of decades too early. His ideas would find a ready reception now.

Woodworth and his environment-learning rats are of more than historical interest, however. Contemporary ideas about situated cognition, having come to us from anthropology, are heavily loaded with human cultural concerns. They are connected with Vygotsky and his belief in the social origins of cognitive structures. As several of the chapters in this volume indicate, situated cognition has also gotten itself connected with feminism and educational radicalism. We tend to forget that animal cognition is situated as well.

In order to consider such an idea as ‘overcoming’ situated cognition, it is important to appreciate that, no matter how deeply enmeshed in culture human cognition may have become, its situatedness has a biological basis. Rat cognition is situated in a world of mazelike tunnels formed by the framing of building walls and floors. Learning such environments is what rat learning is geared to. Rat intelligence involves ‘leveraging off’ the structural properties of such environments for purposes of escape, gaining entry, food gathering, and so on. Of course, rat cognition is situated in a world that also contains other rats—but it is a stretch to say that it originates on a social plane. Our brain, like the rat's, evolved as an adaptation to our environment. The environment, of course, changes, and many of the most striking changes have been brought about through our own doing. But

those changes have occurred too recently to have had much effect on our own evolution, which for most of its period took place in a world little affected by our presence.

Outgrowing Animal Cognition

Thus, we do not come by situated cognition through a cultural or learning process. Our brains evolved to deal with situations in which we find ourselves. We have, however, managed to transcend our animal heritage in certain ways, and in this chapter I argue for the value of viewing these as ways of overcoming the situatedness of cognition. Like other adventures in 'overcoming' nature, overcoming the situatedness of cognition has risks as well benefits and is, in a fundamental sense, illusory. But identifying the risks and benefits and separating illusion from reality is part of the program I am advocating here.

The most obvious way in which we humans transcend our animal heritage is through transforming physical environments and creating new social structures and practices along with them. The second way is through acquiring expertise, which enables us to function in a novel environment much as if we had evolved within it (Bereiter & Scardamalia, 1993).

Environmental transformation and expertise are to be found in every society; but there is a third way of transcending biological givens that is much less common and that represents a far more radical departure from the kinds of cognitive adaptations we share with other species. This is the kind of departure most dramatically exemplified by science. It amounts to creating a world of immaterial knowledge objects and acquiring expertise in working with them. Although these knowledge objects may refer to spatially and temporally located situations, they are not bound to those situations. Thus, this third way represents a stronger sense in which humans may be said to overcome the situatedness of cognition. The third way, furthermore, greatly extends the other two. A modern city is a physical environment within which human beings have developed many new forms of practice and expertise, but this environment could not exist were it not for centuries of development of abstract knowledge now put to use in the construction of tall buildings, electrical power

grids, heating and air-conditioning systems, and the countless other technological underpinnings of a modern city.

These three kinds of advances beyond animal cognition map nicely on to Karl Popper's metaphoric schema of three worlds—World 1 being the material world of inanimate and animate things (including human beings), World 2 the subjective world of individual mental life, and World 3 the world of immaterial knowledge objects (Popper, 1972). Lacking other handy labels, I will use Popper's terms, without implying a necessary commitment to other aspects of Popper's epistemology.

Situated cognition researchers have contributed substantially to our understanding of the relations between Worlds 1 and 2, arguing that these are much more directly and intimately connected than previous cognitive theories had supposed. But they have not done the same for World 3. Instead of according knowledge existential status in its own right, as epistemologists have traditionally done, they have tried to account for it in terms of (a) the practices of particular groups, such as scientists or mathematicians; (b) concrete embodiments of knowledge, such as books and instruments; and, occasionally, (c) as content in individual minds—as mental models, for instance (Greeno, 1994).

Although these are important aspects of knowledge, they seem to me to miss the core. That core is represented in the metaphor of World 3—a world, wholly created by the human intellect, which enables us, for better or worse, to escape the situational embeddedness of cognition. Without that core, formal education becomes meaningless (as, indeed, some advocates of situated cognition seem to believe it is). Again, for better or worse, formal education is our individual escape route from the confines of situated cognition.

Can There Be Nonsituated Cognition?

James G. Greeno, a leading exponent of situated cognition, has expressed dissatisfaction with the term. It seems to refer to a type of cognition and thus to imply that there also exists a type of cognition that is not situated (Greeno, 1994). Situativity theorists (to use Greeno's suggested replacement term) deny this: All cognition is situated and

couldn't be any other way. If we accept this premise, then the title of this chapter, "Situated Cognition and How to Overcome It," is an oxymoron.

But, in fact, there is nonsituated cognition, and situativity theorists have devoted a lot of effort to criticizing it. The catch is that nonsituated cognition is not found in nature (at least not in nature as it is known to earthlings), it is found only in machines. Most of artificial intelligence has been constructed according to a model radically at variance with the kind suggested by situativity theory. In situated cognition, people (or other agents) carry on activity in the world, adapted to the constraints and affordances of the environment.

Cognition is the individual or collective process by which people negotiate these constraints and affordances, according to their individual or collective purposes. Machine intelligence of the classic AI variety is not like that. Cognition is an entirely internal process of symbol manipulation (Vera & Simon, 1993). Interaction with the outside world is done by means of transducers that translate inputs from sensors into symbols that the machine can manipulate or that translate symbols into actions. Thus, a robot controlled by AI of this kind will contain a plan—for getting from point A to point B, let us say—which controls how data from its visual sensors are translated into symbols that its program can then convert into instructions to its servomechanisms so that the plan is executed through physical movement. The robot may also contain a program for revising the plan, in case of mishap. One very important line of argument in favor of situated cognition comes from roboticists, who find that this kind of robot cognition does not work very well (Beer, 1991). It is too slow and crude and prone to catastrophic failure. But these kinds of criticisms acknowledge that there is such a thing as nonsituated cognition, making the case that it is not a very good kind of cognition for getting around in the world.

The existence of nonsituated cognition, albeit artificially created, is, I believe, profoundly important for understanding human cognition and its situatedness. For one thing, it allows us to talk about advantages and disadvantages of situated cognition, which would make no sense if there were nothing to compare situated cognition to. It also affords

the possibility of identifying degrees of situatedness. With such possibilities in view, it is no longer absurd to talk about ‘overcoming’ situated cognition. The questions are, why would anyone want to and how could it be done?

Advantages and Disadvantages of Situated Cognition

If we take rule-based AI¹ as exemplifying nonsituated cognition, then looking at what it does well and poorly (compared to human beings) may offer us insights into the advantages and disadvantages of situated cognition. Using rule-based programs, computers are much better than we are at carrying out long chains of reasoning and at exhaustively searching memory (Anderson, 1985). Thus, they excel at chess. The best programs, which can beat all but a few human experts, succeed by reasoning farther ahead and along more paths than their human opponents (Charness, 1991). We are abysmally bad at searching memory in a listwise fashion. Try naming the 50 United States—or some other familiar set of about that size, if the states are not familiar enough. Almost everyone misses one or two and has a terrible time finding the missing ones, whereas computers have no trouble with this sort of task.

What we do remarkably well in comparison to rule-based AI is recognize patterns—for instance, recognizing a face from the past, even though it has aged 20 years since last we saw it. Computers, by contrast, have trouble identifying letters of the alphabet under the normal variations of handwriting and typography, a task many preschool children can handle easily. Another relative strength of human cognition is associative retrieval—for instance, reading a research report and being reminded of a related finding from a decade past, on a slightly different topic. This is a chancy business for us, but rule-based AI cannot do it at all unless the stored items have been appropriately indexed beforehand (Schank, Collins, & Hunter, 1986). Pattern recognition and associative retrieval seem to be the means

¹ The qualifier, ‘rule-based,’ is necessary because there has lately emerged a different kind of AI, connectionism (1989), which functions much more like situated cognition. Throughout this chapter, however, when I refer to AI I always mean the rule-based kind, whether the qualifier is attached or not.

by which we grasp analogies and metaphors, and this gives us a great imaginative edge over the literal-minded machine (Margolis, 1987).

The relative strengths of computer and human cognition directly reflect differences between nonsituated and situated cognition. Rule-based AI works very well when all the necessary information can be explicitly represented and indexed, as is the case with a game of chess or a gazeteer. The rule-based system can then go to work on its stored information and produce a result appropriate to the part of the real situation of which it contains a representation; e.g., it can compute a move appropriate to a real chess game. When all the necessary information is coded into rules or propositions, formal logic comes to the fore as a powerful tool for arriving at decisions, and logical operations are what computers excel at. The trouble is that the great bulk of real-world situations cannot be represented in this way. Chess games can, but as simple a game as tag cannot. This is because chess has a set of rules that allow all the possible moves to be computed, whereas the possibilities inherent in a dozen kids running around on a playground are for practical purposes limitless. Representations are necessarily abstractions. Abstractions based on Newton's laws work for physical situations involving a small number of inanimate objects, but when objects have minds of their own and can twist and dodge the way agile children can, such abstract representations become relatively useless. It is not that Newton's laws cease to hold, of course; it is just that the variables are too numerous and are impractical to measure and compute.

A tag-playing robot would need a different kind of mind from the one rule-based AI would give it. Instead of a mind that works on internal representations of the playground and the participants, it would need a mind more directly attuned to the physical and social environment, responding quickly to opportunities for tagging, switching from one pursuit to another the instant that a more promising target presented itself. In short, it would need situated cognition.

At this point, my account sounds like much of what appears in the situated cognition literature. The situated actions of just plain folks come off as flexible, adaptive, and elegant—in a word, intelligent—whereas action based on formal procedures and principles comes off as brittle, plodding, insensitive to nuance—in a word, stupid. It is time, therefore, to look at the other side. Although nonsituated cognition may not be very good for guiding a robot in a game of tag, it has proved capable of guiding a space vehicle to Mars. Surely there is a lesson for us in that.

The Problem of Transfer

The main weakness of situated cognition is, it seems, precisely its situatedness. In traditional language, the limitations of situatedness are referred to as problems of transfer. What we learn in one situation we often fail to apply in another. Situativity theory helps us to understand why this is so. The progress of situated learning consists of increasingly fine attunement to the constraints and affordances of the particular situation. Thus, as learning proceeds it tends to become less and less generalizable to other situations. In your first job as a store clerk, you will begin by learning many things that are applicable to store clerking in general—how to address customers, ring up sales, bag purchases, watch for shoplifters, and so on. But as the weeks go on, your skills will become more and more specific to the particular store and its merchandise, clientele, management, physical layout, staff, and so on.

Advanced stages of situated learning may, in fact, begin to yield negative transfer, as habits are acquired that will need to be overcome in a new situation. There is a deeper problem of transfer, however. Elsewhere I have tried to show that what mainly fails to transfer is learned intelligent behavior (Bereiter, 1995). The course of situated learning typically has the aspect of a progression from being inept and prone to stupid mistakes to being competent and smart. Although important parts of what is learned in one situation may transfer to a new one, the part that does not transfer is likely to include the being smart. Again, this makes sense in light of situativity theory, for being smart just means becoming so nicely attuned to specific constraints and affordances of a situation that you can

effortlessly cope with whatever problems arise.² In a new situation, you are liable to have to start over being stupid.

If, categorically, learned intelligent behavior is not transferrable from old situations to new, this has grave implications for education, if not for humankind in general, what with the accelerating pace of change. But there we have an irony worth pondering. The ‘accelerating pace of change’ is increasingly driven by technological innovations, virtually every one of which is an instance of transfer of intelligent behavior from one situation to another. I do not want to make space travel out to be the highest achievement of human intelligence, but it is surely our most colossal example of transfer of learning. No amount of situated cognition or ‘legitimate peripheral participation’ would get people to the moon and back.³ It took something more to produce that kind of transfer, and we must try to pin down what that is. Failing in that, we may face a future in which a small number of people have caught on to some secret of transferrable learning and thus are able to keep creating and adapting to new situations, while the rest of us find it increasingly difficult to cope.

Greeno, Smith, and Moore (1993) have offered a situated account of transfer, attributing it to constraints or affordances that are common across situations. This account differs from older, nonsituated accounts, in that constraints and affordances are not characteristics of either the environment or of the person, considered separately, but of the relationship between person and environment. Thus, transfer is a matter of the same kind of relationship coming into play in different situations. I think that is a valuable conceptual advance on this timeworn problem, but I also believe that it raises issues that can only be dealt with by introducing ideas of nonsituated cognition.

² This is how Dreyfus and Dreyfus (1986) characterize expertise—a characterization that we criticize as failing to distinguish experts from experienced nonexperts (Bereiter & Scardamalia, 1993).

³ This is not to deny that situated learning played a part. Astronauts did, after all, do part of their training near Sudbury, Ontario, where years of mineral extraction had produced a terrain closely approximating what was expected to be found on the moon.

To constitute an interesting case of transfer, the constraints and affordances of two situations must appear different on the surface. In one of the classic transfer problems, people learn a clever way of beaming x-rays at a tumor and then are given the opportunity to apply this knowledge to figuring out a clever way for Crusaders to attack a castle. Without prompting, people seldom exhibit transfer; but if told that the solution to the first problem should suggest a solution to the second, more people will make the connection. Thus, transfer in such cases is anything but automatic. People have to be looking for a relationship. And what kind of relationship is it? The word that comes to mind is ‘abstract.’ It is a relationship based on formal, structural, or logical correspondences. In the tumor and castle example, the abstract idea is focus—things following converging paths to produce maximum effect where they meet.⁴ To discover such abstract relationships, however, one has to carry out something quite different from ordinary situated action. One has to create symbolic representations of situations and carry out operations on those symbols. In other words, one has to act more like a computer and less like a creature of nature.

Learning Beyond What the Situation Calls For

Flora and Dora are both ‘A’ students in Algebra I. The next year they take Algebra II. Flora again acs the course, while Dora finds herself at a loss and just manages to scrape by with the ‘C minus’ awarded to students who try hard but don’t get it. Here we have an apparent case of the same learning transferring for one person but not for another. Few mathematics educators would buy that interpretation, however. They would conjecture that Flora and Dora learned quite different things in Algebra I. What Flora learned evidently provided a good basis for Algebra II, whereas what Dora learned did not. So it is not that something failed to transfer for Dora, it is that she failed to learn what was transferrable. How are we to account for these different learnings, and how do they relate to situated cognition?

⁴ Thus, converging x-rays destroy the tumor but have less effect on the surrounding tissue; converging tunnels enable large numbers of troops to reach the castle simultaneously.

In the last 20 years there has been quite a bit of research on the Floras and Doras of the world. Even without the research, it is easy to divine that Flora probably understood the mathematics presented in Algebra I and that Dora didn't, getting by instead on rote procedures. What the research has done is give us an idea of what Flora did differently from Dora in order to produce this result. Stepping down from algebra to arithmetic, a nice example of this research comes from Resnick and Neches (1984). They examined a practice common in elementary arithmetic, in which children carry out operations with concrete objects that mirror such symbolic operations as borrowing and carrying (or regrouping, as it is now called). Most children, although they were able to carry out both the concrete and the symbolic operations, failed to make a connection between them. Some children did make the connection, however. On interviewing the children, it was found that the children who made the connection reported that they were trying to make a connection.

That may not sound very surprising, but consider it from the standpoint of situated cognition. Clearly, the situation was designed to afford experimenting with the relationships between concrete and symbolic quantitative operations. But any situation affords innumerable opportunities for inquiry. Why we would exploit one and not another—or any at all, for that matter—depends partly on our own goals and partly on, to put it broadly, 'what the situation calls for.' Generically, the situation in which the children in Resnick and Neches's study found themselves was that of schoolwork (Doyle, 1983). A schoolwork situation is rather like that in a garment factory. Although a number of workers may be present in the same room, each one is independently engaged in carrying out a task specified by the supervisor or teacher. What the situation calls for is defined by the task constraints. The tasks are usually defined in such a way as not to put undue strain on the capacities of those performing them. There usually are time constraints, however, and so there is motivation to find ways of satisfying task requirements that economize on time. Another characteristic shared by schoolwork and garment work is a very limited time horizon. Although the teacher or supervisor may have long-range objectives, the students or

workers are not expected to look beyond the immediate day's task. (School 'projects' are an exception, as their name implies.)

Given these characteristics of schoolwork, it then becomes remarkable that some children would take it upon themselves to try to discover a logical connection between the concrete and symbolic components of the task they were assigned to carry out. The task assignment did not require it. The task components were easily enough executed that it was unlikely that an impasse would drive them to deeper analysis. And the time horizon, suggesting that it would all be over soon, offered no reason to think ahead to the possible relevance of the current task to future situations.

Returning to Flora and Dora, their first-year algebra class probably provided opportunities—through textbook explanations and worked examples, class discussions, and problems—to develop a basic understanding of algebraic functions. But what the situation actually called for was just solving lots of linear equations. By learning a few procedures and applying them carefully, an assiduous student could solve the equations without any need for conceptual understanding. That, we may surmise, is what Dora did. It worked well through Algebra I. But then she got to Algebra II and encountered an explosion of different types of equations and complications in procedures for solving them. Try as she might, she made frequent errors. Having no sense of how algebra related to arithmetic, she never checked her answers with trial values. (Perhaps arithmetic did not make much sense to her, either, and therefore provided no basis on which to build an understanding of algebra.) Consequently her errors went uncorrected and the marks on her schoolwork plummeted.

Flora, we surmise, did acquire an understanding of algebra in the course of her first year. But on the basis of related research, we may further surmise that this did not just happen. Despite the fact that the situation did not actually call for it, Flora must have expended effort in trying to understand what algebra was about and how it connected with what she already knew. This is what we have elsewhere defined as intentional learning (Bereiter & Scardamalia, 1989). Intentional learning is primarily a matter of goals rather

than strategies. Examining the goal-related statements of people studying computer programming, Ng and Bereiter (1991) were able to identify three levels of goals. The first and most common are task completion goals. In the Flora and Dora case, these would be goals associated with correctly completing assigned algebra problems. At the next level are instructional goals. These are goals related to what the teacher or textbook is trying to teach. They can vary greatly in how explicitly and saliently they are put before the student. In a typical algebra textbook they would be discernible from section headings and the like, but they could easily be ignored in the pursuit of task completion goals. Finally, and rather rare, are knowledge-building goals, which pertain to the learner's personal agenda for constructing knowledge. Among other things, the three kinds of goals differ greatly in their time horizons. In a conventional algebra class, the time horizon for task completion goals is usually the next day. The time horizon for instructional goals is likely to be the next examination, or at most the end of the course. The time horizon for knowledge-building goals, by contrast, may extend indefinitely far into the future, and may also extend into the past, encompassing a history of past learning that is consciously built upon in the present.

To succeed in Algebra II, the student must at least have achieved some of the instructional goals implicit in Algebra I. Merely achieving task completion goals, even doing very well at them, would not suffice. To be on track for becoming a mathematician, it is probably also necessary to be pursuing knowledge-building goals—goals that extend beyond the instructional goals of the immediate course and that involve, in effect, reconstructing mathematical knowledge in one's own way and following out its implications (cf. Popper & Eccles, 1977, p. 461).

These three levels of goals differ in their level of abstractness. Correspondingly, we may say that they differ in their degree of situatedness. Action in pursuit of task completion goals is highly situated, being directly linked to manifest constraints of the situation. The pursuit of instructional goals is less so. And when we get to pursuing knowledge-building goals, we are talking about action that is only weakly connected with the immediate situation,

that consists largely of mental work on symbolic objects, some of which are abstracted from the current situation but others of which originate quite outside it.

Situativity theorists might concede that something like this continuum of abstraction exists, but they would argue that the more abstract kind of mathematical activity is just as situated as the more concrete. This is where the relational character of situativity becomes important. Although Flora and Dora may be in the same physical environment, they are in different situations. The affordances and constraints are different, reflecting their differing motives and capacities. But even though one may be plodding through a workbook assignment while the other is reflecting on the nature of mathematical functions, each is engaged in a cultural practice that is adapted to situational constraints and affordances.

I do not want to argue the contrary. With a concept as elastic as situatedness, it would be impossible to exclude any activity from it categorically. What I do want to argue for, nevertheless, is the value of thinking of situatedness as varying along a continuum. Both ends of the continuum are limits that might be approached but could never be reached. At the extreme situated end would be organisms so fully and rigidly adapted to one particular environment that they could never survive in any other. The organisms that live in the steaming pool at the base of Old Faithful come to mind as possible candidates for this end of the continuum, but no doubt even they have some range. To mark the other end of the continuum—to make the other end of the continuum even imaginable—we must invoke rule-based artificial intelligence. With such a continuum in mind, we can grant that Flora and Dora, along with all animal organisms, demonstrate situated cognition, but at the same time we can make something of the idea that in their mathematical behavior Flora is farther toward the nonsituated end than Dora. In the next sections I will try to show what can be made of this idea.

Learning to Simulate a Computer

The pioneers of artificial intelligence did not suppose that they were creating something to stand at the opposite pole from natural cognition. On the contrary, they thought, and

many still maintain, that the computational systems they created were reasonable simulations of human intelligence. These systems were, after all, based on what common sense and 2000 years of philosophy told us is the nature of human intelligence—that it consists of the ability to reason and figure things out on the basis of information coming from the outer world. AI systems did that, and did it with kinds of tasks that challenge human reasoning abilities: playing chess and other intellectual games, solving logical and mathematical puzzles, diagnosing illnesses. Moreover, fine-grain comparisons between machine and human behavior indicated that machines tended to carry out their tasks in human-like ways, encountering human-like difficulties, and even following time courses similar to those of human reasoners (e.g., Anderson, 1983; Newell & Simon, 1972).

How could it have come about, as situativity theorists suggest, that these AI pioneers were all barking up the wrong tree, designing something radically at variance with what natural cognition is really like? To answer that question, we had better open it up and ask as well how common sense and 2000 years' worth of philosophers could have made the same mistake. The answer, I think, lies in introspection.

We are conscious of only a small part of our cognitive activity. In routine activities, such as driving a car or reading a newspaper, we process vast amounts of information, but we tend to be conscious of it only when some mental effort is involved—when a problem arises or when something comes up that 'makes us think.' Thus, when philosophers or just plain folks reflect on the nature of cognition, their introspective database consists of these rather exceptional events. This is like doing an ethnography of a people based solely on observing riots and demonstrations.

During these exceptional cognitive events, we find ourselves formulating problems and hypotheses, recalling rules and facts, and engaging in spurts of if-then reasoning. The great fallacy of traditional theory of mind has been to assume that the rest of cognition consists of these same sorts of things, carried out unconsciously (Dreyfus & Dreyfus, 1988). A more realistic and productive view, which situativity theorists have helped advance, is that these

symbolic processes are themselves exceptional and by no means representative of the great bulk of cognitive activity. Sitativity theorists have helped us to appreciate the great unconscious part of cognition.⁵ My purpose here is to look again at the exceptional conscious part, to reflect on what it is for and how it relates to unconscious situated cognition.

The following is no doubt an overstatement, but it will serve to launch the discussion: The mental processes that we are conscious of are processes that do not come naturally (which probably has much to do with why they are conscious). They are acquired processes, culturally mediated, which enable the brain to act as if it were a different kind of device from what it had evolved to be. The device we simulate in our conscious thinking is a logic machine—a machine which finally, in the current century, people have been able actually to build. Thus, according to this conjecture, the computer does not simulate human cognition. Instead, human cognition is able to simulate the computer—although, until this century, the computer existed only as an idealization of a certain form of intelligence and not as a physical reality.

According to Kitto (1954), the ancient Greeks believed they had discovered logical reasoning, and they treated it as a sort of information technology to be exploited and disseminated. According to Piaget (Inhelder & Piaget, 1958), contemporary adolescents often have the same notion when they first become able to carry out formal logical operations. They think they can now figure out everything. Such enthusiasms are consistent with a cognitive achievement that, at least subjectively, brings new powers. They are also, however, a signal to us to be on guard against overextended claims.

⁵ Although Freud is justly credited with winning a place for the unconscious in psychology, he held to the classic view that unconscious structures and processes are similar to conscious ones. It is the *content* of the unconscious that Freud treated as distinct, as well as the processes by which certain content is rendered unconscious.

Young children and tribal peoples unfamiliar with the syllogism have often been regarded as lacking in reason, and much of modern research in anthropological and developmental psychology has gone into demonstrating that this is not the case. Logicality, it turns out, is to be found in all peoples and at very young ages, provided one looks for it in situations where the people have had a chance to develop some knowledge and competence (Donaldson, 1978; Scribner, 1979). So what is it, then, that the Greeks thought they had discovered and that adolescents keep rediscovering? Piaget seems to have had it basically right. It is one thing to think logically about concrete reality; it is something else to think logically about propositions (which may or may not refer to concrete reality).

In a famous study that was originally taken to demonstrate lack of logical reasoning in uneducated people, Luria (1976) presented Russian peasants with syllogistic reasoning problems such as the following:

In the Far North, where there is snow, all bears are white. Novaya Zemlya is in the Far North. What color are bears there?

Luria found that the peasants tended to waffle on such problems, protesting that they were really not all that familiar with conditions in Novaya Zemlya and advising Luria to go check with someone who had been there. By contrast, collective farm workers who had had some education promptly responded that the bears in Novaya Zemlya are white.

Without speculating too deeply about thought processes, we can say with confidence that the more educated workers did not actually know that the bears in Novaya Zemlya were white. What they knew, rather, was that the two given propositions imply that the bears in Novaya Zemlya are white. Either proposition could be false or there could be no such place as Novaya Zemlya or there might be no bears there at all. These may have been some of the considerations that bothered the peasants, for they are legitimate concerns when it comes to thinking about the real physical world—Popper's World 1. But in the kind of language game that we call logical reasoning, none of these considerations count. Certain propositions are taken as given, and at issue are propositions that follow from them. This is

action carried out in World 3, which is a world that was probably unknown to the illiterate peasants.

By itself, World 3 is just a world of language games carried out by what might as well be arbitrary rules. The great power of World 3, the thing that makes it exciting to people when they discover it, is that when one works with premises that are true with respect to World 1, then valid conclusions from these premises turn out also to be true with respect to World 1. Conversely, if valid conclusions turn out not to be true with respect to World 1, this is a signal that something is wrong with the premises. Disciplined movement back and forth between World 1 and World 3 gives us the hypothetico-deductive method and opens up vast possibilities for theory development, problem solving, and design.

Such heady talk is, of course, the basis of positivism, and it is subject to all the criticism that has been leveled against that overly optimistic philosophy. It furthermore overlooks the large and vital role of informal, unarticulated, impressionistic, and embedded knowledge (Bereiter & Scardamalia, 1993). This is not the place for a critique of pure reason, however. The point I am trying to make is that a significant and empowering change takes place when one catches on to treating ideas as objects of inquiry. Once we get good at it, it seems perfectly natural and we tend to forget that there was ever a time when we did not do so. But it is always effortful, open to serious error. It is not how we function most of the time, and it may well be that a large part of the human population never functions that way.

In human development, several theorists have characterized cognitive development along similar lines. Donaldson (1978), after documenting the logicity of young children's behavior, proposes that during subsequent cognitive development this logicity becomes less situation-dependent—thus, in some sense, less situated. Karmiloff-Smith (1992) has proposed a theory even more closely aligned with the present argument. She portrays cognition as going through a series of “representational redescription,” such that knowledge that was originally implicit in skills finally comes to be rendered sufficiently

explicit that it can be applied in different domains and in the service of goals different from those that gave rise to it.

Schooling and Knowledge Work

Although situated cognition researchers have taken a lively interest in learning, both in and out of school, they have not come up with anything that could be called a new educational vision. Instead, situativity theorists have tended to endorse various innovations of a social constructivist cast, interpreting them within their own frameworks. As has been pointed out (Wineburg, 1989) and acknowledged (Brown, Collins, & Duguid, 1989), however, the educational ideas coming from situativity theorists have not advanced notably beyond those of John Dewey. The main difficulty, I would suggest, is that situativity theory has not been able to provide a cogent idea of the point of schooling. This difficulty, in turn, derives from a serious confusion between product and process.

If the work of a community of practice is manufacturing paint, for instance, there is no difficulty in separating the process from the product. One may view the manufacturing of paint as situated activity without also having to regard the paint as situated. The paint will be used in all kinds of remote and unknown situations. The fact that the paint may also be used by the people who make it—to paint their shop walls, for instance—still introduces no confusion. But when the product of an activity is knowledge, and the knowledge is mainly of use to the people who produce it, confusion can be well nigh total.

The source of the confusion is that knowledge production, like any kind of human activity, takes place in some physical and social situation, and accordingly situated knowledge also develops. This is knowledge constituted in the practice of the community and intimately involved with the affordances and constraints of the situation. But this is not the same as the knowledge that is the *product* of the situated activity, anymore than the situated knowledge of the workers in the paint factory is the same as the paint they produce.

Some knowledge-production situations are less confusing than others, however, so let us consider one of those first. A forensic chemistry laboratory produces knowledge of a

particular kind. Through analysis of materials obtained at a crime scene, knowledge is produced that contributes to creating an account of what went on at the scene. In this case, it is not difficult to distinguish between the knowledge embedded in the practice of the chemists and the knowledge that they deliver to the detectives. The two kinds of knowledge relate to entirely different situations. The distinction becomes trickier if the chemists are doing basic research. In this case, the knowledge they produce relates to their own practice as well as to others'. But with a little effort the distinction can be maintained. It is not so much different from the paint-makers being users of the paint they produce. If, however, chemical research is being carried out by students in a school laboratory, then the distinction becomes even less obvious. This is because the students are likely to be the only users of the knowledge they produce. Nevertheless, I believe that the school situation, like the other situations in which knowledge is created, can best be understood by striving to distinguish knowledge implicit in the process from knowledge that is the product of the process.

No such distinction is normally made in education, even with the popularization of constructivist ideas. The results of knowledge construction are thought of as entirely internal—internal to the minds of individual students under most construals, or internal to the distributed cognition of the classroom community under construals influenced by sociocultural theories (Cobb, 1994). Accordingly, constructivism becomes more or less synonymous with learning by discovery, and it competes—not always successfully—with direct instruction (Harris & Graham, 1994).

But students can produce knowledge objects—theories (or theorylike conjectures, at any rate), interpretations, historical accounts, problem statements, defenses based on evidence, and so on. These may be embodied in reports or presentations, but not necessarily. When students use a networked computer database as a discourse medium, they can produce quite substantial knowledge objects without the need for any more tangible product than the electronic record of the discourse itself (Scardamalia, Bereiter, Cassells & Hewitt, in press; Scardamalia, Bereiter, Lamon, 1994).

The observable goings-on in this activity that we call ‘collaborative knowledge building’ fall easily within the spectrum of what others might call ‘constructivist learning,’ ‘cognitive apprenticeship,’ ‘inquiry learning,’ or ‘talking science.’ The distinctiveness is conceptual; it is a matter of how teachers and students conceive of what they are doing and the effect this has on efforts to do it better. One thing that must be recognized about the many exciting experiments in educational uplift that are going on (it is true of all the ones I have knowledge of, and so I confidently generalize to the rest) is that reality falls well short both of the ideal and of the exemplary episodes reported in the literature. Hence, in pedagogy as in science, improbability is of the essence. The following are ways in which a knowledge-building conceptualization of schooling offers advantages over other approaches that regard both knowledge construction and the knowledge produced as situated:

1. The focus of classroom activity shifts from improving students’ minds to improving their theories or other knowledge objects. This is a clearer objective and one that students and teachers alike can more readily track.

2. A developmental continuum may be recognized that runs from unconscious learning in early childhood (Montessori, 1967) to self-aware, intentional learning (Bereiter & Scardamalia, 1989) and then to inquiry that is focused on the external world and finally to inquiry that is focused on World 3 objects as they relate both to the external world and to one’s own purposes (Scardamalia, Bereiter, & Lamon, 1994). Helping students advance along this continuum then becomes a meaningful educational objective.

3. Production of knowledge objects inevitably involves building on or otherwise dealing with existing knowledge objects (hence Newton’s avowal that he stood upon the shoulders of giants). Consequently, familiarity with culturally significant World 3 objects—the goal of ‘cultural literacy’ (Hirsch, 1987)—comes about naturally rather than through a didactic regimen.

4. The problem of ‘inert knowledge’ (Whitehead, 1929) is effectively circumvented. Progressive education sought to avoid inert knowledge by having learning come about

naturally through the social life of the community. But the social life of school communities does not naturally give rise to much learning of an abstract or theoretical nature. The most immediate and obvious use of knowledge objects is in creating new ones—in creating new understanding either of particular phenomena or of a class of phenomena. Students experience the power of concepts in science and other disciplines by using them to help solve problems in their own knowledge-building efforts.

5. Knowledge building is not in competition with instruction. In real life, people occasionally take time out from their work to learn something—often learning something that will help them in their work. Professional associations often include training sessions and tutorials in their annual meetings. Similarly, in schools there is no reason why time cannot be taken out from knowledge building, to whatever extent is judged necessary, and devoted to explicit learning activities. The more immediately relevant these are to students' knowledge building the better, of course, but there is no reason why they have to be carried out in a way that is ideologically consistent with the school's approach to knowledge building (Bereiter & Scardamalia, in press).

6. Knowledge building provides a natural basis for involving people outside the school who are engaged in related activities—scientists, curators and librarians, experts in various trades and professions, and so on. Bringing such people into learning activities is problematic because they are too peripheral to the curriculum to have a good sense of what role they should take and they may furthermore lack pedagogical skills that the role requires. But contributing to other people's research is something they probably already do. This involvement of talents beyond the classroom is almost obligatory if students are free to follow a knowledge-building project wherever it leads them (Scardamalia & Bereiter, 1994).

6. The knowledge objects students produce in school will tend naturally to be ones of very basic and general applicability. This is because (a) there is no particular job that the knowledge must serve (as there is, for instance, in the forensic chemistry laboratory) and (b) students' questions, when freely generated, tend to be 'why' questions that lead toward

deep principles (Scardamalia & Bereiter, 1992). Consequently, the knowledge that students produce is the kind that serves broadly to overcome the limits of situated cognition.

7. The situated learning that does occur is learning how to function in a community of practice whose work is work with knowledge. The transferrability of this learning to 'knowledge work' in out-of-school situations is, of course, chancy; but it seems reasonable to assume that students who have had years of experience in explicitly working with knowledge will have an advantage over ones whose experience has been limited to the traditional kinds of scholastic learning and doing in which knowledge, as such, is seldom the object of attention.

This last point warrants elaboration, in view of the much-heralded movement toward a knowledge-based society (Drucker, 1994)—a society in which the main wealth-producing work is knowledge work. If the term has any distinctive meaning, it must imply work that is focused on knowledge objects themselves. All work, even that conventionally categorized as unskilled, involves a great deal of knowledge (Vallas, 1990). What distinguishes knowledge work is not using knowledge but creating or adding value to it.

Of course, knowledge work is situated activity and has all the characteristics of other situated activity, including tool use and leveraging off physical affordances of the environment; but it is distinctive in two respects:

1. The knowledge that is being worked with is not situated knowledge. It is knowledge that has been transformed into objects that can be treated or used in an unlimited variety of situations. Thus, knowledge is no longer bound to the situations in which it was constituted.
2. In order to work effectively with knowledge objects, people have to master the practices of nonsituated cognition. This means learning to carry out the sorts of unnatural cognitive actions performed by logic machines. This does not mean becoming less human; it means acquiring a special set of skills to use wisely or

unwisely, imaginatively or ploddingly, as we do with the many other intellectual, practical, and social skills that constitute human competence.

These two characteristics mark a cultural divide. A major social issue for our time is whether the world will be run by an expert elite on one side of the divide, while the bulk of humanity remains on the other. It seems to me that today's schools are on the wrong side of the divide. That bodes ill for prospects of moving much of the population to the postindustrial side.

One of the most disturbing indicators that I encounter comes from my experiences in speaking publicly about the ideas discussed in this section. People in modern businesses understand what I am talking about immediately. Educators usually do not. They think I am just talking about active learning. Educators are immersed in World 3, but they are like the proverbial fish immersed in water. They cannot see it. They do not conceive of knowledge as something that can be manufactured, modified, worked with, and in some cases even packaged and sold. Unfortunately, the rise of situated cognition theory does not help in this regard. It has contributed greatly to our understanding of the kind of knowledge that is implicit in practice, but by treating all knowledge as situated it renders the world of knowledge objects invisible.

In a famous statement, Sir Isaac Newton likened himself to a child finding pretty stones on the shore "whilst the great ocean of truth lay all undiscovered before me." Those pretty stones, however, were the foundation of the modern world. We need schools in which students learn to work with pretty stones. As for the "great ocean of truth," all we can say with confidence is that, if it ever is discovered, it will not be by fishes.

References

- Anderson, J. R. (1983). The architecture of cognition. Cambridge, MA: Harvard University Press.
- Anderson, J. R. (1985). Cognitive psychology and its implications (2nd ed.). San Francisco, CA: W. H. Freeman.
- Beer, R. D. (1991). Intelligence as adaptive behavior. Cambridge, MA: MIT Press.
- Bereiter, C. (1995). A dispositional view of transfer. In A. McKeough, J. L. Lupart, & A. Marini (Eds.), Teaching for transfer: Fostering generalization in learning. Hillsdale, NJ: Erlbaum.
- Bereiter, C., & Scardamalia, M. (1989). Intentional learning as a goal of instruction. In L. B. Resnick (Eds.), Knowing, learning, and instruction: Essays in honor of Robert Glaser (pp. 361-392). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Bereiter, C., & Scardamalia, M. (1993). Surpassing ourselves: An inquiry into the nature and implications of expertise. La Salle, IL: Open Court.
- Bereiter, C., & Scardamalia, M. (in press). Rethinking learning. In D. Olson & N. Torrance (Eds.), Handbook of education and human development: New models of learning, teaching and schooling. Cambridge, MA: Basil Blackwell.
- Bereiter, C., Scardamalia, M., Cassells, C., & Hewitt, J. (in press). Postmodernism, knowledge-building, and elementary science. Elementary School Journal.
- Brown, J. S., Collins, A., & Duguid, P. (1989). Debating the situation: A rejoinder to Palincsar and Wineburg. Educational Researcher, 18, 10-12, 62.
- Charness, N. (1991). Expertise in chess: The balance between knowledge and search. In K. A. Ericsson & J. Smith (Eds.), Toward a general theory of expertise: Prospects and limits (pp. 39-63). Cambridge: Cambridge University Press.
- Cobb, P. (1994). Where is the mind? Constructivist and sociocultural perspectives on mathematical development. Educational Researcher, 23(7), xx.
- Donaldson, M. (1978). Children's minds. London: Croom Helm.

- Doyle, W. (1983). Academic work. Review of Educational Research, 53, 159-199.
- Dreyfus, H. L., & Dreyfus, S. E. (1986). Mind over machine. New York: Free Press.
- Dreyfus, H. L. (1988). The Socratic and Platonic basis of cognitivism. AI and Society, 2, 99-112.
- Drucker, P. F. (1994, November). The age of social transformation. Atlantic Monthly, p. 53-80.
- Gardner, H. (1985). The mind's new science: A history of the cognitive revolution. New York: Basic Books.
- Greeno, J. G. (1994). Understanding concepts in activity. In C. Weaver, C. R. Fletcher, & S. Mannes (Eds.), Discourse comprehension: Essays in honor of Walter Kintsch xx:
- Greeno, J. G., Smith, D. R., & Moore, J. L. (1993). Transfer of situated learning. In D. K. Detterman & R. J. Sternberg (Eds.), Transfer on trial: Intelligence, cognition, and instruction (pp. 99-167). Norwood, NJ: Ablex.
- Harris, K. R., Graham, S. (Eds.) (1994). Special Issue on Implications of Constructivism for Students with Disabilities and Students at Risk: Issues and Directions. Journal of Special Education, 28(3), 233-378.
- Hirsch, E. D., Jr. (1987). Cultural literacy: What every American needs to know. Boston, MA: Houghton Mifflin.
- Inhelder, B., & Piaget, J. (1958). The growth of logical thinking from childhood to adolescence. New York: Basic Books.
- Karmiloff-Smith, A. (1992). Beyond modularity: A developmental perspective on cognitive science. Cambridge, MA: MIT Press.
- Kitto, H. D. F. (1954). The Greeks. Baltimore, MD: Penguin Books.
- Luria, A. R. (1976). Cognitive development: Its cultural and social foundations (M. Lopez-Morillas & L. Soltaroff, Trans.). Cambridge, MA: Harvard University Press.
- Margolis, H. (1987). Patterns, thinking, and cognition. Chicago: University of Chicago Press.

- Montessori, M. (1967). The absorbent mind. New York: Holt, Rinehart, & Winston.
- Newell, A., & Simon, H. A. (1972). Human problem solving. Englewood Cliffs, NJ: Prentice-Hall.
- Ng, E., & Bereiter, C. (1991). Three levels of goal orientation in learning. The Journal of the Learning Sciences, 1(3 & 4), 243-271.
- Popper, K. R. (1972). Objective knowledge: An evolutionary approach. Oxford: Clarendon Press.
- Popper, K. R., & Eccles, J. C. (1977). The self and its brain. Berlin: Springer-Verlag.
- Resnick, L. B., & Neches, R. (1984). Factors affecting individual differences in learning ability. In R. J. Sternberg (Ed.), Advances in the psychology of human intelligence (pp. 275-323). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Rumelhart, D. E. (1989). The architecture of mind: A connectionist approach. In M. I. Posner (Eds.), Foundations of cognitive science (pp. 133-159). Cambridge, MA: MIT Press.
- Scardamalia, M., & Bereiter, C. (1992). Text-based and knowledge-based questioning by children. Cognition and Instruction, 9(3), 177-199.
- Scardamalia, M., Bereiter, C., & Lamon, M. (1994). CSILE: Trying to bring students into world 3. In K. McGilley (Ed.), Classroom lessons: Integrating cognitive theory and classroom practice. (pp. 201-228). Cambridge, MA: MIT Press.
- Scardamalia, M., Bereiter, C., Hewitt, J., & Webb, J. (in press). Constructive learning from texts in biology. In K. M. Fischer & M. Kirby (Eds.), Relations and biology learning: The acquisition and use of knowledge structures in biology. Berlin: Springer-Verlag.
- Schank, R. C., Collins, G. C., & Hunter, L. E. (1986). Transcending inductive category formation in learning. Behavioral and Brain Sciences, 9, 639-686.
- Scribner, S. (1979). Modes of thinking and ways of speaking: Culture and logic reconsidered. In R. O. Freedle (Ed.), New directions in discourse processing (pp. 223-243). Norwood, NJ: Ablex Publishing Corporation.

Tolman, E. C. (1949). There is more than one kind of learning. Psychological Review, 56, 144-155.

Vallas, S. P. (1990). The concept of skill: A critical review. Work and Occupations, 17, 379-398.

Vera, A. H., & Simon, H. A. (1993). Situated action: A symbolic interpretation. Cognitive Science, 17, 7-48.

Whitehead, A. N. (1929). The aims of education. New York: Macmillan.

Wineburg, S. S. (1989). Remembrance of theories past. Educational Researcher, 18, 7-10.

Woodworth, R. S. (1958). Dynamics of behavior. New York: Henry Holt