

Draft of chapter to appear in E. De Corte, L. Verschaffel, N. Entwistle, & J. van Merriënboer (Eds.), *Unravelling basic components and dimensions of powerful learning environments*. EARLI Advances in Learning and Instruction Series.

Revised: 16 February, 2003

Please do not quote without permission.

Learning to Work Creatively With Knowledge

Carl Bereiter and Marlene Scardamalia

OISE/University of Toronto

It is easy to produce a list of “soft skills” needed for work in the Knowledge Age. More than 140,000 Web pages contain the phrase, and a sampling among the first 4,000 indicates most of them advertise training—communication skills training, thinking skills training, human-relations skills training, and so on. Without presuming to evaluate these multitudinous efforts, we think a fair generalization is that they are relatively short-term interventions that do not so much develop new skills as sharpen or repurpose old ones. The ability to write a coherent paragraph, to utter a cogent statement, to make sense of numerical data, to mathematize and solve a quantitative problem, to read people’s intentions and moderate personal goals in the interests of team success—these and many other soft or semi-soft skills take years to develop. People look to the schools to do the long-term work. The soft skills are all very difficult to teach and their transferability to new situations must always be questioned. Curriculum standards and guidelines, in echoing the business world’s list of soft skills, tend to treat them as if effective means of teaching them were readily at hand. Even learning scientists often lend support to this unwarranted and facile optimism, claiming their approaches teach a range of 21st-century skills, similar to those in the curriculum guidelines and in the 140,000 Web pages.

We take it that the actual belief among most learning scientists is that only in the basic academic skill areas of reading, writing, and elementary mathematics do we know how to teach cognitive skills with fair confidence that they will transfer to a wide range of situations. Yet we must respect the demands of society for the schools to turn out people who, in addition to being proficient in these basic skills, will be prepared to learn new things, collaborate in the solution of novel problems, and produce innovations in areas that presently may not even exist. In the absence of tested methods, how do we do this? The time-honored and still the only promising way is *immersion*. If you do not have an effective way of *teaching* a foreign language, then place the students in an environment where that is the dominant language and trust that their natural adaptive

abilities will lead them to master the language. By the same reasoning, if we want students to acquire the skills needed to function in knowledge-based, innovation-driven organizations, we should place them in an environment where those skills are required in order for them to be part of what is going on. That, as we understand it, is the main reason the focus of this conference book is on powerful learning *environments* rather than on powerful instructional *methods*. For a range of very important objectives we have no effective instructional means and therefore must rely on environmental immersion.

Belief Mode and Design Mode

Both schooling and knowledge work are about knowledge, and so it would seem that they ought to have a great deal in common. However, there is a strong intuition, implicit in the various pronouncements about education for the 21st-century, that present-day schools deal with knowledge in quite different ways from the ways it is dealt with in the working world. We suggest that the difference is a difference between two modes that characterize our dealing with knowledge in all kinds of contexts; the problem is that schooling almost exclusively emphasizes one of these modes, whereas knowledge work in the real world mainly emphasizes the other. These two modes we call the *belief mode* and the *design mode*.

When in belief mode, we are concerned with what we and other people believe or ought to believe. Our response to ideas in this mode is to agree or disagree, to present arguments and evidence for or against, to express and try to resolve doubts. When in design mode, we are concerned with the usefulness, adequacy, improvability, and developmental potential of ideas. Switching back and forth between modes is common. The main work of a planning committee, for instance, is in design mode, but the planners may frequently pause to consider such belief mode issues as the reliability of certain data or the soundness of a certain assumption. However, it is possible for discourse to go on indefinitely within belief mode, never venturing into design. This is

characteristic of much academic discourse (in fact, that is what we imply when we say a question is “purely academic”).

Good educational programs generally do a good job of equipping students to think in the belief mode. They teach them to turn a critical eye on beliefs, to use evidence and logic, to resist propaganda, and they provide them with the background knowledge needed to evaluate truth claims. Bad educational programs also equip students to function in the belief mode, except that they do it badly. They turn out students who are unquestioning of authority or too fond of their own opinions (or both), who have little sense of how to apply evidence, and who lack sufficient knowledge to form intelligent judgments in most areas. Good or bad, however, educational programs in all their familiar variations operate almost exclusively in the belief mode *as far as ideas are concerned*.

Activity in the design mode is not absent from schools. It is to be found in crafts, dramatic productions, creative writing, and the increasingly ubiquitous “project.” In many of these the focus is on creating artifacts, but the artifacts are not *conceptual* artifacts (cf. Bereiter, 2002). They may embody ideas, reflect ideas, use ideas, but the artifacts are not themselves ideas. They are not theories, proofs, problem formulations, interpretations, or things of that sort. When ideas are presented for consideration, they are almost always presented in belief mode. The focus is on whether the idea is true or warranted. If experiments are conducted, their purpose is to validate, to provide an empirical basis for accepting the idea. Questions that would be asked in design mode—questions that would be asked in a real-world knowledge-based organization—are questions like the following:

What is this idea good for?

What does it do and fail to do?

How could it be improved?

If, in the school context, someone has an idea about a fund-raising event, those design mode questions will be raised. But if the idea under consideration is a theoretical idea—an idea like static force, natural selection, oxidation, or capital—it is treated quite differently. Students learn to treat such ideas as fixed entities, to be accepted or rejected and sometimes to be applied. They do not learn to treat them as improvable objects and to carry out the kind of design work with ideas that leads a knowledge-based society forward.

The essence of the design mode is idea improvement. This is obvious when the design is about physical machines, tools, or other artifacts. No one even thinks of designing the ultimate computer or refrigerator. One only thinks of improvement—advancing the state of the art. There is no ultimate computer or refrigerator, because with each advance new possibilities arise for further advances. But the same is true of the ideas that emerge in science, history, and other disciplines. Modern practitioners of those disciplines do not look for final answers but for continual idea improvement. Again, each advance in the state of the art opens up possibilities for further advance. Somehow, if the schools are to enculturate students into the Knowledge Age, they must introduce this dynamic of continual idea improvement. They must bring the design mode into the heart of the educational program instead of relegating it to extracurricular or peripheral activities.

Contemporary Efforts to Bring the Design Mode into the Academic Curriculum

In the remainder of this chapter we examine four constructivist educational approaches, considered as environments for initiating students into practices that characterize knowledge work in the design mode. Customarily, these approaches are considered from the standpoint of activity structures, and from this standpoint they appear as procedural variations on a common constructivist theme. Considered as environments, however, we find deeper differences that amount to quite different

kinds of experience in working with knowledge. Our conclusion is that each of the approaches offers a different and significant way of bringing design-mode activity into the academic curriculum, but that one of them, Knowledge Building, offers the possibility of integrating all the approaches into an overarching learning environment that provides fuller and more authentic immersion in the actual life of a knowledge society.

The four approaches are Learning by Design, as developed at Georgia Tech; Project-Based Science, as developed at the University of Michigan; Problem-Based Learning, as developed at Southern Illinois University; and Knowledge Building, as developed at the Ontario Institute for Studies in Education/University of Toronto. We attach institutional names to these movements, not because they are unique to those institutions but because the conceptual labels are easily appropriated and sometimes used indiscriminately. A Web search on any of these labels will bring up a number of examples that are inconsistent with the defining characteristics of the exemplary approaches. Henceforth we will capitalize the labels when we refer to the exemplary versions and will use lower case when we refer to the broader range of activities that use the same labels.

Although the four approaches have been used in various subjects, their common ground is science education, and so we will focus our comparisons in that area. They have all demonstrated effectiveness in teaching science content. Our purpose here, however, is to consider them as potential environments within which students may experience working with ideas in design mode.

As a framework for this examination of design-mode approaches to science learning, we propose a continuum that runs from context-limited to context-general work with ideas. In the adult world, much creative work with ideas is of a context-specific nature. Materials engineers working for a floor covering manufacturer, who may be developing a new bonding agent for floor tiles, have a very limited context

within which ideas are created, evaluated, and improved. It will be only a lucky accident if their inventions prove to have application elsewhere. The manager of a health care organization, who is trying to design a more productive way for the professional staff to interact, is working with ideas that are applicable within the context of a particular organization. In this case, however, there is more likelihood that ideas proving fruitful in one organization will have some generalization to other, related contexts, and so this instance of knowledge work would occupy a more intermediate position on the continuum. At the extreme context-general end of the continuum we have what is commonly known as “basic research”—efforts to develop theories of universal application.

School activities can similarly be ordered according to this continuum. As preparation for adult life, it is desirable for students to have experience in design-mode work with ideas across the whole continuum. However, it is in the nature of general education that its main concern is with knowledge of high generalizability. We want students to understand the nature of photosynthesis and how this differs from animal nutrition; we will have failed if all they have learned is how to sprout beans in a pan of dirt. Therefore it is significant that most of the approaches we shall examine appear to emphasize activities toward the context-limited end of the continuum. By default, that leaves the context-general end to be dealt with in belief mode. An important conclusion we shall urge is that this need not be the case—that students can engage in creative design work with fundamental ideas of wide generality.

Learning by Design™

In Learning by Design, as described by Holbrook and Kolodner (2000, p.),

Science learning is achieved through addressing a major design challenge (such as building a self-powered car that can go a certain distance over a certain terrain).... To address a challenge, class members develop designs, build prototypes, gather performance data and use other resources to provide

justification for refining their designs, and they iteratively investigate, redesign, test, and analyze the results of their ideas. They articulate their understanding of science concepts, first in terms of the concrete artifact which they have designed, then in transfer to similar artifacts or situations, and finally to abstract principles of science.

Design projects with this general intent are now fairly common, planning a trip to Mars being the most frequently encountered example in North American schools. However, in the approach taken by the Georgia Tech group, the students are challenged to design something that they can actually build and test—unlike those planning a Mars landing module. Moreover, the design challenges are planned so that faulty science will lead to performance failure.

The bridge between context-specific engineering problems and context-general science is a narrow one, however. The researchers found that teachers want to teach the science first and then introduce the design problems, whereas the intent is that the design problems should motivate inquiry into the underlying science. Kolodner (2002) offered the following example of scientific inquiry motivated by product design. The design task was to maximize the distance traveled by a toy car driven by air expelled from a balloon through a straw. The inquiry pursued by the students was to determine how variations in the length and diameter of the straw affected performance. It is not clear how this context-limited inquiry could bring students into contact with underlying scientific ideas—such as Newton's Third Law.

There are cases in which a design problem leads to engagement with deep scientific principles. They are cases in which the design must simulate or intervene in a natural process. Examples are designing an artificial organ (one of the early ventures in Learning by Design used this as a problem), designing a specific antibiotic, and building a rainmaking machine. By their nature, design problems of these kinds lie beyond the scope of what students could actually build and test. Consequently, activity methods in

general, and Learning by Design is no exception, depend on rather forced connections between the activity and basic ideas. The result, whether the ideas are taught first or afterward or in the course of the design project, is that the ideas are dealt with in the belief mode, in parallel with but not intrinsic to activity in the design mode.

Learning by Design has obvious relevance to creative knowledge work. Engineering and design play leading roles in modern life and students should gain experience in doing them, learn intelligent ways of going about them, see the relevance to them of different kinds of knowledge and skill. We only argue that one should not expect Learning by Design to provide direct engagement in the creation and improvement of underlying theoretical ideas.

Project-Based Science

As defined by Marx, Blumenfeld, Krajcik, & Soloway (1997, p. 341),

Project-based science focuses on student-designed inquiry that is organized by investigations to answer driving questions, includes collaboration among learners and others, the use of new technology, and the creation of authentic artifacts that represent student understanding.

Unfortunately, much of what is currently promoted and practiced under the label of “project-based learning” does not fit this definition but instead represents the traditional “project” or research report dressed up in modern technology. The crucial difference is in the phrase “investigations to answer driving questions.” The traditional project is defined by a topic, not a question—much less a “driving” question. This traditional approach is exemplified in a widely circulated video produced a few years ago by a computer company to demonstrate the great new things that digital media will make possible. It showed an elementary school pupil preparing a report on volcanoes, collecting material from the Web to use in a presentation. Then it showed the student at the front of the room giving his presentation, which ended with a view of a volcano erupting, shown on a screen that extended across the width of the classroom

and accompanied by a roaring sound that had classmates gasping in awe. Great educational experience if you are going to be a television producer. But great science education? No, probably a shade worse than what would come from a conventional report, where ideas would count for a little more and their absence could not so easily be hidden behind media glitz.

Examples of “driving questions” related to volcanoes are:

What causes volcanoes?

What is the risk that a volcano will erupt near here?

Are volcanoes related to earthquakes?

Why do new volcanic islands keep coming up in Hawaii?

Such questions could give rise to Project-Based Science as envisaged by the Michigan researchers. Ideally, such a Project-Based Science unit will include simulations with computers or physical models, so that students are not limited to authoritative sources for answers to the question. If directed toward sufficiently deep questions, Project Based Science can indeed engage students in creative work with important ideas.

Therein lies a shortcoming of Project Based Science, however. How does one get students engaged with sufficiently deep questions? The four driving questions listed above are in a roughly increasing order of depth. The first two questions are the ones most likely to occur spontaneously to students, and the first question, “What causes volcanoes?”, is a good starting point for inquiry. The last two questions are likely to arise only after students have made progress in understanding and become acquainted with the formation and location of volcanic islands. For the teacher to pose these questions at the outset would not be a good idea, because the students would not be in a position to recognize their significance. They are questions that should *emerge* as a consequence of progress in work with ideas. It is characteristic of successful work with ideas in design mode that the questions or problems evolve as well as the answers and

solutions and that there is a creative interplay between them. Project-based methods tend to curtail this iterative process.

The curtailment is in large part caused by another stipulation. The original question “drives” a process that is directed toward the production of a concrete artifact—a report, a performance, a model, a letter to the city council, or whatever. This provides closure for the project but, unfortunately, it also gives closure to the advancement of knowledge. In some instances, pursuit of the “driving question” gets completely derailed as students’ attention focuses on producing the “authentic artifact” (e.g., Yarnall & Kafai, 1995). Even in the more successful instances, however, the project structure discourages the progressive reformulation of questions and problems, for to do so sets back design and production of the “authentic artifact.”

Problem-Based Learning

Problem-Based Learning is often treated as synonymous with project-based learning. Although the two have more in common than a shared acronym, Problem-Based Learning grows out of a different tradition with a different focus. Originating in a medical school, Problem-Based Learning was designed to teach medical knowledge and skills by engaging students in solving problems similar to those they could expect to encounter in practice—that is, problems of diagnosis and treatment involving simulated cases. Students are expected to determine among themselves what information they need to solve the problem, dig out and share the information, and work together toward a satisfactory solution, with guidance but little direct help from the instructor. As typically employed in medical schools, problem-based work is run according to a tight schedule and fixed procedures, with only limited opportunity for iterative idea improvement; but these are not essential features of the approach and are not mentioned among the minimum requirements for Problem-Based Learning at the Problem-Based Learning Initiative’s website (http://www.pbli.org/pbl/generic_pbl.htm). Unlike Project-Based Science, Problem-

Based Learning is not focused on a tangible end product. The end product is a problem solution—a purely conceptual artifact. Thus iterative idea improvement is, at least in principle, something that Problem-Based Learning could promote. In medical education, the goal is to turn out doctors who can solve diagnostic problems instead of relying on rote procedures and typical cases. A similar goal can motivate its application to education in general.

When applied to school learning, however, Problem-Based Learning immediately runs into a problem with one of its main principles. Whereas in medical education it is possible to select cases that engage students with problems very similar to those they will encounter in later work, this is not a realistic possibility in general education. The closest schooling can come to solvable problems that all students will likely encounter in later life are home management and personal finance problems, but these are usually of little interest to children and adolescents. There is one major exception, however. Schooling can engage students in solving the kinds of problems they will encounter in later education. Given that most students in wealthy nations do pursue education beyond secondary school and that more would probably do so if they were better prepared, Problem-Based Learning that succeeded in equipping students to deal creatively with future educational challenges would deserve a central place in any school curriculum.

To serve this purpose, however, Problem-Based Learning must depart even farther from its medical school roots. Case-based problems are inherently context-limited. A case acquires general significance when it is considered in the context of a more fundamental inquiry. Deep inquiry may start with an intriguing case: In an example provided by Hunt and Minstrell (1994), inquiry starts with the problem of what happens to an object on a spring balance as the air is evacuated from around it. But attention quickly shifts to the students' explanatory ideas and to the testing and revision of these ideas through experiments that broaden the inquiry to general concerns about

gravity, the difference between weight and mass, and weight of the atmosphere. In this shift from context-limited to context-general, the nature of the problem itself undergoes transformation, and this is something that Problem-Based Learning has not been designed to address.

Knowledge Building

“Knowledge Building” may be defined simply as “creative work with ideas that really matter to the people doing the work” (Scardamalia & Bereiter, in press). It is not confined to education but applies to creative knowledge work of all kinds. Whether they are scientists working on an explanation of cell aging or first-graders working on an explanation of leaves changing color in the fall, knowledge builders engage in similar processes with a similar goal. That goal is to advance the frontiers of knowledge as they perceive them. “As they perceive them” is an important qualification when Knowledge Building is undertaken in educational contexts. Identifying frontiers and judging what constitutes an advance are essential parts of Knowledge Building, which students need to learn to carry out themselves, not depend on a teacher or a textbook to do for them.

As a constructivist approach, Knowledge Building shares many characteristics with the other constructivist approaches discussed earlier. There are, however, important differences in strategy. The following are some of the more salient ones that have evolved through a decade and a half of work by teachers and researchers. A fuller account is presented in Scardamalia (2002).

- *A focus on idea improvement.* Idea improvement is a primary and sustaining goal. Students are encouraged to put forth their own ideas early. From then on their task is to work together on improving those ideas, and to use the full range of available strategies and resources to improve them.
- *Problems versus questions.* Although problems are often expressed as questions, we have found that pursuing solutions to problems rather than answers to questions best encourages knowledge building. Answers have a certain finality to

them, whereas problem solutions are generally continually improvable. Whereas comparing answers to questions puts students into the belief mode, solutions to problems, including solutions to knowledge problems, can be carried out in design mode—judging what different solutions do and do not accomplish, what new problems a solution raises and what problems need to be solved in order to progress in solving the main problem. Knowledge Building pedagogy differs from Problem-Based Learning in that the preferred problems are ones of considerable generality.

- *Knowledge of value to the community.* In Knowledge Building, conceptual artifacts themselves—developed ideas—are an important product, used by community members—primarily as tools that enable further knowledge advances. Thus the progressive character of modern sciences and disciplines is also characteristic of knowledge building pedagogy. This does not mean that the students are expected to produce an original theory of gravitation to stand alongside Newton's. Rather, what they produce would likely be consistent with Newton but enriched by insights that made gravitational theory come alive for them and made it something they could apply to new problems of understanding. Considerable original scientific thought on the part of the students will have gone into such a product. (Centuries elapsed between Newton and Einstein, yet creative scientific work was being done throughout that period; so we should not expect students to do better than that in a school term.)
- *Emergent goals and products.* Students in Knowledge Building classrooms typically produce tangible or visible products—reports, multimedia presentations, playlets, demonstrations—but these are not predetermined and made the ultimate objective of a “project.” They may emerge at any point in the iterative knowledge building process. They may take various forms and serve various purposes, from highlighting a problem to disseminating results. It is important to note, however,

that in many instances there is no tangible product other than the computer record of the online work that was done in achieving the knowledge advances (Bereiter, et al., 1997).

- *Constructive use of authoritative sources.* When functioning in belief mode, contemporary educators have a good deal of trouble deciding how to treat authoritative sources such as textbooks and encyclopedias. Should the material be treated as truth, as a current best guess, as personal opinion, or as what some authoritative group has agreed to treat as true? Each way has its drawbacks, and so the choice tends to be made on ideological grounds. In a Knowledge Building design mode, however, all ideas are treated as improvable; ideas that have been the result of extensive work and development are treated with respect and judged according to what they can contribute to the group's current problem-solving effort. This does not eliminate problems of truth or validity, but it puts them in a context where they can be dealt with along with other difficulties. One does not expect to achieve perfect understanding any more than one expects to produce the perfect computer; one only expects to improve on the present state of things.

To educators acclimated to approaches like those previously discussed, Knowledge Building may sound arid and abstract, suitable perhaps for a graduate seminar but not for the schoolroom. However, Knowledge Building research has accumulated compelling demonstrations that school-age children, even the youngest, find joy in working with ideas.¹ Their ideas are as real to them as physical objects, and their interest intensified as they watch their ideas take new forms and engage others. Furthermore, the hands-on activities that enliven modern classrooms are present in Knowledge Building classrooms as well. What distinguishes them is that they are enlisted in the creation and improvement of theories and "big" ideas. Students in a Knowledge Building classroom may build and test paper airplanes, but they are not, as one project-based version of this activity has it, trying to isolate the variables that affect

how far a paper plane will fly. Instead, the students are trying to understand the physics of lift—what keeps airplanes aloft. Control of variables will come into play, as they create and improve their ideas, but the focus is on theoretical ideas, not variable-testing. By the time they are finished they may also understand water skiing and hydroplaning, how propellers work, and may have investigated whether Bernouli's principle applies under water.

All four of the approaches discussed here—along with virtually all serious approaches to education in the schools—aim to bring about understanding of the big ideas that make a knowledge society possible in the first place. Knowledge Building additionally aims at the fullest possible immersion in the work by which such ideas are created and improved.

Comparisons and Possibilities of Synthesis

The differences among the four approaches are often minimized, lumping them together as constructivist. Kolodner (2002), for instance, explicitly identifies Learning Science by Design as a variety of Project-Based Science, as defined by the Michigan group. From a Knowledge Building perspective we see an important distinction. Learning Science by Design, according to the description set out by Kolodner, lacks the main defining characteristic of Project-Based Science: It is not “organized by investigations to answer driving questions” but instead is organized around efforts to meet challenges in design of physical devices. In terms of the continuum that runs from context-limited to context-general, Learning by Design lies farthest toward the context-limited end. It is explicitly focused on building concrete things that actually work. Ideas enter the process insofar as they are relevant to producing a device that fulfills its intended function (Kolodner, 2002). Project-Based Science, in contrast, is explicitly concerned with finding answers to significant questions; this presumably places it farther along the continuum in the context-general direction. However, to the extent that projects settle into explanations limited to a particular context (a survey of local

water quality is a popular example) they lean in the context-limited direction. In Problem-Based Learning, the products are problem solutions, but the problems themselves are context-limited insofar as they focus on particular cases—diagnosing the illness of a particular real or simulated patient, devising a plan of action for a particular situation, explaining a particular observed phenomenon. Of the four approaches, only Knowledge Building explicitly supports design-mode work at the context-general end of the continuum—that is, design work directly aimed at creating and improving broadly significant theories, problem formulations, interpretations, and the like.

Considered as learning environments, each of the approaches offers a distinct kind of experience in working with ideas in design mode. An obvious strategy is to ensure that the curriculum engages students, at various times, in each of the four approaches. This has several drawbacks, however. In real life, creative knowledge work does not stay focused at one point on the context-limited to context-general continuum but moves back and forth according to needs and opportunities. Even the most nuts-and-bolts engineering projects can profit from a consideration of basic ideas and their implications, and it is rare for theoretical inquiry to be carried on in the absence of hands-on experimentation, model building, and the like. A curriculum that has students doing Learning by Design one month and Problem-Based Learning on an unrelated topic the next month does not provide the necessary flexibility in the pursuit of long-term goals.

A more serious drawback to an eclectic or mixed approach is that it does not provide for socializing students into a Knowledge Age culture, such as exists in leading university centers of research and innovation-oriented companies. A powerful learning environment, after all, should be something that students are *in* rather than just something that structures their activity for a certain period of time. It should, to a certain extent, foster a *way of life* that manifests itself throughout the school day and beyond. Accordingly, an environment for developing Knowledge Age soft skills ought

to be one in which creative work with ideas is, as Peter Drucker (1985, p. 151) put it, “part and parcel of the ordinary, the norm, if not routine.”

Of the four approaches, Knowledge Building is the least bound to particular activity structures and is explicitly defined as an approach to knowledge that “is not confined to particular occasions or subjects but pervades mental life—in and out of school” (Scardamalia, 2002, p. 81). Thus it provides a comprehensive framework for incorporating other approaches, while at the same time preserving a coherent emphasis. Within a knowledge building framework, design of physical devices, experiments, information-gathering and problem-solving efforts would all be marshaled to the basic knowledge building objective—to advance the frontiers of knowledge as these are perceived by the students. Toy building, contrived projects, and problems whose main purpose is to develop skills would give way to constructions and inquiries designed by the students to advance their ideas and to make worthwhile contributions to the “state of the art” as it exists in the classroom.

A comprehensive knowledge building environment would provide a means of initiating students into a knowledge-creating culture—to make them feel a part of humankind’s long-term effort to understand their world and gain some control over their destiny. Knowledge would not be seen as something handed down to them from dead White males. Rather, they would look on those dead White males—and other intellectual forbears of different race and gender—as fellow workers whose work they are carrying forward. The Knowledge Society, as it is taking shape today, seems headed toward a very sharp separation between those who are in it and those who, whether they live a continent apart or on the same street, are on the outside looking in. A knowledge building environment should provide all students an opportunity to be on the inside looking out.

Technology for Knowledge Building

All of the approaches we have been discussing feature collaboration, and all make some use of ICT (information and communication technology) to support the collaboration. However, only in Knowledge Building is a software environment the principal environment in which the collaborative work goes on. In Learning by Design, the principal environment is the physical workbench where something is being built; the software environment is an adjunct for sharing ideas and information and developing plans. In Project-Based Science, the principal environment is that in which data are being collected. It could be a stream where water samples are gathered; it could be a laboratory where experiments are run. The principal software is that used in data collection and analysis, but with a discussion environment serving again as an adjunct. In Problem-Based Learning, the principal environment is the tutorial room. ICT figures primarily as a means for information gathering and for communication among participants between tutorial sessions. In Knowledge Building, however, a software environment—Knowledge Forum® or a derivative of it—is the principal environment in which work with ideas goes on. It is where ideas are set forth, discussed, revised, organized, combined, and so on. If, as in Project-Based Science, real-world or experimental data are collected, they are put to use by being brought into the work in Knowledge Forum. Other learning technologies, such as simulations and video recordings, are integrated as closely as possible into the Knowledge Forum work, preferably by running inside Knowledge Forum notes.

Discussion is, of course, central to collaborative work with ideas, but if a software environment is to serve as a workspace for knowledge building it must provide for more than the usual concatenation of discussion threads. If you have ever tried to do serious collaborative work in a conventional discussion, chat, or bulletin board environment you will have experienced how frustrating it is. It is all one thing after another, with no way to bring about integration and coherence. The essence of creative (as distinct from argumentative) work with ideas is making connections. But current

designs of collaborative learning environments do not take account of this. Instead, they represent variations on what is essentially a message-passing model. This model may be fine for the opinion-stating and question-answer dialogue that prevails in public online forums, but it is grossly inadequate for what we have called “progressive discourse”—discourse that gets somewhere, advances on a problem, produces a conceptual artifact.

A Knowledge Building environment needs multiple ways of representing and organizing ideas and flexible ways to link them. It should provide not only for subordination (which threaded discourse typically provides) but also for superordination—a new idea subsuming or integrating previously recorded ideas. An elementary requirement for idea improvement is that it should be possible to revise a previous formulation—something that message-based systems generally forbid. The technology needs to do more than that, especially if it is to be used for educational purposes (cf. Scardamalia, 2002), but our purpose here is not to lay out design specifications or to describe Knowledge Forum in detail. Instead, it is to argue that technology intended to develop skills for the knowledge age needs to be designed on the basis of some articulated conception of what it means to create and work with knowledge.

Conclusion

Alfred North Whitehead (1929, p. 83) said that education must equip students with the ideas and the capacities that would enable them “to appreciate the current thought of their epoch.” In our epoch, which is coming to be known as the Knowledge Age, it is impossible to appreciate the “current thought” and to take part in it unless one thinks of ideas as serving purposes and as being continually improvable so as to serve those purposes better. The traditional concerns with truth and warranted belief are not abandoned, but they are subordinated to the mission of advancing the frontiers of understanding and efficacy on all fronts. A powerful learning environment for the

present age, accordingly, is an environment that immerses students in the effort to advance the frontiers of knowledge as they perceive them.

In a series of reports from classrooms organized around this Knowledge Building effort, we and our co-workers have provided evidence that students at all ages can in fact do creative work with knowledge, while at the same time performing well according to the usual criteria of school achievement (e.g., Bereiter, et al., 1997; Hewitt, 2002; Oshima, 1977; Scardamalia, 2002; Scardamalia, Bereiter, & Lamon, 1994). The creative work students do is not just creative work with concrete artifacts and media, although they do that, too. It is creative work with ideas that are at the heart of the sciences, scholarly disciplines, and knowledge-creating enterprises.

To extend Whitehead's statement, we may say that you cannot appreciate the current thought of this epoch unless you are an active participant in advancing it, because continual advancement of knowledge in all spheres is what constitutes the current thought of the Knowledge Age. The schools, even those most attuned to the new informational media and to the ideas of constructivism and active learning, have by this standard not yet entered the Knowledge Age. In order for them to enter, we need a radical innovation in educational experience that brings ideas and idea improvement to the center.

References

- Bereiter, C. (2002). *Education and mind in the knowledge age*. 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.
- Drucker, P. (1985). *Innovation and entrepreneurship: Practice and principles*. New York: Harper and Row.
- Hewitt, J. (2002). □From a focus on tasks to a focus on understanding: The cultural transformation of a Toronto classroom. □In T. Koschmann, R. Hall, & N. Miyake

- (Eds.) *Computer Supported Cooperative Learning Volume 2: Carrying forward the conversation* (pp. 11-41). Mahwah, New Jersey: Lawrence Erlbaum Associates.
- Hmelo, C. E., & Evensen, D. H. Introduction. In D. H. Evensen and C.E. Hmelo (Eds.), *Problem-Based Learning, A Research Perspective on Learning Interactions* (pp.185-195). Mahwah, NJ: Lawrence Erlbaum Associates.
- Holbrook, J., & Kolodner, J. L. (2000). Scaffolding the development of an inquiry-based (science) classroom. In B. Fishman & S. O'Connor-Divelbiss (Eds.), *Fourth International Conference of the Learning Sciences* (pp. 221-227). Mahwah, NJ: Lawrence Erlbaum Associates.
- Hunt, E., & Minstrell, J. (1994). A cognitive approach to the teaching of physics. In K. McGilley (Eds.), *Classroom lessons: Integrating cognitive theory and classroom practice*. (pp. 51-74). Cambridge, MA: MIT Press.
- Kolodner, J. L. (2002). Learning by Design™: Interations of design challenges for better learning of science skills. *Cognitive Studies*, 9(3), 338-350.
- Kolodner, J.L., Crismond, D., Gray J., Holbrook, J., & Puntambekar, S. (1998). Learning by Design from Theory to Practice. Proceedings Third International Conference of the Learning Sciences '98, pp.16 - 22.
- Marx, R. W., Blumenfeld, P. C., Krajcik, J. S., & Soloway, E. (1997). Enacting project-based science. *Elementary School Journal*, 97, 341-358.
- Oshima, J. (1997). Students' construction of scientific explanations in a collaborative hyper-media learning environment. In N. M. R. Hall N. Enyedy (Ed.), *Computer Support for Collaborative Learning '97*, . Toronto:
- 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., & Bereiter, C. (in press). Knowledge building. In *Encyclopedia of Education*, Second Edition. 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.
- Smith Lea, N. & Scardamalia, M. (1997). Progressive curriculum and knowledge building uses of networked multimedia resources. Paper presented at the Annual Meeting of the American Educational Association, Chicago, March 24, 1997 Session 5.13. Available at http://csile.oise.utoronto.ca/abstracts/aera_pcn_kb_uses/
- Whitehead, A. N. (1929). *The aims of education*. New York: Macmillan.
- Yarnall, L., & Kafai, Y. (1996, April). *Issues in project-based science activities: Children's constructions of ocean software games*. Paper presented at the annual meeting of the American Educational Research Association. Available at http://www.gse.ucla.edu/kafai/Paper_Kafai%2FYarnall.html.

¹ References to this research, including many online documents, are available at www.ikit.org.

Background ARC is committed to developing their creative learning offer for young people. To date the prevailing barrier to this expansion has been a significant lack of high quality and diverse artists living and working in the Tees Valley. At times this issue has prevented ARC from being able to deliver creative learning activity that is both relevant to ARC's artistic programme and in demand; examples include young people requesting music production, spoken word, film, set design and breakdance. In response to this we developed How to Work Creatively with Young People, a professional development programme. Knowledge workers are workers whose main capital is knowledge. Examples include programmers, physicians, pharmacists, architects, engineers, scientists, design thinkers, public accountants, lawyers, and academics, and any other white-collar workers, whose line of work requires one to "think for a living". Knowledge work can be differentiated from other forms of work by its emphasis on "non-routine" problem solving that requires a combination of convergent and divergent thinking. But despite the amount