

ED 308 213

TM 013 536

AUTHOR Bereiter, Carl
 TITLE The Role of an Educational Learning Theory:
 Explaining Difficult Learning.
 PUB DATE Mar 89
 NOTE 12p.; Paper presented at the Annual Meeting of the
 American Educational Research Association (San
 Francisco, CA, March 27-31, 1989).
 PUB TYPE Reports - Evaluative/Feasibility (142) --
 Speeches/Conference Papers (150)

EDRS PRICE MF01/PC01 Plus Postage.
 DESCRIPTORS *Cognitive Psychology; *Difficulty Level;
 *Educational Psychology; Knowledge Level; Learning
 Strategies; *Learning Theories; *Problem Solving;
 Protocol Analysis
 IDENTIFIERS Analogical Reasoning

ABSTRACT

The possibility of developing a learning theory that is designed to insure its relevance to educational problems is discussed. It is suggested that the constitutive problem for an educational psychology of learning is how one learns things that are difficult to learn. Behaviorist learning theories fail almost entirely to explain why anything is harder to learn than anything else. Questions concerning the means by which learners learn to solve difficult problems restate at a higher level the same issue that all cognitive learning theorists must contend with--the means by which one can design a learning system that works without the need for an executive that is already knowledgeable. Problem solving seems to approach a satisfactory means of modeling this learning process. Findings pertinent to this model, however, are quite preliminary and based mostly on case studies. Thinking-aloud protocols show that learners use four kinds of knowledge: (1) knowledge about knowledge; (2) domain-specific knowledge; (3) analogy to more familiar domains; and (4) expectations about the level of promise that a particular path of inquiry might provide. It is concluded that problem solving links cognitive psychology to learning theory. (TJH)

 * Reproductions supplied by EDRS are the best that can be made *
 * from the original document. *

- This document has been reproduced as received from the person or organization originating it.
 Minor changes have been made to improve reproduction quality.

- Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

CARL BEREITER

The Role of an Educational Learning Theory: Explaining Difficult Learning¹

Carl Bereiter
Centre for Applied Cognitive Science
Ontario Institute for Studies in Education

During the period when learning theory was the dominant concern of experimental psychology, there were a number of theories competing fairly directly with one another and trying to answer the same questions. (See, for instance, the early editions of Hilgard's Theories of Learning.) At that time, accordingly, it was reasonable to treat learning theories generically and to ask the question, "What can learning theory contribute to education?" Many psychologists and educators tried their hands at that question, with results that were invariably disappointing.

From the broader perspective of contemporary cognitive science, however, it is possible to envision a range of learning theories addressed to different levels of phenomena and not competing in the sense of making contradictory claims. Dennett (1978) has sketched the three main levels at which such theories might be framed--the physical level (as in neurological theories of learning), the design level (the characteristic level of psychological theorizing), and the intentional level (a level at which more philosophical or humanistic theories may be found but a level which Dennett [1983] argues can also accommodate scientifically testable theories). Sweeping generalizations about the relevance of learning theory to education are therefore no longer supportable. Indeed, the possibility now arises of creating a theory of learning that is designed from the beginning so as to ensure its relevance to educational problems. In this paper I want to offer some ideas about how that might go.

Theorizing enjoys such prestige these days that people seldom ask what we need a theory for. But if we are to be serious about developing theory that addresses a different level of problem from that addressed by existing cognitive theories, then we should be able to

¹Paper presented at the Symposium titled "Toward a Unified Approach to Learning as a Multi-Source Phenomenon," annual meeting of the American Educational Research Association, San Francisco, March 28, 1989.

ED308213

1013536

point to some things that need explaining. I believe the whole-person and phenomenological theories that periodically capture the fancy of educators usually fail on this account. They do not identify any explanatory problems and so they end up just providing a nice way of talking about things people already know.

The constitutive problem for an educational psychology of learning, I would suggest, is how we learn things that are hard to learn. Schools only exist because some things are hard to learn. If everything were easy to learn, libraries and other resource centers would suffice; educationists, to the extent that they existed at all, would be concerned only with the what of learning and not with the how. A theory of learning would have about as much educational significance as a theory of respiration, both dealing with essential processes but ones that can normally be taken for granted.

Behaviorist learning theories failed almost totally to explain why anything was harder to learn than anything else. They could explain why some paired associates were harder to memorize than others, but not, for instance, why minuends with zeros in them are so much harder than ones without. Cognitive theories have made big gains in being able to explain learning difficulties of the latter sort. But why is mathematics as a whole so difficult compared to other subjects? Why does almost everyone eventually fail, even if they did well in earlier courses? And how do we explain those who do succeed? Those are questions. I would suggest, that cannot be answered at the level of production systems or planning nets or neural nets, but they are nonetheless legitimate questions for an educationally relevant theory of learning to tackle.

Theories at different levels of description are usefully interconnected if the lower-level (i.e., more reductive) theory constrains the solutions that are admissible at the higher level and if the higher-level theory constrains the goals that are pursued at the lower level. That kind of useful connection exists, albeit imperfectly, between neuroscience and cognitive psychology. It does not exist, for instance, between cognitive science and holistic or personological theories. A successful educational learning theory should be

such that it is rigorously constrained by more basic cognitive theories and it should in turn help to define what it is that needs explaining at the more basic level. Educational research seems to have failed sadly in this latter respect. Instead of looking to educational research for facts that need explaining, cognitive theorists like John Anderson and Allen Newell have looked to the findings of traditional experimental psychology. Anderson, recognizing the limitations of that database, has set about doing his own educational research.

The question of how we learn things that are hard to learn restates at a higher level the same question that cognitive learning theorists all contend with: How can you design a learning system that works without need for an executive that is already knowledgeable? (In other words, how can a learning system work without having a teacher inside it?) And this question in turn is a version of the philosophical or metatheoretical question, how is it possible for a more complex structure to emerge from less complex ones? These are all versions of what I have discussed elsewhere as the "learning paradox" (Bereiter, 1985). At one level the question is how learning is possible in principle, at another level it is how learning can be modeled under realistic constraints, and at the educational level it is how pragmatically real barriers to learning are overcome. Insofar as these pragmatically real barriers to learning are related to the absence of a knowledgeable teacher inside the learner's head, there is the possibility for a useful theoretical connection between the educational level and the basic cognitive learning theory level. That connection is more obvious if we narrow our constitutive question somewhat, to the following:

How do people learn when they have little relevant prior knowledge?

This question is obviously pertinent to school problems. Students are always being led into areas where their prior knowledge is scanty. At the same time it is closely related to the basic problem of emergent cognitive structures--how a new knowledge state can be built that is more complexly structured than the knowledge states that preceded it.

Learning as Problem Solving

According to John Anderson's theory of skill acquisition (Anderson, 1987), skills start life as problem solutions which later become compiled into procedures. The gap between declarative and procedural knowledge is crossed, according to his theory, by weak problem-solving heuristics that make use of declarative knowledge to solve procedural problems. This idea can be carried to a higher plane if we adopt the premise that learning itself can be treated as a problem, and that the gap between present knowledge and some more advanced state of knowledge can also be bridged by applying weak problem-solving heuristics to existing knowledge.

There is evidence that this strategy actually works. In an experiment carried out by Marlene Bird (Bereiter & Bird, 1985), elementary school students were trained in several comprehension strategies, one of which was to take an experienced difficulty in comprehension, state it as a problem, and then try to solve the problem. No strategies were suggested for solving the comprehension problems--just state a problem and try to solve it. Yet trained students showed a significant increase in their use of this method and showed significant gains in reading comprehension performance compared to controls.¹

If you went to an oracle for advice and were told that what you should do is state your problem and then go off and solve it, you would probably feel you had not gotten much value for your money. But the research on children's comprehension indicates that, although they often show signs of vague awareness that something has gone wrong in comprehension, they seldom identify a problem such that they could attempt to solve it. Human beings have a lot of general problem-solving resources available, but they do not always use them. Evidently just learning to apply problem-solving strategies that they already have to problems of comprehension is an advance for children.

In several studies we have examined individual differences in approaches to learning. The most striking difference is between people who approach learning as a routine task

and those who approach it within a problem-solving framework. An analogy to dish washing may clarify the distinction. An automatic dishwasher has no problem-solving capabilities. It runs through a preset routine and the dishes come out however they come out. Some humans wash dishes by hand in the same way. Others, however, approach dish washing as problem solving. They inspect the item to be washed, select an appropriate procedure, examine the result, and apply back-up procedures as needed to achieve a desired result. When interviewed about how they will go about learning something, adults are much more likely to take a problem-solving approach than children, but among students of the same age, large differences may be found. Margaret Ogilvie (Ogilvie & Bereiter, 1989) reports surprising consistency of individual differences in approach to three quite different kinds of learning--learning of a typical academic sort, learning to juggle, and learning in a computer game involving pattern recognition. The students who take a problem solving approach to learning show a strong tendency to do so in all three tasks and to show superior achievement in all three.

We are trying to get some insight into how problem solving works in difficult learning situations. Findings here are quite preliminary and are based mostly on case studies, but ones that cover quite a range, from elementary school students trying to understand difficult scientific texts to advanced piano students trying to learn a culturally unfamiliar piece of music and medical students going through their first experiences of clinical training. The differences between those who approach learning as a routine and those who approach it as a problem are equally striking at all levels. The following seems to be what problem-solvers do to bridge the gap from present knowledge to more complex knowledge: They set up a problem space with the desired learning as the goal. The goal state can be only very vaguely specified at first, however, because to do otherwise would require that they already possess the knowledge they are trying to achieve. The question to be answered is, how can the goal state be represented at all in such a way as to make it possible to apply means-end analysis and other heuristic search strategies as opposed to just trusting to conventional learning procedures, as the

nonproblem-solvers do?

Using Knowledge to Achieve Knowledge

To begin with, it should be recognized that problems in which goals are initially vague and emerge gradually through the course of problem solving itself are by no means limited to learning. They constitute the large class of what Greeno (1978) has called composition problems. They include writing, visual arts, and theory construction, to name a few. From another standpoint, they comprise a large part of what society recognizes as creative work. It does not help our explanatory effort much to say that learning is not just problem solving but creative problem solving, since creativity itself is so poorly understood. But identifying learning with creative problem solving does make it more obvious why learning, when carried out in a problem-solving mode, is difficult to explain.

People working on learning problems must constrain the goal in some way and add further constraints as they proceed. Thinking-aloud protocols show learners to be making use of four kinds of knowledge in doing this:

1. Knowledge about knowledge. If you are a sophisticated learner venturing into an unfamiliar domain, you can nevertheless make certain assumptions based on previous learning experience:

- That the sheer quantity of knowledge to be acquired is larger than you are aware of
- That the knowledge has a structure and that it is more complex than your current envisionment
- That you are going to have trouble judging the importance of information (and so had better err on the side of overestimating importance)
- That familiar words may have special meanings in the domain
- That you must be on the watch for complicating factors
- That what makes the domain itself interesting and important cannot be fully appreciated until you have acquired more knowledge of it

These assumptions enable you to constrain the learning goal initially, at least to an extent that makes possible the use of means-end heuristics. Novice learners often have no inkling of these things and therefore cannot act strategically.

2. Domain-specific knowledge. Learners are seldom totally ignorant of a domain.

Because of the considerations noted in the preceding point, sophisticated learners are cautious about using scanty prior knowledge for interpreting new information. However, they make good use of this knowledge in order to constrain learning goals. In her study of pianists (masters thesis in progress), Pam'la Ghent provided learners with text material describing Indonesian wayang music and a tape recording of it being played on traditional instruments, before giving them a piano transcription to learn. The problem-solving learner made extensive use of this background information in order to set goals for himself in constructing an interpretation of the piece, whereas the nonproblem-solver largely ignored this domain-specific information and relied instead on already-learned solutions. The result was that the problem-solver produced a performance that an expert critic identified as faithful to the character of wayang music, whereas the nonproblem-solver produced something the critic identified as a French Impressionist style imposed on the music.

3. Analogy to more familiar domains. In one of our interview protocols, someone who has previously learned knitting contemplates learning weaving. She uses her previous experience to infer kinds of things she will need to learn, obstacles she is likely to encounter, and so on.

4. Knowledge of promisingness. As in other creative tasks, learners must embark on paths without knowing for sure where they are heading. Choices are constrained by rather impressionistic judgments of promisingness. Creative experts, we believe, develop a large repertoire of schemas enabling them to recognize, albeit with imperfect accuracy, the deep structural features of promising questions, hypotheses, lines of movement, and so on in their domains. Expert learners (this is an inference; I cannot at this time point

to evidence) must also develop schemas permitting them to evaluate the promisingness of alternative paths in learning. These schemas are acquired through experience in solving learning problems. Nonproblem-solvers, relying on routine learning procedures, probably gradually improve their routines, but they do not acquire the knowledge of promisingness required for creative pursuit of learning goals.

In summary, the picture we get of the problem-solving learner is of someone who makes active use of general knowledge about knowledge, domain specific knowledge, analogies to past learning experiences, and accumulated impressionistic knowledge about promisingness. Using this knowledge not only makes possible the successful pursuit of learning goals under conditions of low prior knowledge but it also contributes to the further development of learning expertise.

Toward a Psychology of Difficult Learning

The picture we have just sketched gives quite a different view of "learning to learn" than has been prevalent in the learning skills literature. Work on learning skills has been concerned with providing students with a better repertoire of routines to apply in learning situations. Nowadays these routines are called "strategies," but generally speaking they lack the conditionality required of strategies. Our analysis suggests this is inevitable, given the emergent character of learning goals. Sophisticated "learning to learn" means learning how to use several varieties of limited prior knowledge in creative problem solving aimed at the attainment of only vaguely specifiable knowledge goals.

Students who rely on routine procedures do learn, however. How do such students manage to learn difficult things? It seems that they learn them the way they learn easy things, only with more time and less success. That is not much of an answer, but I fear it is the best answer that basic cognitive science currently provides.

Basic cognitive science provides us with some powerful theories about skill acquisition, memory, and comprehension. These theories are most successful, however, in accounting for easy learning--that is, learning in situations where what is already available in

memory provides abundant resources for tagging, interpreting, or otherwise processing incoming information. For understanding difficult learning, the best learning theory we have available seems to be problem-solving theory, and that is what I have mainly relied on in the preceding discussion.

I proposed at the outset the idea of an educational learning theory that would stand in relation to basic cognitive psychology as cognitive psychology stands in relation to neuroscience. This requires connecting links that allow for mutual constraint. I am now suggesting that problem solving, as applied to learning itself, provides such a link. It should by no means be the only link. But it is a major link and a promising one. Basic problem-solving theory provides important conceptual tools for examining problematic learning. Educational research, for its part, should be able to probe deeply into the problem-solving processes of learners and into the kinds of knowledge that constitute expertise in learning, identifying a new layer of phenomena that cognitive science must eventually endeavor to explain.

References

- Anderson, J. R. (1987). Skill acquisition: Compilation of weak-method problem solutions. Psychological Review, 94, 192-210. ,
- Bereiter, C. (1985). Toward a solution of the learning paradox. Review of Educational Research, 55, 201-226.
- Bereiter, C., & Bird, M. (1985). Use of thinking aloud in identification and teaching of reading comprehension strategies. Cognition and Instruction, 2, 131-156. /
- Dennett, D. C. (1978). Brainstorms. Montgomery, VT: Bradford Books.
- Dennett, D. C. (1983) Intentional systems in cognitive ethology: The "Panglossian paradigm" defended. Behavioral and Brain Sciences, 6, 343-390.
- Greeno, J. G. (1978). Natures of problem-solving abilities. In W. K. Estes (Ed.), Handbook of learning and cognitive processes (Vol. 5, pp. 239-270). Hillsdale, NJ: Erlbaum.
- Ogilvie, M. F., & Bereiter, C. (1989, March). The role of generalized strategies in the development of learning across domains. Paper presented at the meeting of the American Educational Research Association, San Francisco.

Notes

¹Because several strategies were taught together, it cannot be demonstrated that this particular one contributed to overall performance gains. However, it is sufficient for the present argument that they were able to learn the strategy and use it with apparent success for solving immediate comprehension difficulties.