Interchange, Vol. 17, No. 2 (Summer 1986), 10–19

Educational Relevance of the Study of Expertise

Carl Bereiter and Marlene Scardamalia
Centre for Applied Cognitive Science, OISE

Certain kinds of knowledge come very easily and naturally to human beings. The most notable example is language, but recent research indicates that elementary mathematical knowledge also appears universally and at an early age (Gelman & Gallistel, 1978). Human beings are further distinguished, however, by their ability to acquire expertise—high levels of knowledge and skill that do not come about naturally but require a special and sustained investment in learning. Although civilization may be based on natural human abilities to communicate, calculate, and plan, modern civilizations also depend on enormously varied expertise—in manual skills, in science, in management, and of course in teaching, which is a means of propagating the other kinds of expertise.

Cognitive scientists have been active in studying both naturally developing abilities and expertise. Obviously, both are important in understanding the human mind. But it should be equally obvious that expertise is the more educationally relevant since it presents the more serious challenges to education. Strangely, many educators take the opposite position. They take natural learning, particularly language development, as the model for what all human learning is or should be. But if all learning were like language development, there would be no need for institutions dedicated to education. Some educators, of course, pursue the natural learning argument to this extreme conclusion; but the fact that more or less formal educational procedures have been independently invented by societies around the globe strongly suggests that there is some inherent need for them. That need, we would argue, is precisely the need for expertise—the need for knowledge and skills that do not arise naturally through experience.

In this article, we shall try to summarize the findings on expertise that are of educational interest. Some common characteristics distinguish experts from nonexperts (or "novices," as they are called in the cognitive science trade) in a variety of domains. These common characteristics are worth attending to at two levels. At the first level, they suggest something about what the course of instruction should be for developing expertise in any particular area. Even though elementary and secondary school educators do not usually think of their work as producing experts, they should be moving students in the direction of expertise, and so it is important to understand what that direction is. Our own research, however, is concerned with characteristics of expertise at a second level. This is the level that may be described as "being expert at becoming an expert." We all enter a new domain as novices, but some people know better than others how to go about becoming expert in the new domain. Acquiring this second level of expertise, we argue, is not simply a matter of "learning how to learn." It is a matter of learning how to learn the kinds of things that will lead to expertise rather than learning the kinds of things that will restrict one to being a novice.

What Experts Have in Common

Expertise, as we shall use the term here, is relative. Some of the earliest and most influential cognitive research on expertise compared chess grand masters with chess masters. Both groups would be considered expert by ordinary standards, yet a major difference was found between members of the two groups in their ability to remember chessboard configurations. A later study compared children who were skilled chess players with adults who were not and found the same kind of difference favoring the more expert group. Differences can be found between children who are good and poor mathematical problem solvers, and the same kinds of differences can be found between physicists and undergraduate physics students and between medical specialists and residents who are beginning their specialization. Therefore, it is reasonable to talk about characteristics that distinguish the relative expert from the relative novice all along a scale ranging from very low to very high levels of attainment.

The following are major characteristics that distinguish relative experts from relative novices in any domain:

Complexity of skills. By definition, experts are better at what they do than are nonexperts. But this is not simply a matter of greater proficiency in executing the same process. Typically, the process has a different and more complex structure for the expert. The expert's process responds to more kinds of information. Whereas the novice marksman may be able only to point at the target and adjust crudely for its direction of movement, the skilled marksman can additionally take account of windage and the target's distance and speed of movement. The expert can also exert voluntary, strategic control over parts of the process that are carried out automatically by the novice. The skilled singer or speaker, for instance, exerts kinds of control over breath, posture, and shape of the oral cavity that are unknown to the novice. As a result of this more complex structure of control over performance, the expert is able to set and pursue goals that the novice cannot pursue at all, let alone pursue successfully.

Amount of knowledge. That experts know more about their areas than novices do might seem too obvious to be worth mentioning, but educators have tended to undervalue this aspect of expertise. One continually encounters the pronouncement that students' heads should not be filled with information but that instead students should be taught how to find the information they need. This is one of the silliest of persisting ideas in education. A related but more plausible educational belief is that students should not be taught isolated facts; rather, they should be taught central concepts that render facts meaningful. Research by Voss and others has shown, however, that the ability to grasp and retain new information, to link it meaningfully to central ideas, depends greatly on the extent of one's existing knowledge. People with little knowledge in a domain are almost compelled to deal with knowledge as isolated fragments, only vaguely relatable to higher level principles. The more you know, the more you are able to learn and the more meaningful the learning is likely to be.

Knowledge structure. Experts not only know more, but their knowledge is organized in more coherent and usable ways. Novices tend to have what may be described as a "shallow" knowledge structure—a few general ideas and a lot of details connected to the general ideas but not to each other. Experts, on the other hand, have deep or multilevel knowledge structures, with many connections between and within levels. The result is knowledge that can be put to a variety of uses, whereas novice knowledge tends to be what A. N. Whitehead referred to as "inert" knowledge.

Problem representations. To qualify as an expert rather than merely a skilled performer, one must display an ability to solve novel problems in one's domain. It has therefore seemed reasonable to suppose that experts have superior general strategies for problem solving. The extent to which this is true remains unsettled at present. What has been established in a variety of domains, however, is that experts represent or perceive problems differently from novices. Novices are mainly influenced by the concrete content of problems—whether a problem is about falling objects or objects rolling along surfaces, whether it is about money changing hands or about quantities of material. Experts, by contrast, respond to the abstract structure of problems; for instance, a money problem and a materials problem might be seen as similar if both involved separating a quantity into equal parts. A notable characteristic of experts is that they have in memory a large number of abstract problem types; they can often solve problems quickly and effortlessly by recognizing the problem as one of a type and then applying the procedures they have learned for problems of that type.[1] Thus, ironically, experts have less need of general problem-solving strategies than do novices, because to experts few problems are really novel.

How Do Novices Become Experts?

As the characteristics separating experts from novices come to be understood, it becomes increasingly puzzling how it is that anyone ever manages to become an expert. If you start out with scanty and shallow knowledge, what is there to connect new knowledge to? It would seem that most new information either would not be grasped in the first place or would be quickly forgotten or else assimilated to the existing shallow knowledge structure. So how could expert understanding gain a foothold? And if, at the outset, you can only deal with the superficial aspects of phenomena, how will you ever start to acquire the repertoire of abstract problem types that make the expert a powerful problem solver?

This is not just a chicken-and-egg type of paradox. It is clear that the early stages of acquiring expertise are indeed often long and difficult. Beginners in a discipline do understand very little and forget most of what they learn. If it is understood that this is natural and largely inevitable, we should be able to take a more tolerant view of both students and teachers—for it is often teachers who take the blame when students reveal novice-like knowledge and behavior.

Nevertheless, the world is full of experts, and all of them started out as novices.

When we say it is natural for human beings to acquire expertise we imply that people are naturally disposed to struggle through the early stages of learning and to overcome the limitations that their initial forms of knowledge and skill impose on them. This is more than just a matter of practice making perfect. A phenomenon that has been observed even with young children is that once they have mastered a procedure they do not simply go on practising it and gradually honing their skill. Instead, once the procedure becomes automatic enough that it no longer requires all their attention in order to function, children begin turning attention to the procedure itself, revising or transforming it into a more sophisticated one (Karmiloff-Smith, 1979). Without instruction, young children have been observed to devise a series of increasingly sophisticated addition algorithms, to invent grammatical markers not actually present in their native language, and to devise written forms that gradually approximate phonetic spelling. Similarly, one can observe in children's incessant "why" questions an effort to transform fragmentary information into causally related knowledge.

But people differ a great deal in the extent to which they carry out this expertise-building work. At one extreme are people who appear to be in the process of becoming experts at everything. This does not mean, of course, that they achieve high levels of expertise in everything; rather, it means that in many different areas they are continually trying to get beneath the surface of the phenomena they are dealing with and to build up a coherent and usable body of knowledge. Developing expertise seems to be their normal way of behaving in most situations. At the other extreme are people for whom learning appears always to be incidental. Something else has to motivate them—curiosity, a desire to make money, the need to solve a practical problem, a yearning to enjoy homemade beer—and learning comes about as a by-product of the ensuing activity.[2]

Such differences in disposition to develop expertise may be noted even among schoolchildren. In schools, however, it is common to attribute such differences to interest. There is no doubt that interest is an important factor in carrying people through the travail of becoming an expert. But we also believe that some people are much more able than others to *get themselves interested* in anything they set their minds to. Interest is a major problem in schools precisely because most students seem to have no control over their interest. Interest is something that happens to them, not something they bring about through deliberate mental activity. For the more active type of learners, on the other hand, interest is better thought of as something they have to *invest* rather than as something that already exists in the subject matter awaiting their response.

Expertise in Common Academic Skills

Although elementary and secondary schooling are not concerned with producing specialized experts, they are concerned with producing people who can read and write well and make effective use of mathematics. In these basic skills, differences can be observed that turn out, on deeper investigation, to be similar to the general expert-novice differences previously discussed. The study of expert-novice differences, we believe, sheds important light on what is required for upgrading literacy and numeracy.

Expertise in Writing

Of the traditional three R's, writing is the one in which expertise is most easily recognizable. It is obvious that there are people whose skill in writing is of an altogether different order from that of the ordinary literate person. Although most people who use writing in their occupations develop adequate skill in the mechanics of writing, there is a widespread impression that most people are inexpert writers. We have spent much of the past decade studying the nature of competence in written composition (Scardamalia & Bereiter, 1985). It turns out that the general differences discussed previously between relative experts and relative novices show up clearly when we compare the average grade 6 or 10 student with the average university English major. The structure of the writing process is much simpler for the relative novice. It consists essentially of telling what one knows about a designated topic, while conforming to the requirements of a text type—argument, factual essay, or the like. The relatively more expert writers, on the other hand, construct much more complex representations of the writing task, plan and revise plans, check results against goals, and consider a variety of kinds of information. Like other novices, novice writers deal with surface phenomena. Their revisions of texts tend to be cosmetic, whereas the most expert writers often make more basic changes in form and content, changes that involve rethinking plans and purposes.

An issue of considerable educational importance has come to light through this research. The novice way of writing—what we call the "knowledge-telling" strategy—turns out to be a highly efficient strategy for handling routine writing assignments and to be capable of considerable refinement. Accordingly, there is little incentive for students to develop more expert-like strategies. In the short term, at least, the gains are greater from simply becoming a more proficient novice. Perhaps that is why so many people remain poor writers in spite of extensive practice. The long-term losses go deeper than that, however. By sticking to the surface-level approach to writing, the lifelong novice misses out on the gains in knowledge and understanding that more expert writers obtain from the composing process itself.

Expertise in Reading

As with writing, most people who read a lot become proficient at it, but this does not necessarily mean that they are expert readers. Much of our own time and that of our colleagues seems to go into dealing with failures of reading comprehension on the part of students, reviewers, colleagues, critics, and even occasionally on the part of ourselves. We are not talking about disputable interpretations but about plain failures to notice what was said.

The commonest failure is to believe that the text says what it was expected to say. In line with our earlier description of expert-novice differences, this failure usually involves applying to the text a problem representation simpler than is required to grasp the intended meaning. This is commonly referred to as "shallow" reading of the text, and it again relates to the novice tendency to deal with surface features. Studies of novice reading strategies, some of which we have conducted ourselves, reveal a process that is (not surprisingly) similar to the novice composing process (Scardamalia & Bereiter, 1984). Text statements are considered one at a time, within a framework provided by the text type; they are retained in memory if they are sufficiently salient and rejected as untrue or unimportant if they are not. If retained, the statements are related to the topic of discourse but not much related to one another. Thus they take a place in the shallow knowledge structure we described as characteristic of novices.

A more expert approach to reading can occasionally be found, however, even among school-age students. The outstanding characteristic of this more expert approach is a sustained effort to figure out what the text is trying to say. This effort is revealed by having readers think aloud while they read. The more expert readers tend to summarize text meaning as they go along, to look back in the text to recover needed information, to search actively for relationships (for effects to go with causes and vice versa, for examples, definitions, and the like), and to formulate comprehension difficulties as specific problems that they try to solve. Inexpert readers do much less of this and instead tend to comment on each item of text information separately. These differences suggest that for the expert reading is a way of reconstructing and building more complex personal knowledge, whereas for the nonexpert reading is a way of adding isolated facts without substantial alteration in what is already believed.

Expertise in Using Elementary Mathematics

Using mathematics means constructing a mathematical model of a situation and then carrying out mathematical operations, the results of which can be applied via the model to the real-life situation. For instance:

> You have just opened a 3-litre can of olive oil. In order to keep it from turning rancid before it is all used up, you decide to decant and seal it into wine bottles. How many 750 ml wine bottles will you need? You reason that 750 ml is ¾ litre, and so the problem is "How much is 3 divided by ¾?" You know that 3 divided by ¾ is the same as 3 times 4/3, which is 4. So you conclude that you will need 4 wine bottles to hold the olive oil.

In the above example, "3 divided by ¾" is the mathematical model you construct of the situation and "3 times 4/3" is an alternative model mathematically equivalent to the first model, although it does not fit the real-life situation the way the first one does. Nevertheless, you confidently carry out the calculation using the second model and apply the result to the real-life situation. This is quite a complex bit of model using, which employs no mathematics above the level normally taught by grade 6.

Many real-life situations require more advanced mathematics, but a large part of the world's work can be done with no mathematics beyond what is normally taught in high school. In the ability to construct and apply models based on such mathematics, however, extreme variations in expertise can be observed. The mathematical portion of the Scholastic Aptitude Test, for instance, requires no knowledge beyond elementary algebra and geometry, yet it can be used to identify both the mathematically gifted, at one extreme, and, at the other, substantial numbers of students who score below chance level, meaning that they would do better by blind guessing.

Inexpert users of mathematics are most clearly distinguished by the tendency to deal with concrete or surface aspects of problems and not to grasp their abstract form. Instead of constructing mathematical models of problem situations, they try to apply rules that directly translate problem elements into mathematical operations. This can lead to bizarre mistakes such as solving the olive oil problem by dividing 3 into 750, thereby obtaining an answer of 250 bottles. Since the problem talks about pouring something into different containers, they conclude that the operation is division. They look for numbers to put into a division algorithm and naturally select the smaller number as the divisor and the larger number as the dividend.

Mathematics educators have been concerned for many years about the tendency of students to lose hold of reality when they enter the domain of mathematics. It would seem that a student with even the crudest sense of the magnitude of 3 litres and of the size of wine bottles ought to realize that 250 is an absurd answer to the number of wine bottles required to hold 3 litres of oil. Yet such absurdities are commonplace. To keep students in closer touch with reality, educators have devised many ingenious concrete representations of symbolic operations. One of the commonest has been blocks representing units, tens, hundreds, and so on, which can be used to enact concretely the symbolic operations of carrying and borrowing used in arithmetic. These reality-linking devices have never seemed to work as well in practice as in theory, however.

Recently, Resnick and her colleagues have begun investigating what goes on in young students' minds when they deal with concrete analogues of symbolic operations in mathematics (Resnick & Neches, 1984). A difference among students appears that is very similar to the other differences we have observed between relative experts and relative novices. Some students continually strive to connect what they are doing symbolically to corresponding real-world changes in quantity. Others dutifully carry out the matching operations with numbers and with blocks, but they make no effort to construct a coherent interpretation of what is happening. Thus, even though the teacher has carefully set up a situation in which the connections between symbols and reality are evident, the connections are not constructed in the mind of the learner—unless the learner makes an effort to construct them. But it seems that it is just such constructive effort, such model building, that constitutes expertise in the uses of elementary mathematics. Thus, mathematics educators have succeeded in presenting ideas in ways that are useful to the relative expert, but they have not found ways to foster the expertise itself.

Expertise in Becoming an Expert

Everyone knows that some young people are better writers, readers, or mathematical problem solvers than others, and nothing much is gained by labelling these children "relative experts." In the preceding brief analysis we have tried to advance two ideas that go beyond mere labelling. One idea is that the relative experts are not merely better at doing the same things that others do; they do things differently, and the same differences appear in various domains. The other idea is that the relative experts are distinguished by efforts they invest in becoming expert. It is this latter point that we want to develop in the remainder of this paper.

In the three areas that we have considered—writing, reading, and mathematics—expertise is seen to arise from doing more than the immediate learning situation requires. But it is not doing more in any of the conventional senses of academic performance. That is, it does not involve doing extra tasks or putting extra care into doing things correctly. There is not, in fact, anything very visible to it at all, and that is partly why it has taken cognitive research to reveal what expertise consists of.

In cognitive domains, the work of becoming an expert is concerned with the mental structures that lie behind performance. Novices, by contrast, deal with the overt requirements of tasks. In writing, the overt task is usually to produce a piece of text presenting information on a certain topic and conforming to certain stated or conventional requirements. Thinking-aloud transcripts of novice writers show them to be occupied almost entirely with these overt requirements. When they get stuck, their typical recourse is to reread the assignment for further clues about what to do. More expert writers, on the other hand, construct goals of their own that elaborate on the explicit requirements, and when they get stuck they typically re-examine their own goals. The resulting composition may look like simply a very good job of carrying out the explicit assignment, but a large part of the mental work went into constructing and pursuing goals that were not given in the assignment. The main value of this extra work for the student, we suggest, does not lie in the immediate literary product but in the expertise that is being built up through this behind-the-scenes effort.

The overt point of reading, in most school situations, is to store up information to be used for answering questions. We find it alarming that the novice reading strategy described in the preceding section appears to be tailored to this very task. A shallow mental structure in which details are stored under topical headings seems to be ideal for quick retrieval in response to unpredictable questions—a function that is valuable in school and hardly anywhere else. To become an expert reader, one must go well beyond this kind of task and focus instead on integrating text information with one's knowledge of the world. It seems likely that schooling, with its continual emphasis on the answering of questions as the immediate goal of learning, may be defeating the development of expertise in all but the most hardy students.

In mathematical problem solving, the overt task is of course to produce correct answers to problems. It is not so obvious what there might be above and beyond this task to occupy the more expert student. The most visible clue to what else there is has been that the most capable problem solvers tend to pause and reflect *after* they have solved a problem. What seems to be going on is that the students are trying during these pauses to extract general knowledge from the particular experience. The general knowledge in this case would most likely be knowledge of problem types—abstract schemas that, as we noted previously, distinguish experts in all domains and figure prominently in their superior ability to solve problems.

In each instance, we find the expert-like student carrying out mental activities that have future expertise as their most important consequence. This does not necessarily mean that the student is conscious of such a connection, but it does imply that the student has somehow acquired a commitment to the autonomous pursuit of learning that is sustained in spite of lack of support from the immediate situation.

Practical Research Directions

In this paper we have focussed on research that seeks to *understand* expert-novice differences. Several lines of applied research at OISE are concerned with designing educational approaches to *promote* expertise and expert-like abilities in students. These lines of research have complex rationales and research foundations, and so they cannot be described but only mentioned in this brief article.

Case (1978) has developed a theory and technology of instruction based on analysis of expert and novice procedures and design of successive stages for transforming the novice's procedures into more expert-like ones. In current work, Case and Sandieson have modified the approach so that the focus is on relative expertise, specifically on identifying the procedures of students who are only one stage more advanced than those who are to be taught.

Our own applied research, as previously suggested, is concerned with promoting self-directed efforts that lead to expertise. In the area of writing, this has meant developing ways to help students move away from straight "knowledge telling" to more reflective processes in composing (Scardamalia & Bereiter, 1985). In the Intentional Learning Project, we are trying to find ways both to foster active learning and to alter the school practices that encourage passive learning. The same goals underly the Computer Supported Intentional Learning Environments (CSILE) project: we see much of the current work on computer microworlds as providing excellent opportunities for the student who already knows how to go about the pursuit of expertise, but many students need a more supportive computer environment if they are to function as truly active learners.

Finally, we have hopes for a cooperative project among several cognitive scientists at OISE that will bring modern knowledge about expertise and its development to bear on the problem of learning disabilities. A learning disability may be defined as *ascribed incapacity to develop expertise in a certain domain.* If expertise is taken as relative, however, no child should have to suffer the devastating consequences of such an ascription. But it will be a major challenge to find ways to stop treating children with learning difficulties as patients and to begin treating them as builders of their own expertise.

Notes

1. This also explains why, on occasion, the talented amateur is able to solve a problem that has stumped the experts. But for every such instance, dear to creativity buffs, there must be thousands of instances in which experts solve with ease problems that would baffle the amateur.

2. Notice that the distinction here is not between intrinsic and extrinsic motivation. The activities of the second type of person may be intrinsically motivated, but the intrinsic motivation is not *motivation to develop expertise.*

References

Case, R. (1978). A developmentally based theory and technology of instruction. *Review of Educational Research,* **48,** 439–463.

Gelman, R., & Gallistel, C. R. (1978). *The child's understanding of number.* Cambridge, MA: Harvard University Press.

Karmiloff-Smith, A. (1979). Micro- and macrodevelopmental changes in language acquisition and other representational systems. *Cognitive Science,* **3,** 91–118.

Resnick, L. B., & Neches, R. (1984). Factors affecting individual differences in learning ability. In R. J. Sternberg (Ed.), *Advances in the psychology of human intelligence* (Vol. 2, pp. 275–323). Hillsdale, NJ: Lawrence Erlbaum Associates.

Scardamalia, M., & Bereiter, C. (1984). Development of strategies in text processing. In H. Mandl, N. Stein, & T. Trabasso (Eds.), *Learning and comprehension of text* (pp. 379–406). Hillsdale, NJ: Lawrence Erlbaum Associates.

Scardamalia, M., & Bereiter, C. (1985). Research on written composition. In M. Wittrock (Ed.), *Handbook of research on teaching* (3rd ed.) (pp. 778–803). New York: Macmillan Education Ltd.

Voss, J. F. (1984). On learning and learning from text. In H. Mandl, N. L. Stein, & T. Trabasso (Eds.), *Learning and comprehension of text* (pp. 193–212). Hillsdale, NJ: Lawrence Erlbaum Associates.

Interchange, Vol. 17, No. 2 (Summer 1986), 19–24

DISCUSSION

David Olson, Carl Bereiter, Marlene Scardamalia,
Clive Beck, David Stern, Harvey Graff, Sharon Bailin

DAVID OLSON: I'd like to pick up on one of the points to which Carl and Marlene alluded, and that is the use of models, like novice-expert models, for describing human competence. In the past, of course, we had other kinds of explanations—ability explanations, methodological explanations, and so on. But what I find important about the use of the models is the attempt to account for the mental structures that children use in constructing their own representations of the world. It's not nearly enough to say that a child has a high level of ability or a low level of ability; you want to know what procedures the child is using and how he or she is organizing them in order to carry out his/her goals. And it seems to me that if we pursue the use of models like novice-expert models and move away from the traditional ability theory, we also move away from traditional developmental theory. Carl mentioned that the young expert is rather similar to the older expert. Now, that's not to say that developmental differences are unimportant; of course they continue to be extremely important. But there is an important assumption here about development (which I would attribute initially, perhaps, to Piaget), namely, the notion that the child is building a representation of the world. You have no such thing as an ignorant child or an ignorant learner. You have someone like a novice, an adult novice who has conceptions of the world and a conception of the problem and he or she will use those ideas in interpreting any tasks assigned. Carl and Marlene have obviously struggled with just how models change from novice to expert, and I can't say either how one model gives place to another. Carl mentioned