Scardamalia, M., & Bereiter, C. (2006). Knowledge building: Theory, pedagogy, and technology. In K. Sawyer (Ed.), ***Cambridge Handbook of the Learning Sciences*** **(pp. 97-118). New York: Cambridge University Press.**

Knowledge Building: Theory, Pedagogy, and Technology

Marlene Scardamalia and Carl Bereiter

There are substantial similarities between deep learning and the processes by which knowledge advances in the disciplines. During the 1960s efforts to exploit these similarities gave rise to learning by discovery, guided discovery, inquiry learning, and *Science: A Process Approach* (American Association for the Advancement of Science, 1967). Since these initial reform efforts, scholars have learned a great deal about how knowledge advances. A mere listing of keywords suggests the significance and diversity of ideas that have come to prominence since the 1960s: Thomas Kuhn, Imre Lakatos, sociology of science, the "Science Wars," social constructivism, schema theory, mental models, situated cognition, explanatory coherence, the "rhetorical turn," communities of practice, memetics, connectionism, emergence, and self-organization. Educational approaches have changed in response to some of these developments; there is a greater emphasis on collaborative rather than individual inquiry, the tentative nature of empirical laws is more often noted, and argumentation has become an important part of some approaches. But the new "knowledge of knowledge" has much larger educational implications: Ours is a knowledge-creating civilization. A growing number of "knowledge societies" (Stehr, 1994), are joined in a deliberate effort to advance all the frontiers of knowledge. Sustained knowledge advancement is seen as essential for social progress of all kinds and for the solution of societal problems. From this standpoint the fundamental task of education is to enculturate youth into this knowledge-creating civilization and to help them find a place in it.

In light of this challenge, traditional educational practice—with its emphasis on knowledge transmission—as well as the newer constructivist methods both appear to be limited in scope if not entirely missing the point.

Knowledge building, as elaborated in this chapter, represents an attempt to refashion education in a fundamental way, so that it becomes a coherent effort to initiate students into a knowledge creating culture. Accordingly, it involves students not only developing knowledge-building competencies but also coming to see themselves and their work as part of the civilization-wide effort to advance knowledge frontiers. In this context, the Internet becomes more than a desktop library and a rapid mail-delivery system. It becomes the first realistic means for students to connect with civilization-wide knowledge building and to make their classroom work a part of it.

The distinctiveness of a knowledge building approach was encapsulated for us by the comment of a fifth-grader on the work of a classmate: “Mendel worked on Karen’s problem” (referring to Gregor Mendel, the great 19th century biologist). Not “Karen rediscovered Mendel” or “Karen should read Mendel to find the answer to her problem.” Rather, the remark treats Karen’s work as continuous with that of Gregor Mendel, addressing the same basic problem. Furthermore, the Mendel reference is offered to help Karen and others advance their collective enterprise. In our experience, young students are delighted to see their inquiry connect with that of learned others, past or present.

Rather than being overawed by authority, or dismissive, they see their own work as being legitimated by its connection to problems that have commanded the attention of respected scientists, scholars, and thinkers.

In this chapter we elaborate six themes that underlie a shift from treating students as learners and inquirers to treating them as members of a knowledge building community. These themes are

- Knowledge advancement as a community rather than individual achievement
- Knowledge advancement as idea improvement rather than as progress toward true or warranted belief
- Knowledge *of* in contrast to knowledge *about*
- Discourse as collaborative problem solving rather than as argumentation
- Constructive use of authoritative information
- Understanding as an emergent

One important advantage of knowledge building as an educational approach is that it provides a straightforward way to address the contemporary emphasis on knowledge creation and innovation. These lie outside the scope of most constructivist approaches whereas they are at the heart of knowledge building.

Community Knowledge Advancement

In every progressive discipline one finds periodic reviews of the state of knowledge or the "state of the art" in the field. Different reviewers will offer different descriptions of the state of knowledge; however, their disagreements are open to argument that may itself contribute to advancing the state of knowledge. The state of knowledge is not what everyone in the field or the average person in the field knows, but neither is it what the most knowledgeable people in the field know, except in some collective sense. Fundamentally, a description of the state of knowledge is not about what is in people's minds at all. If we look back at prehistoric times, using archaeological evidence, we can make statements about the state of knowledge in a certain civilization at a certain time, without knowing anything about any individuals and what they thought or knew.

An implicit assumption in state-of-the-art reviews is that the knowledge in a field does not merely accumulate but advances. There is the implicit image of a moving body, taking in new information and ideas at its leading edge and leaving behind solved or abandoned problems and disproved or outmoded ideas. Creative knowledge work may be defined as work that advances the state of knowledge within some community of practice, however broadly or narrowly that community may be defined.

Knowledge building pedagogy is based on the premise that authentic creative knowledge work can take place in school classrooms—knowledge work that does not merely emulate the work of mature scholars or designers but that substantively advances the state of knowledge in the classroom community and situates it within the larger societal knowledge building effort. This is a radically different vision from contemporary educational practice, which is so intensely focused on the individual student that the notion of a state of knowledge that is not a mental state or an aggregate of mental states seems to make no sense. Yet in knowledge creating organizations it makes obvious sense. People are not honored for what is in their minds but for the contributions they make to the organization's or the community's knowledge.

One component of knowledge building is the creation of "epistemic artifacts," tools that serve in the further advancement of knowledge (Sterelny, 2005). These may be purely conceptual artifacts (Bereiter, 2002), such as theories and abstract models, or "epistemic things" (Rheinberger, 1997), such as concrete models and experimental set-ups. Epistemic artifacts are especially important in education, where the main uses of knowledge are in the creation of further knowledge. When we speak of engaging students in "the deliberate creation and improvement of knowledge that has value for a community"(Scardamalia & Bereiter, 2003) the main value is this epistemic one—a feedforward effect, in which new knowledge gives rise to and speeds the development of yet newer knowledge. In this context, student-generated theories and models are to be judged not so much by their conformity to accepted knowledge as by their value as tools enabling further growth.

Idea Improvement

Engineers and designers do not think in terms of a final state of perfection (Petroski, 2003). Advances in a technology open up new problems to be solved and new possibilities for further advancement, so there is no end in sight. But many people still think of knowledge as advancing toward (though perhaps never reaching) a final state, which is truth: how the universe actually began, the true history of the invasion of Iraq, and so on. But advances in theoretical and historical knowledge always raise new problems and open new possibilities, just as do advances in technology. Except in a few areas such as disease control, progress is measured by comparison to what has gone before, rather than by distance to a predetermined end-point.

As a criterion for evaluating individual performance, "improvement" is a familiar although by no means universally accepted notion in educational assessment. But improvement as a criterion for assessing knowledge itself is virtually unheard of. Here is an example of what this would mean in a learning context. Analysis of a grade 5/6 Knowledge Forum database showed that most of the students initially conceived of gravity as a substance residing within objects rather than as a relation between objects (as is typical: Chi, Slotta, & deLeeuw, 1994). By the end of a unit on gravity, most students still treated gravity as a substance. Thus, measured as distance from the goal—to teach students that gravity is a relationship between two masses—there had been little progress. However, comparing students' end-of-unit writings on gravity with their initial ones, changes could be detected. The students appeared less comfortable with the substance conception and more aware that there were other conceptions, even though they had not yet grasped them. There was also an awareness that gravity is everywhere and not just a property of large celestial bodies. One student wrote: "I need to understand. I know we are a mass our self but then why aren't little parts of dust and small objects attracted to us?…We are much bigger than a small ripped up pieces of paper, but yet you don't see the paper fly across the room or even a small distance to us. WHY?"

We noted similar patterns in another grade 5/6 class that had been studying evolution. Natural selection had not taken hold as the key explanatory concept, although there was a growing recognition that it had something to do with evolution. More tellingly, there was a growing recognition that *some* mechanism of evolution was required, that evolutionary adaptation could not merely be accepted as a primitive—the view that Ohlsson (1991) found characteristic of university undergraduates.

In knowledge building, idea improvement is an explicit principle, something that guides the efforts of students and teachers rather than something that remains implicit in inquiry and learning activities (Scardamalia, 2002). The direct pursuit of idea improvement brings schooling into much closer alignment with creative knowledge work as carried on at professional levels. Generating ideas appears to come naturally to people, especially children, but sustained effort to improve ideas does not. We believe that developing a disposition to work at idea improvement should be a major objective in the education of scholars, scientists, and designers, for without such a disposition the likelihood of a productive career is slight.

To propose idea improvement as an alternative to progress toward truth may suggest a relativist, anti-foundationalist, or extreme social-constructivist theory of knowledge. The point we want to make here, however, is that you need not take a position on this issue in order to adopt a knowledge building pedagogy with idea improvement as a core principle. You can hold that there are preexisting truths and that, short of revelation, idea improvement is our only means of working toward them; or you can hold that what pass for truths are just conceptual artifacts that have undergone a successful process of development. All that is necessary is to adopt as a working premise that all ideas are improvable— or, at any rate, all interesting ideas.

An educational program committed to idea improvement has to allow time for iterations. Iterative idea improvement is in principle endless; in practice, the decision whether to continue a particular line of knowledge building or shift to another is a judgment call, taking into account the progress being made and measuring it against competing demands and opportunities. Ideally (although it is difficult in a graded school system), a student cohort should be able to pick up a thread of inquiry at a later time— even years later. For instance, elementary school students studying electricity often develop a good qualitative understanding of circuits, resistance, and conductance. They may be able to formulate and test interesting hypotheses about why some materials conduct electricity and others apparently do not. But they are unlikely to be able to grasp what electric current actually is. Instead of starting over in high school science, they could reconsider their earlier speculations in light of the more sophisticated concepts now available to them. Electronic media make such continuity technically feasible and could help to bring school knowledge building into closer alignment with the way knowledge advances in the disciplines.

One distinctive characteristic of students in knowledge building classrooms reflects epistemological awareness. When asked about the effects of learning, students in regular classrooms tend to say that the more they learn and understand, the less there remains to be learned and understood (a belief that accords well with the fixed curriculum that directs their work). Students in knowledge building classrooms, however, tend strongly toward the opposite view, as expressed by one fourth-grade student: "By researching it [a particular knowledge problem] you can find other things that you want to research about. And so you realize that there is more and more and more things that you don't know… so, first you know this much [gestures a small circle] and you know there is this much [gestures a large circle] that you don't know. Then you know this much [gestures a larger circle] but you know there is this much [gestures an even larger circle] that you don't know, and so on and so on."

Knowledge *of* in Contrast to Knowledge *about*

Since the 1970s, cognitive scientists largely focus on two broad types of knowledge, declarative and procedural (Anderson, 1980). This distinction now pervades the cognitive literature as well as educational psychology textbooks that take a cognitive slant. The declarative-procedural distinction has proven useful in rule-based computer modeling of cognitive processes, but its application to education and knowledge creation is questionable (Bereiter, 2002, ch. 5). From a pragmatic standpoint, a more useful distinction is between knowledge *about* and knowledge *of* something. Knowledge *about* sky-diving, for instance, would consist of all the declarative knowledge you can retrieve when prompted to state what you know about sky-diving. Such knowledge could be conveniently and adequately represented in a concept net. Knowledge *of* sky-diving, however, implies an ability to do or to participate in the activity of sky-diving. It consists of both procedural knowledge (e.g, knowing how to open a parachute and guide its descent) and declarative knowledge that would be drawn on when engaged in the activity of sky-diving (e.g., knowledge of equipment characteristics and maintenance requirements, rules of particular events). It entails not only knowledge that can be explicitly stated or demonstrated, but also implicit or intuitive knowledge that is not manifested directly but must be inferred (see Bransford et al., this volume). Knowledge *of* is activated when a need for it is encountered in action. Whereas knowledge *about* is approximately equivalent to declarative knowledge, knowledge *of* is a much richer concept than procedural knowledge.

Knowledge *about* dominates traditional educational practice. It is the stuff of textbooks, curriculum guidelines, subject-matter tests, and typical school "projects" and "research" papers. Knowledge *of*, by contrast, suffers massive neglect. There is instruction in skills (procedural knowledge), but it is not integrated with understanding in a way that would justify saying "Alexa has a deep knowledge *of* arithmetic"—or chemistry or the stock market or anything else. Knowledge about is not entirely useless, but its usefulness is limited to situations in which knowledge *about* something has value independently of skill and understanding. Such situations are largely limited to social small talk, trivia games, quiz shows, and—the one biggy—test taking.

To be useful outside the limited areas in which knowledge *about* is sufficient, knowledge needs to be organized around problems rather than topics (Bereiter, 1992). Of course, topics and problems often go together, but in the most interesting cases they do not—for example, when the connection of knowledge to a problem is analogical, via deeper underlying mechanisms rather than surface resemblance. Such connections are vital to invention, theorizing, and the solving of ill-structured problems. For instance, it is useful for learners' knowledge of water skiing to be activated when they are studying flight, because it provides a nice experiential anchor for the otherwise rather abstract "angle of attack" explanation of lift. Ordinarily the teacher is responsible for making such connections, but in the out-of-school world people need to be able to do this themselves if they are to succeed as knowledge-builders. Making this connection promotes the realization that Bernoulli's principle is not the whole story in explaining what keeps airplanes aloft. Ordinarily the teacher is responsible for alerting students to such connections, but in the out-of-school world people need to be able to do this themselves if they are to succeed as knowledge-builders.

Across a broad spectrum of theoretical orientations, instructional designers agree that the best way to acquire what we are calling knowledge *of* is through problem solving—as in the *driving questions* of project-based learning (Krajcik & Blumenfeld, this volume) and in inquiry learning more generally (Edelson & Reiser, this volume). Research on transfer makes it clear, however, that solving problems does not automatically generate the deep structural knowledge on which analogical transfer is based (Catrambone & Holyoak, 1989). Problem-based learning environments fall somewhere on a continuum between context-limited to context-general work with knowledge (Bereiter & Scardamalia, 2003; in press). At the context-limited extreme, students' creative work is limited to problems of such a concrete and narrowly focused kind that they do not raise questions about general principles. Accordingly, the more basic knowledge (of scientific laws or causal mechanisms, for instance) that the curriculum calls for is often left to be conveyed by conventional instructional means. This raises concern that the deep knowledge that is most useful for transfer will not be connected with problems but will remain as knowledge *about* the relevant principles or laws. In knowledge building, students work with problems that result in deep structural knowledge *of*.

Knowledge-Building Discourse

In the view of science that flourished 50 years ago and that is still prominent in school science, discourse is primarily a way of sharing knowledge and subjecting ideas to criticism, as in formal publications and oral presentations, and question-and-answer sessions after these presentations. Lakatos (1976) challenged this idea, showing how discourse could play a creative role—actively improving on ideas, rather than only acting as a critical filter. Recent empirical studies of scientific discourse support Lakatos's view. For example, Dunbar

(1997) showed that the discourse that goes on inside research laboratories is fundamentally different from the discourse that goes on in presentations and papers—it is more cooperative and concerned with shared understanding. Public discourse and collaborative discourse serve complementary functions, and practitioners of a discipline need to be proficient in both (Woodruff & Meyer, 1997). However, cooperative discourse oriented toward understanding is much more relevant to learning (Coleman, Brown, & Rivkin, 1997).

There are weak and strong versions of the claim that collaborative discourse plays a role in knowledge advancement. The weak version holds merely that *empirical findings and other products of inquiry only become contributions to community knowledge when they are brought into public discourse*. This version is compatible with the conventional view of discourse as knowledge sharing. The strong version asserts that *the state of public knowledge in a community only exists in the discourse of that community, and the progress of knowledge just is the progress of knowledge-building discourse*. If, as we argued earlier, the state of knowledge of a community is not something in the minds of individual members of the community, then there is no place else it can exist except in discourse. The weak version holds that the advance of knowledge is *reflected in* the discourse, whereas the strong version holds that there is no advance of community knowledge *apart from* the discourse. (Note that this is not a declaration about *what* knowledge is; it is only a self-evident statement about *where* public knowledge is.)

Both versions require that discourse be treated as having content, that it cannot be all form and process, and that this content can be described and evaluated outside the discourse in which it is constituted. Thus there has to be the possibility of a metadiscourse that takes the *content* of the first-order discourse as its subject. Knowledge building discourse, as we conceive of it, is discourse whose aim is progress in the state of knowledge: idea improvement. It involves a set of commitments that distinguish it from other types of discourse (Bereiter, 1994, 2002):

- a commitment to *progress*, something that does not characterize dinner party conversation or discussions devoted to sharing information and venting opinions
- a commitment to *seek common understanding* rather than merely agreement, which is not characteristic of political and policy discourse, for instance
- a commitment to *expand the base of accepted facts*, whereas, in court trials and debates, attacking the factual claims of opponents is common

By these criteria, argumentation and debate, as currently promoted in schools, falls short. Its emphasis on evidence and persuasion, while admirable in other respects, does not generate progress toward the solution of shared problems of understanding. Knowledge-building discourse in the classroom has a more constructive and progressive character (Bereiter, Scardamalia, Cassells, & Hewitt, 1997).

Constructive Use of Authoritative Information

The use of authoritative information has presented problems for educators ever since the advent of student-centered and constructivist education. On the one hand, we do not want students to meekly accept authoritative pronouncements. “Because I say so” and “because the book says so” are no longer regarded as acceptable responses to students’ skeptical queries. On the other hand, it is impossible to function in society without taking large amounts of information on authority. Even when it comes to challenging authoritative pronouncements, doing so effectively normally depends on bringing in other authoritative information as evidence.

A focus on knowledge building alleviates even if it does not solve the problems associated with authoritative information. Information of all kinds, whether derived from first-hand experience or from secondary sources, has value insofar as it contributes to knowledge building discourse. Quality of information is always an issue, but its importance varies with the task. If the task is one where faulty design will put lives at risk (design of a new drug or of a suspension bridge, for instance), a much higher standard of information quality will be required than if less is at stake or if self-corrective measures can be built into the design. Judging the quality of information is not a separate problem from the knowledge building task, it is part of the task. Judgment may involve argument, but it is argument in the service of the overall idea improvement mission.

Emergent Understanding

How are complex new concepts acquired? Indeed, how is it logically possible to learn "a conceptual system richer than the one that one already has" (Fodor, 1980, p. 149)? The "learning paradox, "as it has come to be called (Pascual-Leone, 1980; Bereiter, 1985), poses a fundamental problem for constructivism: If learners construct their own knowledge, how is it possible for them to create a cognitive structure more complex than the one they already possess? Dozens of articles have appeared claiming to resolve the paradox but in fact failing to address the fundamental problem. The only creditable solutions are ones that posit some form of self-organization (Quartz, 1993; Molenaar & van der Maas, 2000). At the level of the neural substrate, self-organization is pervasive and characterizes learning of all kinds (Phillips & Singer, 1997). As Grossberg (1997, p. 689) remarked, "brains are self-organizing organs par excellence." Explaining conceptual development, however, entails self-organization at the level of ideas—explaining how more complex ideas can emerge from interactions of simpler ideas and percepts.

New conceptual structures, like crystals and ant colonies, emerge through the interaction of simpler elements that do not singly or in combination represent the new concept (Sawyer, 2003). This became evident with the rise of connectionism in the late 1980s (Bereiter, 1991). Connectionist models of learning and development characteristically generate progress from a conceptually impoverished to a conceptually richer system, sometimes by a process analogous to learning from experience and sometimes only by internal self-organization. Connectionist models are examples of the larger class of dynamic systems models, all of which attempt to deal in some rigorous way with emergent phenomena. The emergence of complexity from the interaction of simpler elements is found at all levels from the physico-chemical to the socio-cultural. If learning is paradoxical, so is practically everything else that goes on in the world.

The frequently stated constructivist principle, “Learners construct their own knowledge,” can be restated in dynamic systems terms as “All understandings are inventions; inventions are emergents.” Two obstacles stand in the way of making this more than just a restatement of the same vague principle. First, explanations in terms of dynamic systems are difficult to understand and do not yield the satisfying gestalts that attend narrative explanations. Second—and this is an obstacle much less commonly recognized—a dynamic systems explanation of conceptual growth posits (along with other kinds of interactions) ideas interacting with ideas to generate new ideas. This level of description is common in the philosophy of knowledge and in the history of ideas. The practical import of this discussion is that instructional designers need to think more seriously about ideas as real things that can interact with one another to produce new and more complex ideas. School-age students have shown themselves able to make sense of and profit from computer representations of self-organization at the idea level (Ranney & Schank, 1988).

From Computer Supported Intentional Learning to Knowledge Building Environments

Although the term “knowledge building” is now in wide use (in 125,000 Web documents, as of July, 2005) we were, as far as we can ascertain, the first to use the term in education, and certainly the first to have used it as something more than a synonym for active learning. Prying loose the concept of knowledge building from concepts of learning has been an evolutionary process, however, which continues. An intermediate concept is “intentional learning”(Bereiter & Scardamalia, 1989)—something more than “active” or “self-regulated” learning, more a matter of having life goals that include a personal learning agenda. This concept grew out of research revealing the opposite of intentional learning: students employing strategies that minimize learning while efficiently meeting the demands of school tasks (Brown, Day, & Jones, 1983; Scardamalia & Bereiter, 1987). Although students were responsive to a more “knowledge-transforming” approach (Scardamalia, Bereiter, & Steinbach, 1984), effects dissipated when they returned to ordinary classroom work. Many characteristics of classroom life conspire to discourage intentional learning (Scardamalia & Bereiter, 1996), but a key factor seems to be the structure of classroom communication, in which the teacher serves as the hub through which all information passes. Altering that information flow was one of our goals when we designed the software application we called CSILE—Computer Supported Intentional Learning Environments— first used in early prototype version in 1983 in a university course, more fully implemented in 1986 in an elementary school (Scardamalia, Bereiter, McLean, Swallow, and Woodruff, 1989).

Another motive guiding the design of CSILE was a belief that students themselves represented a resource that was largely wasted and that could be brought into play through network technology (Scardamalia & Bereiter, 1991). Classroom work with CSILE proved this to be true beyond anything we had imagined. The classroom, as a community, could indeed have a mental life that is not just the aggregate of individual mental lives but something that provides a rich context within which those individual mental lives take on new value. CSILE restructured the flow of information in the classroom, so that questions, ideas, criticisms, suggestions, and the like were contributed to a public space equally accessible to all, instead of it all passing through the teacher or (as in e-mail) passing as messages between individual students. By linking these contributions, students created an emergent hypertext that represented the collective rather than only the individual knowledge of the participants. We introduced epistemological markers ("My theory," "I need to understand," "New information," and so on), through "thinking types" that could be integrated into the text of notes, as students chose, to encourage metadiscourse as well as discourse focused on the substantive issues under investigation.

By the 1990s the idea of knowledge building as the collaborative creation of public knowledge had assumed ascendancy, with individual learning as an important and demonstrable by-product (Scardamalia, Bereiter, & Lamon, 1994). In this light, we undertook a major redesign of CSILE to boost it as an environment for objectifying ideas and their interrelationships and to support collaborative work aimed at improving ideas.

In scientific and scholarly research teams, knowledge building often proceeds with no special technology to support it. This is possible because knowledge building is woven into the social fabric of the group and in a sense all the technology used by the group supports it. This becomes evident if we consider successful research laboratories like those studied by Dunbar (1997) in light of the themes previously discussed:

- Knowledge advancement is the defining purpose of the research laboratory, and so it is not difficult to keep this purpose salient; schools, by contrast, have a multiplicity of purposes touching on many different aspects of student development.
- Although publications, speaking invitations, patents, and grants are markers of success in the research world, they all depend finally on idea improvement. You cannot get on the program at a scientific meeting or be awarded a patent by simply repeating last year's successful idea. In schools, by contrast, reproduction of existing ideas figures prominently in learning activities and assessment.
- Expertise in the research world presupposes deep knowledge *of* the problem domain; mere knowledge *about* gains little credit. In the school world, however, knowledge *about* is the basic indicator of academic achievement. A knowledge building technology, accordingly, ought to favor increasingly deep inquiry into questions of *how* and *why* rather than the shallower kinds of inquiry guided by questions of *what* and *when*.
- Discourse within a research group is geared to advancing the group's knowledge building goals. Argumentation about knowledge claims takes place in public arenas. In the classroom, however, discourse can serve a wide range of purposes, from selfexpression to knowledge recitation. Communication technology should help to move discourse along a knowledge building path.
- Constructive use of authoritative information comes naturally to a research organization; original work is almost always built upon previous work, and theories

are tested against data not only from local work but also from published research (Bazerman, 1985). In school, however, authoritative information is most commonly brought forward as *that which is to be learned*. Using it in knowledge building therefore requires a shift in focus, which may require external support. A knowledge building technology should facilitate *using* information, as distinct from learning it. Obtaining, recording, and storing information would become subsidiary functions, designed to serve purposes of knowledge creation.

- Significant advances in knowledge by a research laboratory are obviously emergents; the knowledge didn't pre-exist in anyone's mind nor was it simply there to be read out of the "book of nature." But in schools a major concern is students' acquisition of knowledge that already exists as part of the culture. It needs to be recognized, however, that grasping this knowledge is also emergent, and so knowledge building technology for schools needs to be essentially the same as what would support the work of knowledge creating organizations.

The next generation of CSILE, called Knowledge Forum®, provides a knowledge building environment for communities (classrooms, service and health organizations, businesses, and so forth) to carry on the sociocognitive practices described above—practices that are constitutive of knowledge- and innovation- creating organizations. This is a continuing challenge; Knowledge Forum undergoes continual revision as theory advances and experience uncovers new problems and opportunities. It is an extensible environment supporting knowledge building at all educational levels, and also in a wide range of non-educational settings.

The distinctive characteristics of Knowledge Forum are perhaps most easily grasped by comparing it to the familiar technology of threaded discussion, which is to be found everywhere on the Worldwide Web and also as a part of instructional management systems like Blackboard and WebCT. Threaded discussion is a one-to-many form of e-mail. Instead of sending a message privately to people the sender selects, the sender "posts" it to a discussion site, where all posted messages appear in chronological order, with one exception: a response to a message is shown indented under the original message, rather than in chronological order. Responses to that response are further indented, and so on, forming a "thread" that started with the very first posting. Like e-mail messages generally, a discussion forum message, once "posted," cannot be modified. "Threading" produces a downward-branching tree structure, which is the only structuring of information (besides chronological) that the technology allows. There is no way to create higher-level organizations of information, to comment simultaneously on a number of messages, or to make a connection between a message in one thread and a message in another. Thus the possibilities for knowledge building discourse are extremely limited. In fact, our experience is that threaded discussion militates against deepening inquiry; instead, it is much more suited to rapid question-answer and assertion-response exchanges. Although communities based on shared interests do develop in some threaded discussion forums, this technology provides little means for a group to organize its efforts around a common goal. As the number of postings increases, what appears on the screen becomes an increasingly incoherent stream of messages, leading discussion monitors to impose arbitrary limits on thread length and to erase threads of a certain age. Thus a cumulative advance in the state of knowledge is hardly conceivable.

Knowledge Forum's technological roots are not in e-mail at all. Knowledge Forum is a multimedia database, designed so as to maximize the ability of a community of users to create and improve both its content and organization. Thus the database itself is an emergent, representing at different stages in its development the advancing knowledge of the community. From the users' standpoint, the main constituents of a Knowledge Forum database are *notes* and *views*. A view is an organizing background for notes. It may be a concept map, a diagram, a scene—anything that visually adds structure and meaning to the notes whose icons appear in it. Notes are contributed to views and may be moved about to create organization within views. The same notes may appear in more than one view. Fig. 1 shows several different views of the same notes produced by first-graders in studying dinosaurs.

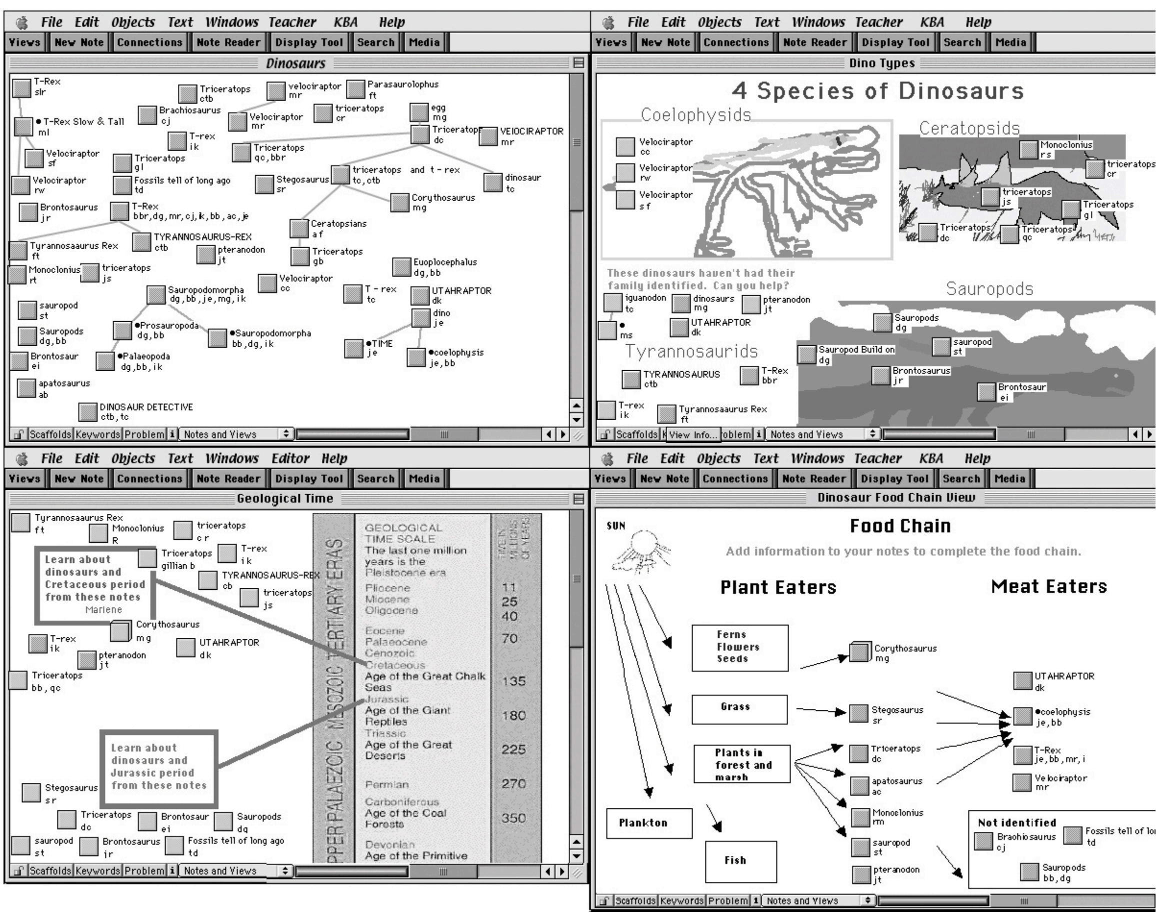

Figure 1: Four different user-generated graphical representations of the same notes illustrate the multiple perspectives, multiple literacies, and teamwork enabled by CSILE/Knowledge Forum.

Wherever one is in a Knowledge Forum database, it is always possible to move downward, producing a lower-level note, comment, or subview; upward, producing a more inclusive note or a view of views; and sideways, linking views to views or linking notes in different views. Notes themselves may contain graphics, animations, movies, links to other applications and applets, and so on.

Knowledge Forum lends itself to a high level of what we call "epistemic agency" (Scardamalia, 2000). Although among philosophers this term denotes responsibility for one's beliefs (Reed, 2001), we use the term more broadly: epistemic agency refers to the amount of individual or collective control people have over the whole range of components of knowledge building—goals, strategies, resources, evaluation of results, and so on. Students can create their own views, as can authorized visitors (telementors) from outside the class. Groups of students may be given responsibility for different views, working to improve their usefulness to the class, to remove redundancies, and so on. Knowledge

Forum provides "scaffolds" to help shape discourse to knowledge building purposes—for instance, a set of theory-building scaffolds that include "My theory," "New information," "This theory explains," and "This theory cannot explain." Similar supports have been used in other collaborative learning software (see Andriessen, this volume; Edelson & Reiser, this volume; Stahl, Koschmann, & Suthers, this volume; Linn, this volume), but typically their use, and sometimes even the order in which they are used, is mandatory. In

Knowledge Forum use of the scaffolds is optional, and they may be modified as knowledge building progresses. One fourth-grade class decided that they were doing too much “knowledge telling” and so they introduced new scaffolds to focus attention on ideas.

We designed Knowledge Forum not simply as a tool, but as a knowledge building *environment*—that is, as a virtual space within which the main work of a knowledge building group would take place (Scardamalia, 2003). It has proved useful not only in formal educational settings but also in other circumstances where groups are striving to become knowledge building organizations—service and professional organizations, teacher development networks, and businesses that are aiming to boost their innovative capabilities. Giving pragmatic support to the idea that the same process underlies both school learning and high-level knowledge creation, the same version of Knowledge Forum has been used without modification at levels ranging from kindergarten to graduate school and professional work.

Of course, students using Knowledge Forum do not spend all their time at the computer. They read books and magazines, have small-group and whole-class discussions, design and carry out experiments, build things, go on field trips, and do all the other things that make up a rich educational experience. But instead of the online work being an adjunct, as it typically is with instructional management systems, bulletin boards, and the like, Knowledge Forum is where the main work takes place. It is where the “state of knowledge” materializes, takes shape, and advances. It is where the results of the various off-line activities contribute to the overall effort. If students run into a problem, they often recommend starting a space in Knowledge Forum to preserve and work out the ideas. At the end of Grade 1, a child moving to a class without Knowledge Forum asked, “Where will my ideas go? Who will help me improve them?” The Grade 2 teacher decided to use Knowledge Forum; the child’s Grade 1 ideas lived on, to be improved along with new ideas generated in Grade 2.

Knowledge Building Pedagogy

A knowledge building pedagogy evolved along with the technology, with teachers’ innovations and students’ accomplishments instrumental in this evolution. Two different progressions in pedagogy over three-year periods are reported by Scardamalia, Bereiter, Hewitt, and Webb (1996) and Messina and Reeve (2004). The goal was not to evolve a set of activity structures, procedures, or rules, but rather a set of workable principles that could guide pedagogy in a variety of contexts. The six themes that have framed the discussion in this chapter reflect this emphasis, as does a more fully elaborated set of 12 knowledge building principles (Scardamalia, 2002). The problem has been that principles— whether framed as goals, rules, beliefs, design parameters, or diagnostic questions—are viewed by some as too abstract to be very helpful and by others as mere redescriptions of things they already do. Movies and examples from student and teacher work are effective in arousing interest in knowledge building and in showing that something different from more familiar constructivist, discovery, and collaborative learning approaches is going on, but the result is a heightened demand for “how to do it” recommendations.

Adhering to a principled rather than a procedural approach has undoubtedly impeded the spread of knowledge building and Knowledge Forum, but the quality and innovativeness of the work carried on by teachers who have assimilated the principles appears to justify the approach. Numerous examples may be found in the posters presented at the annual Summer Institute on Knowledge Building. Abstracts are available at http://ikit.org/summerinstitutes.html. An unanticipated benefit of a principle- based approach is that the students themselves may begin to use knowledge building principles in conceptualizing their own work. We have already mentioned the students who diagnosed their work as “knowledge telling”—a term derived from a cognitive model of immature composing processes (Scardamalia, Bereiter, & Steinbach, 1984). Caswell and Bielaczyk (2001) report students’ productive use of the principle of “improvable ideas.” In another class, elementary school students in an inner city school—identified as one of the neediest in Toronto—have studied and begun to apply such concepts as epistemic agency, pervasive knowledge building, and community knowledge, and to describe their work at the Knowledge Building Summer Institute. These reports are themselves striking illustrations of the principle of turning higher levels of agency over to students. For decades educators have promoted constructivist ideas among themselves whereas their students have been expected to carry out constructivist activities without access to the constructivist ideas lying behind them. There is an internal contradiction there that a principled approach to knowledge building should overcome.

Figures 1 through 6 illustrate elementary school knowledge building in Toronto and Hong Kong, as supported by Knowledge Forum. The notes in Figure 1 were produced by Grade 1-3 students who were contributing information and graphics concerning their favorite dinosaurs. The upper-left view shows what the discourse space looked like after the students had entered their early notes; these notes are not organized in any particular way. Soon after these initial postings were completed, the children discovered classmates who had the same favorite dinosaur (triceratops, brontosaurus, etc.). Several students had produced graphic rather than text notes, and others wanted to link their notes to these graphics. So students used these graphics to draw the background of a new view that organized the notes according to dinosaur type; this new view is shown in the upper-right corner of Figure 1.

At about the same time, students in a university course were provided with access rights to this Grade 1-3 knowledge-building discourse. The university students noted, in reading these same notes, that they contained references to geological time, and they created a new 'geological time' view and entered a geological-timeline graphic from the Internet as a background (see the lower left frame of Figure 1). They then searched the primary students' notes for periods of time (e.g., Jurassic), and the new collection was added at the appropriate point to the geological timeline. When the primary students took a look at this new view, those who had not yet identified the time when their dinosaur roamed the earth quickly extended their research so their note would appear in this new view.

The last pane of figure 1 (lower right frame) demonstrates yet another view of these same notes. A biologist was invited to join the knowledge building collaborative efforts. She signed in from her office and created the 'food chain' view that referenced students' dinosaurs as either plant or meat eaters.

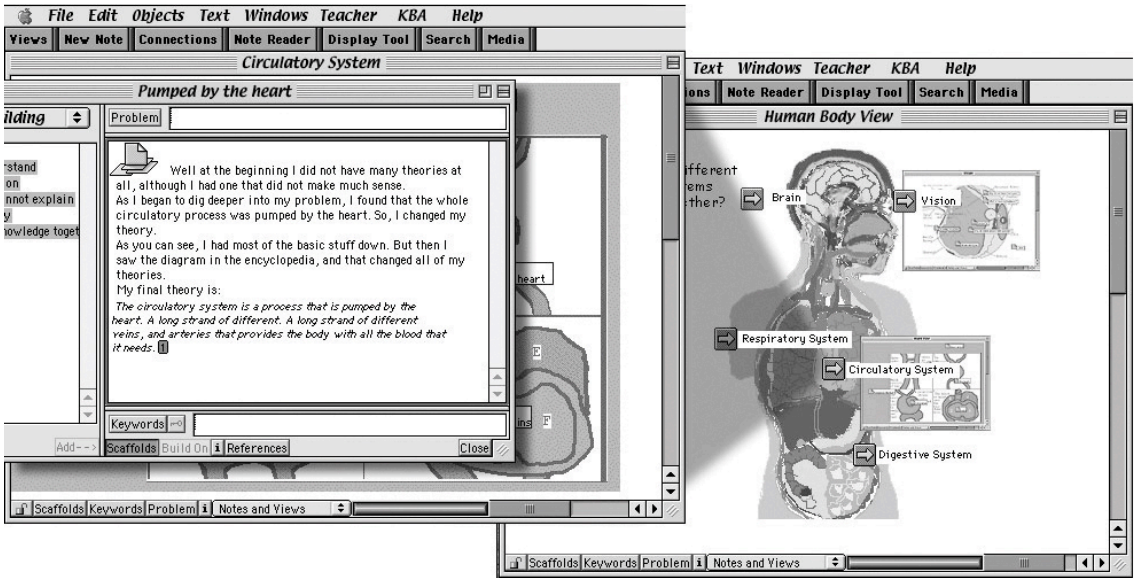

Figure 2: Rise-above and endless improvability of ideas.

Figure 2 is drawn from a Knowledge Forum database from a grade 5/6 class researching "systems of the body." The left side of Figure 2 shows what is called a "rise-above" note—in this case a student's summary of his knowledge advances made over a period of several months. The rise-above note subsumes a number of previous notes, which are now accessible only through this rise-above note. Rise-above notes are also used to synthesize ideas, create historical accounts and archives, reduce redundancy, and in other ways impose higher levels of organization on ideas.

The right side of Figure 2 illustrates the rise-above idea applied to views rather than notes. The smaller pictures are links to separate views created by groups of students working on different body systems. Later, the higher-order "Human Body"view was created to integrate these separate views and to support a new discourse on how different parts of the bodywork together. As this figure suggests, notes and views operate as a form of "zoom in/zoom out," encouraging users to think in terms of relationships.

Endless improvability of ideas is further supported by the following:

- Ability to create increasingly high-order conceptual frameworks. It is always possible to reformulate problems at more complex levels, by creating a rise-above note that encompasses previous rise-above notes, or to create a more inclusive view-of-views.
- Review and Revision. Notes and views can be revised at any time, unlike most discussion environments that disallow changes after a note is posted.
- Published notes and views. Processes of peer review and new forms of publication engage students in group editorial processes. Published works appear in a different visual form and searches can be restricted to the published layer of a database.

Figures 3 through 6 show a progression across grade levels in the kinds of knowledge building achieved when the whole school is committed to it. These examples come from the Institute of Child Study at the University of Toronto, where knowledge building is so embedded in the work of the school that quite a few students have more experience than their teachers and are instrumental in introducing new teachers not only to Knowledge Forum technology but also to the knowledge building culture of the school.

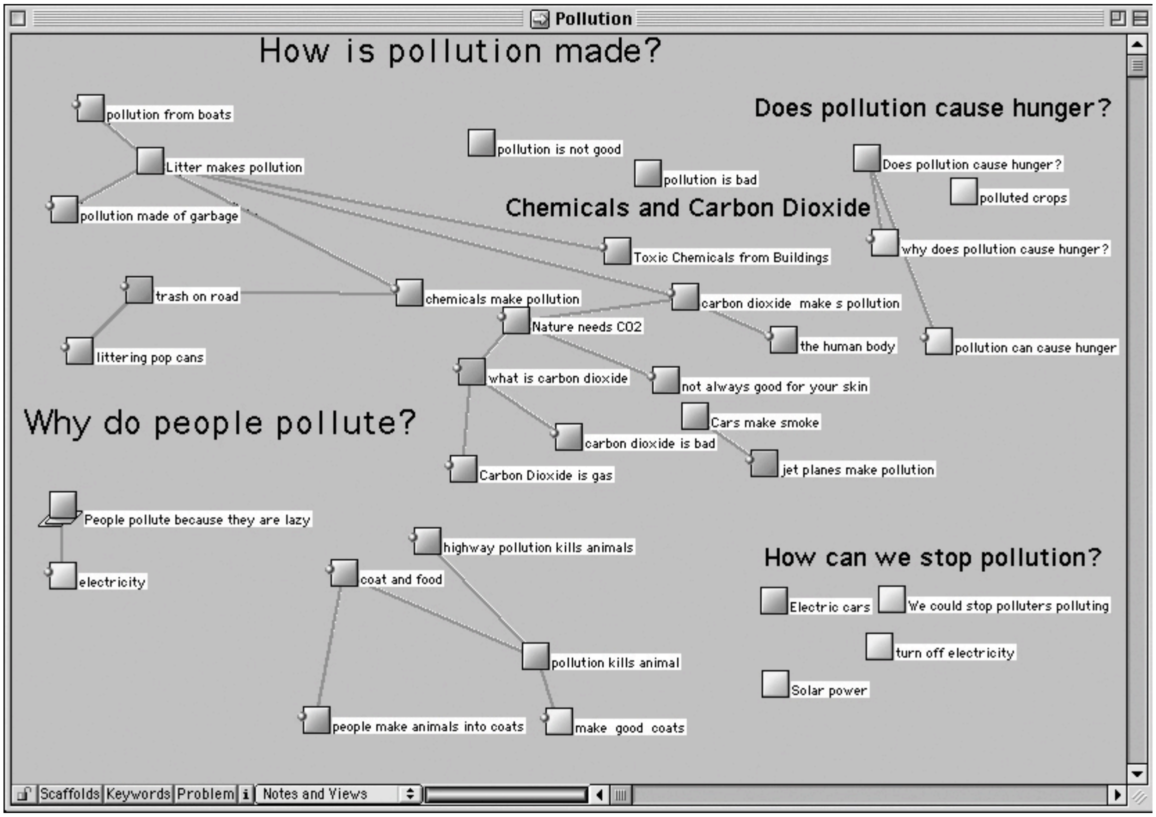

Figure 3: Rise-above view and note by grade 1 students studying pollution.

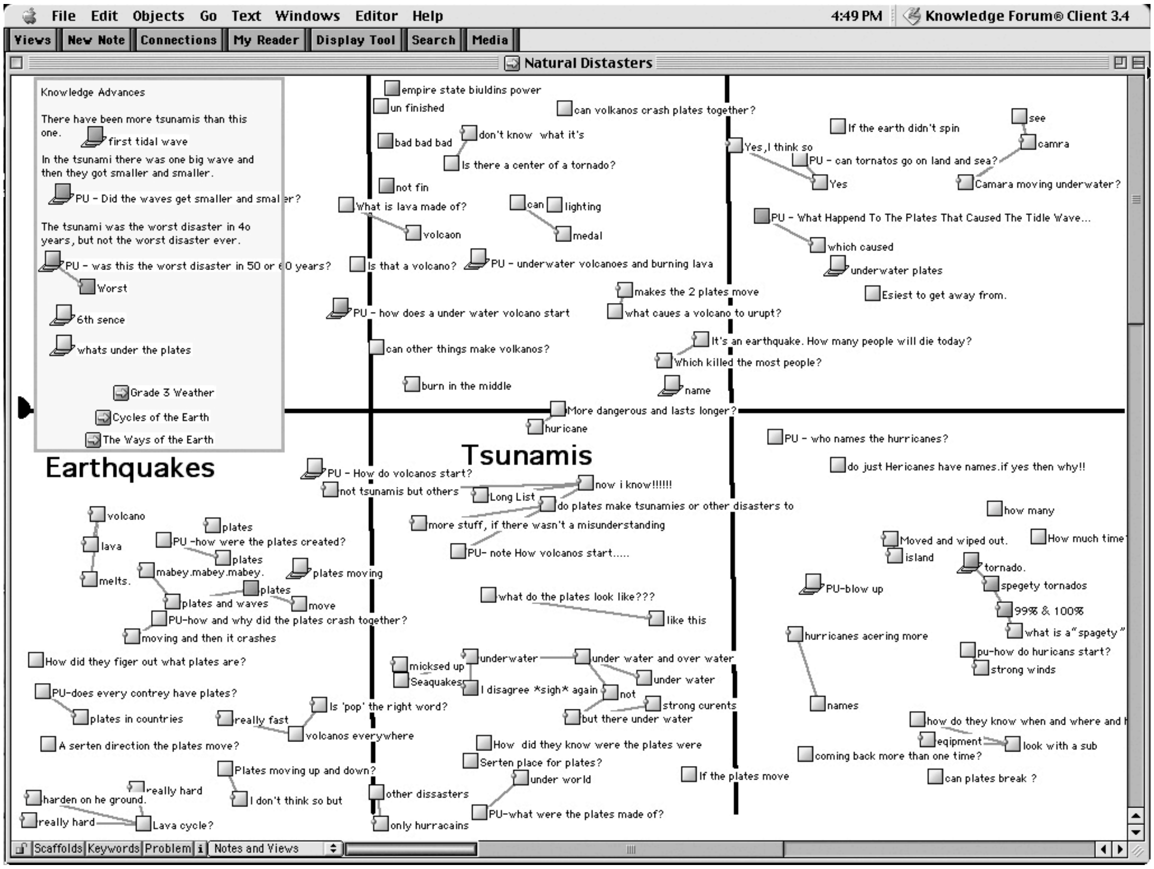

Figure 4: Rise-above view and notes by grade 3 students studying natural disasters.

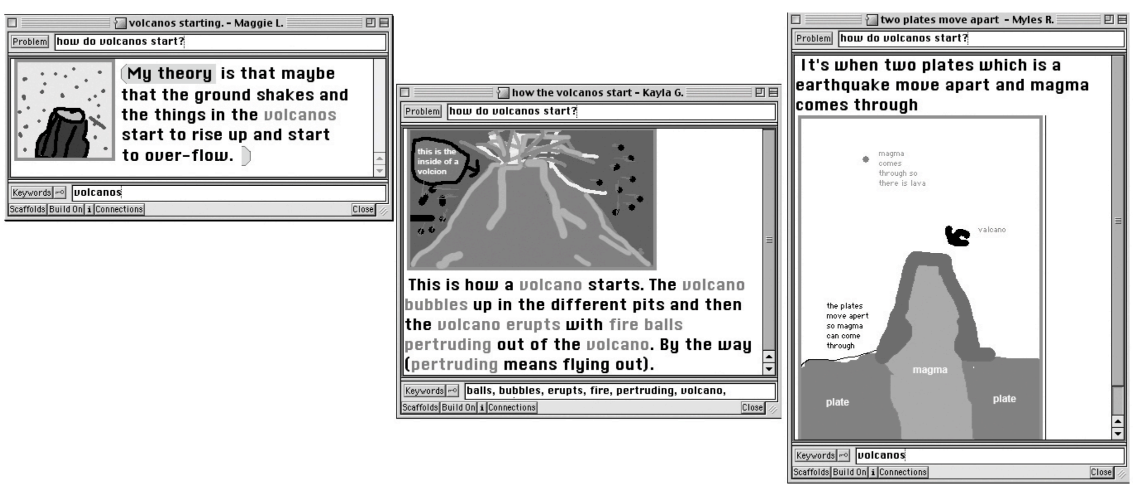

Figure 5: Idea improvement by grade 3 students studying volcanoes, as part of efforts to understand natural disasters.

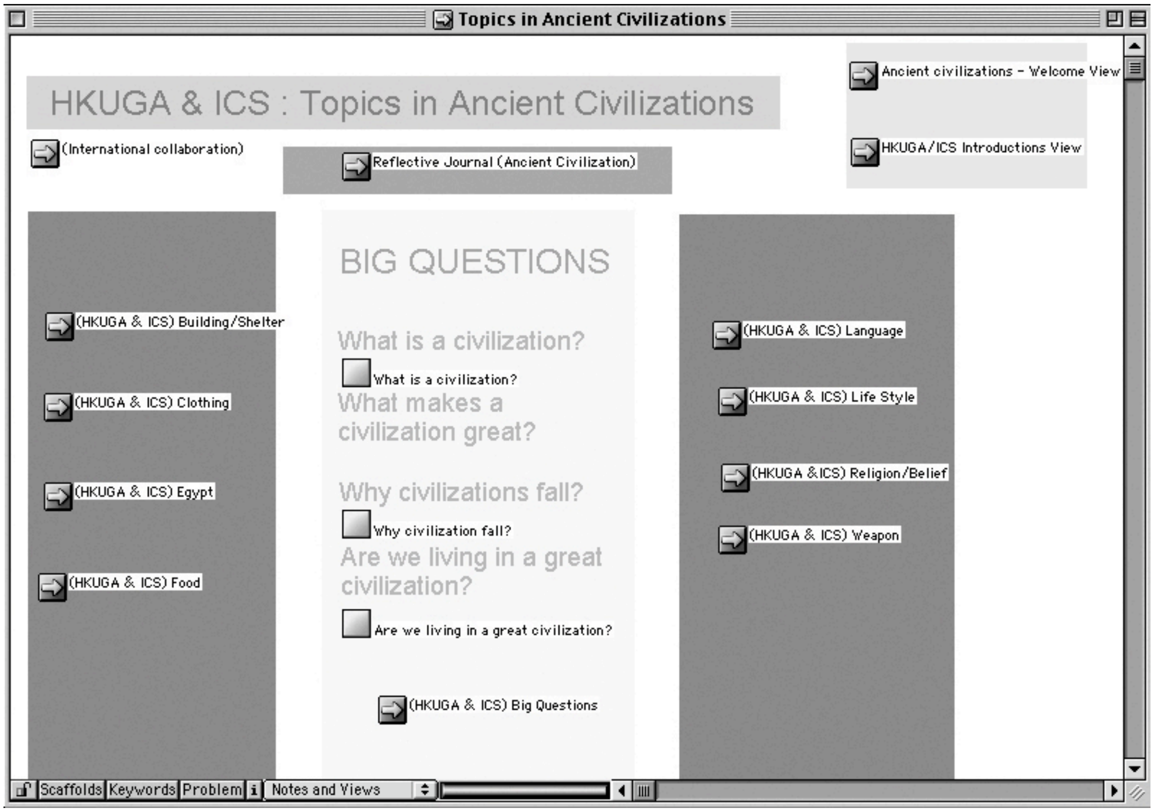

Figure 6: View-of-views by grade 6 students studying ancient civilizations.

Figure 3 shows a view created by Grade 1 students. It represents an overview of their work on pollution. The teacher reports, “This year in grade one we studied ecology as an overarching theme throughout the year …we read newspaper articles, books, and World Wildlife Federation publications … We frequently came across and discussed vocabulary such as pollution, oxygen, carbon dioxide, chemicals, pesticides, endangered, threatened, and so on.” Several students generated the same theory—that pollution is caused by laziness—and the rise-above note in the lower left (basic note icon with leaves underneath) is used to assemble those theories into one note. By Grade 3—see Figure 4—students are engaged in more complex rise-above activity, as indicated by rise-above notes throughout the view. The view itself represents an overview of the work of the class as a whole, with a section in the upper left titled “knowledge advances” providing an even higher-level summary, with links to related views. One of the related views is titled “volcanos.” Figure 5 shows several notes in that view, and efforts to explain volcanoes, starting with surface features and later their “problem of understanding” shifts to trying to figure out what happens below the surface. In Figure 6 we see a view-of-views created by Grade 5-6 Toronto students in collaboration with students in a Hong Kong public school affiliated with the Hong Kong University Graduates Association. They codesigned this view to identify their big questions and to organize their collaborative work.

Conclusion

In education, most of the 20th Century was occupied with efforts to shift from a didactic approach focused on the transmission of knowledge and skills to what is popularly called “active learning,” where the focus is on students’ interest-driven activities that are generative of knowledge and competence. We believe a shift of equal if not greater magnitude will come to dominate educational dialogue in the present century. The 20th Century shift has been aptly characterized by Stone (1996) as a shift from “instructivism” to “developmentalism,” for underlying the shift has been a strong belief in the natural disposition of children to do what is conducive to their personal development—in effect, to know better than the curriculum-makers what is best for them. Dispute over this proposition is by no means settled, but it is rendered moot by a societal shift that puts the emphasis on the ability of organizations and whole societies to create new knowledge and achieve new competencies. In this “knowledge age” context, it cannot be assumed that either the curriculum-makers or the individual students know what is best. The new challenge is initiating the young into a culture devoted to advancing the frontiers of knowledge on all sides, and helping them to find a constructive and personally satisfying role in that culture. The culture-transmission goals of liberal education and the more childcentered goals of developmentalism are not to be ignored, but they are to be realized within an educational environment that is itself an example of and at the same time a legitimate part of the emerging knowledge-creating culture (Smith, 2002). The driving force is not so much the individual interests of children as their desire to connect with what is most dynamic and meaningful in the surrounding society. That, fundamentally, is what knowledge-building pedagogy and knowledge-building technology aim to build upon.

The proof of knowledge building is in the community knowledge that is publicly produced by the students—in other words, in visible idea improvement achieved through the students' collective efforts. Although ascertaining that knowledge building has taken place requires digging into the content of Knowledge Forum databases and recordings of class interactions, it is usually apparent when something is seriously wrong. Pedagogy that is far off the mark will often manifest itself in a Knowledge Forum database that is full of redundancy, that is merely a repository of facts, or that presents a deluge of questions, opinions, or conjectures with no follow-up.

When knowledge building fails, it is usually because of a failure to deal with problems that are authentic for students and that elicit real ideas from them. Instead of connecting to the larger world of knowledge creation, the tasks or problems are mere exercises and are perceived by the students as such. At the deepest level, knowledge building can only succeed if teachers believe students are capable of it. This requires more than a belief that students can carry out actions similar to those in knowledge-creating organizations and disciplines. It requires a belief that students can deliberately create knowledge that is useful to their community in further knowledge building and that is a legitimate part of the civilization-wide effort to advance knowledge frontiers.

Acknowledgements

The authors wish to acknowledge the generous support of the Social Sciences and Humanities Research Council of Canada. We are indebted to the students, teachers, and principals of the Institute of Child Study and Rose Avenue Public School in Toronto and the entire Institute for Knowledge Innovation and Technology team (www.ikit.org), without whose contributions the work reported here would not have been possible. We are also indebted to Keith Sawyer for thoughtful input and help beyond the call of editorial duty.

References

American Association for the Advancement of Science. (1967). *Science: A Process Approach*. Washington, DC: American Association for the Advancement of Science, Commission on Science Education. Distributed by Xerox Corporation.

Anderson, J. R. (1980). *Cognitive psychology and its implications*. San Francisco: W.

Bazerman, C. (1985). Physicists reading physics: Schema-laden purposes and purpose-laden schema. *Written Communication*, *2*, 3-23.

Bereiter, C. (1985). Toward a solution of the learning paradox. *Review of Educational Research*, *55*, 201-226.

Bereiter, C. (1991). Implications of connectionism for thinking about rules. *Educational Researcher*, *20*, 10-16.

Bereiter, C. (1992). Referent-centered and problem-centered knowledge: Elements of an educational epistemology. *Interchange*, *23*, 337-362.

Bereiter, C. (1994). Implications of postmodernism for science, or, Science as progressive discourse. *Educational Psychologist*, *29*(1), 3-12.

Bereiter, C. (2002). *Education and mind in the knowledge age*. Mahwah, NJ: Lawrence Erlbaum Associates.

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. (2003). Learning to work creatively with knowledge. In E. D. Corte, L. Verschaffel, N. Entwistle, & J. V. Merriënboer (Eds.), *Powerful learning environments: Unravelling basic components and dimensions* (pp. 73-78). Oxford: Elsevier Science.

Bereiter, C., & Scardamalia, M. (in press). Models of teaching and instruction in the Knowledge Age. In P. A. Alexander and P. H. Winne (Eds.), *Handbook of educational psychology* (2nd ed.). Mahwah, NJ: Lawrence Erlbaum Associates.

Bereiter, C., Scardamalia, M., Cassells, C., & Hewitt, J. (1997). Postmodernism, knowledge building, and elementary science. *Elementary School Journal*, *97*, 329-340.

Brown, A. L., Day, J. D., & Jones, R. S. (1983). The development of plans for summarizing texts. *Child Development*, *54*, 968-979.

Caswell, B., & Bielaczyc, K. (2001). Knowledge Forum: Altering the relationship between students and scientific knowledge. *Education, Communication & Information, 1,* 281-305.

Catrambone, R., & Holyoak, K. J. (1989). Overcoming contextual limitations on problem-solving transfer. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*, 1147-1156.

Chi, M. T. H., Slotta, J. D., & deLeeuw, N. (1994). From things to processes: A theory of conceptual change for learning science concepts. *Learning and Instruction*, *4*, 27-43.

Coleman, E. B., Brown, A. L., & Rivkin, I. D. (1997). The effect of instructional explanations on learning from scientific texts. *Journal of the Learning Sciences, 6*, 347-365.

Dunbar, K. (1997). How scientists think: Online creativity and conceptual change in science. In T. B. Ward, S. M. Smith, & S. Vaid (Eds.), *Conceptual structures and processes: Emergence, discovery and change* (pp. 461-493). Washington, DC: American Psychological Association.

Fodor, J. A. (1980). Fixation of belief and concept acquisition. In M. Piattelli-Palmerini (Eds.), *Language and learning: The debate between Jean Piaget and Noam Chomsky* (pp. 142-149). Cambridge, MA: Harvard University Press.

Grossberg, S. (1997). Principles of cortical synchronization. *Behavioral and Brain Sciences, 20,* 689-690.

Lakatos, I. (1976). *Proofs and refutations : The logic of mathematical discovery*. New York: Cambridge University Press.

Messina, R., & Reeve, R. (2004). Knowledge building in elementary science. In K. Leithwood, P. McAdie, N. Bascia, & A. Rodrigue (Eds.), *Teaching for deep understanding: Towards the Ontario curriculum we need* (pp. 94-99). Toronto: Elementary Teachers' Federation of Ontario.

Molenaar, P. C. M., & van der Maas, H. L. J. (2000). Neural constructivism or self-organization? *Behavioral and Brain Sciences, 23*, 783

Ohlsson, S. (1991). *Young adults' understanding of evolutionary explanations: Preliminary observations* (Tech. Rep. to OERI No. University of Pittsburgh, Learning Research and Development Laboratory.

Pascual-Leone , J. (1980). Constructive problems for constructive theories: The current relevance of Piaget's work and a critique of information processing simulation psychology. In R. H. Kluwe & H. Spada (eds.), *Developmental models of thinking* (pp. 263-296). New York: Academic Press.

Petrowski, H. (1996). *Invention by design*. Cambridge, MA: Harvard University Press.

Phillips, W. A., & Singer, W. (1997). In search of common foundations for cortical computation. *Behavioral and Brain Sciences, 20,* 657-722.

Quartz, S. R. (1993) Neural networks, nativism, and the plausibility of constructivism. *Cognition, 48*, 223–42.

Ranney, M., & Schank, P. (1998). Toward an integration of the social and the scientific: Observing, modeling, and promoting the explanatory coherence of reasoning. In S. Reed & L. Miller (Eds.), *Connectionist models of social reasoning and social behavior* (pp. 245-274). Mahwah, NJ: Lawrence Erlbaum Associate.

Reed, B. (2001).Epistemic agency and the intellectual virtues. *Southern Journal of Philosophy, 39*, 507-526.

Rheinberger, H-J. (1997). *Toward history of epistemic things: Synthesizing proteins in the test tube*. Stanford, CA: Stanford University Press.

Sawyer, R. K. (2003). Emergence in creativity and development. In R. K. Sawyer, V.

John-Steiner, S. Moran, R. Sternberg, D. H. Feldman, M. Csikszentmihalyi, & J. Nakamura, *Creativity and development* (pp. 12–60). New York: Oxford.

Scardamalia, M. (2000). Can schools enter a Knowledge Society? In M. Selinger and J. Wynn (Eds.), *Educational technology and the impact on teaching and learning* (pp. 6-10) . Abingdon, Eng.: Research Machines.

Scardamalia, M. (2002). Collective cognitive responsibility for the advancement of knowledge. In B. Smith (Eds.), *Liberal education in a knowledge society* (pp. 76-98). Chicago: Open Court.

Scardamalia, M. (2003). Knowledge building environments: Extending the limits of the possible in education and knowledge work. In A. DiStefano, K. E. Rudestam, & R. Silverman (Eds.), *Encyclopedia of distributed learning* (pp. 269- 272). Thousand Oaks, CA: Sage Publications.

Scardamalia, M, & Bereiter, C. (1987). Knowledge telling and knowledge transforming in written composition. In S. Rosenberg (Ed.), *Advances in applied psycholinguistics: Vol. 2. Reading, writing, and language learning* (pp. 142-175). Cambridge: Cambridge University Press.

Scardamalia, M., & Bereiter, C. (1991). Higher levels of agency for children in knowledge-building: A challenge for the design of new knowledge media. *The Journal of the Learning Sciences*, *1*(1), 37-68.

Scardamalia, M., & Bereiter, C. (1996). Adaptation and understanding: A case for new cultures of schooling. In S. Vosniadou, E. DeCorte, R. Glaser, & H. Mandl (Eds.), *International perspectives on the design of technology-supported learning environments* (pp. 149-163). Mahwah, NJ: Erlbaum.

Scardamalia, M., & Bereiter, C. (2003). Knowledge building. In *Encyclopedia of education* (pp. 1370-1373). New York: Macmillan Reference.

Scardamalia, M., Bereiter, C., & Lamon, M. (1994). The CSILE project: Trying to bring the classroom into World 3. In K. McGilley (Eds.), *Classroom lessons: Integrating cognitive theory and classroom practice* (pp. 201-228). Cambridge, MA: MIT Press.

Scardamalia, M., Bereiter, C., & Steinbach, R. (1984). Teachability of reflective processes in written composition. *Cognitive Science*, *8*(2), 173-190.

Scardamalia, M., Bereiter, C., Hewitt, J., & Webb, J. (1996). 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* (pp. 44-64). Berlin: Springer-Verlag.

Scardamalia, M., Bereiter, C., McLean, R. S., Swallow, J., & Woodruff, E. (1989). Computer supported intentional learning environments. *Journal of Educational Computing Research*, *5*, 51-68.

Smith, B. (Ed.). (2002). *Liberal education in a knowledge society*. Chicago: Open Court.

Stehr, N. (1994). *Knowledge societies*. London: Sage Publications.

Sterelny, K. 2005. Externalism, epistemic artefacts and the extended mind. In (R. Schantz, ed) *The Externalist Challenge: New Studies on Cognition and Intentionality* .Berlin: de Gruyter.

Stone, J. E. (1996). Developmentalism: An obscure but pervasive restriction on educational improvement. *Education Policy Analysis Archives, 4*(8). Retrieved from http://olam.edu. asu.edu/epaa/v4n8.html

Woodruff, E., & Meyer, K. (1997). Explanations from intra- and intergroup discourse:

Students building knowledge in the science classroom. *Research in Science Education*, *27*(1), 25-39.