Session 21D1

Constructing Knowledge Bases: A Methodology for Learning to Synthesize

Karl A. Smith
Anothony M. Starfield
Robert Macneal
University of Minnesota
Minneapolis, Minnesota

> "It is strange that we expect students to learn, yet seldom teach them anything about learning. We expect students to solve problems, yet seldom teach them anything about problem solving. and, similarly, we sometimes require students to remember a considerable body of material, yet seldom teach them the art of memory. It is time we made up for this lack, time that we developed the applied disciplines of learning and problem solving and memory. We need to develop the general principles of how to learn, how to remember, how to solve problems, then to develop applied courses, and then to establish the place of these methods in an academic curriculum."[1] (p. [illegible]7)

This paper begins and ends with a discussion of expert systems and is driven by the concern expressed in the quote above by Don Norman. Our work in the area of requiring students to construct explicit knowledge representations began when we introduced the idea of students building small expert systems. Our primary purpose was to familiarize them with this approach in a course on the application of operations research techniques in engineering. The introduction of this idea had several unanticipated side effects. One, the students were much more enthusiastic than we had expected. Two, they mastered content that we had not expected that they would master. Three, they formulated rules for design and decision-making that showed they had not only reviewed a large amount of information, but that they had reviewed it selectively and purposefully. The outcome of this procedure is described in a recent article by Starfield, Butala, England and Smith.[2]

Our experience with this procedure prompted us to examine how we teach (or, more often, neglect to teach) students how to organize information and synthesize what they know. Both students and teachers tend to concentrate on those aspects of learning where there are established, formal vehicles for conveying skills and for testing whether students have mastered them. Thus outlines and text are the main means that students use for organizing knowledge externally and the majority of student learning involves rote memorization or drill and practice. Memorizing and practicing are well defined activities: the students understand what they are expected to do and teachers understand how to test whether the students have done what was expected of them. Both teachers and students tend to be surprised at how poorly students perform when the tests also expect students to criticize, synthesize, or draw conclusions from what they have learned. They should not be surprised, for while the teacher may have "picked up" these skills, it is not often that a deliberate effort is made to develop them in the student. Where such efforts are made, it is usually via a kind of apprenticeship: "watch how I go about it and emulate me." This is the type of teacher-student interaction that takes place in hospital ward rounds and in engineering design courses. Experts often have difficulty teaching novices, however, because of what Paul Johnson[3] has termed the paradox of expertise. The paradox of expertise states that the knowledge we wish to teach others often turns out to be the knowledge we are least able to talk about.

Learning by apprenticeship can be effective, but in practice it has suffered as faculty-student ratios have deteriorated and as curricula have become overcrowded. What is missing, and sorely needed, is an effective methodology for developing and exercising the skills that enable students to organize and use knowledge constructively. It has taken the authors some time to realize that the experiment in which they required students to build a knowledge base for an expert system provided that kind of methodology. It has taken even longer for them to realize that there are other techniques that serve a similar purpose, and that building a knowledge base belongs in the same tool-kit as spatial learning strategies and the potential uses, in teaching, of computer application strategies such as spreadsheets and databases.

This paper will speculate on the appropriate uses of knowledge representation systems and the need to develop control structures for deciding which knowledge base technique to use, how to use it, and when you're done. We will also discuss how this ties in with the concept of metacognition.

Finally, the role of the learning environment in building knowledge representations, especially active learning and student involvement, will be discussed. Research results on cognitive rehearsal, peer tutoring and the development of expertise through talking with others and preparing to teach will be summarized.

CH2240-0/85/0000-0374 $00.75 © 1985 IEEE

A concept map (a spatial learning strategy), highlighting the specific topics covered in this paper, is shown in Figure 1.

Spatial Learning Strategies

Weinstein and Mayer[4] have devised eight categories that describe the variety of strategies students use to selectively attend to new information, to encode it using processes and in a form that will then facilitate storage, and to recall this new information on either a short-term or a long-term basis. Spatial learning strategies are highlighted under two of these, "Elaboration strategies for complex learning tasks" and the "Organizational strategies for complex learning tasks."

Three relatively content-independent spatial learning strategies--mapping, networking, and schematizing--will be described. Because they require representation of relationships between concepts they facilitate abstraction and deep processing. Unlike more content-dependent techniques (matrixing, flowcharting, constructing pictures or graphs, for example) these systems can be used in a wide variety of text.

Mapping

Concept maps, according to Novak and Gowin,[5] are intended to represent meaningful relationships between concepts in the form of propositions. In its simplest form, a concept map would be just two concepts connected by a linking word to form a proposition.

In a mapping strategy described by Armbruster and Anderson,[6] the student learns a set of relational conventions or symbols. The symbols depict seven fundamental relationships between two ideas, for example, A and B. These relationships are (1) B is an instance of A, (2) B is a property or characteristic of A, (3) A is similar to B, (4) A is greater or less than B, (5) A occurs before B, (6) A causes B, and (7) A is the negation of B. Additionally, two special relationships identify A as an important idea or a definition, the connectives and and or are also used.

Networking

Networking involves three activities: (1) identifying key concepts, (2) laying the concepts

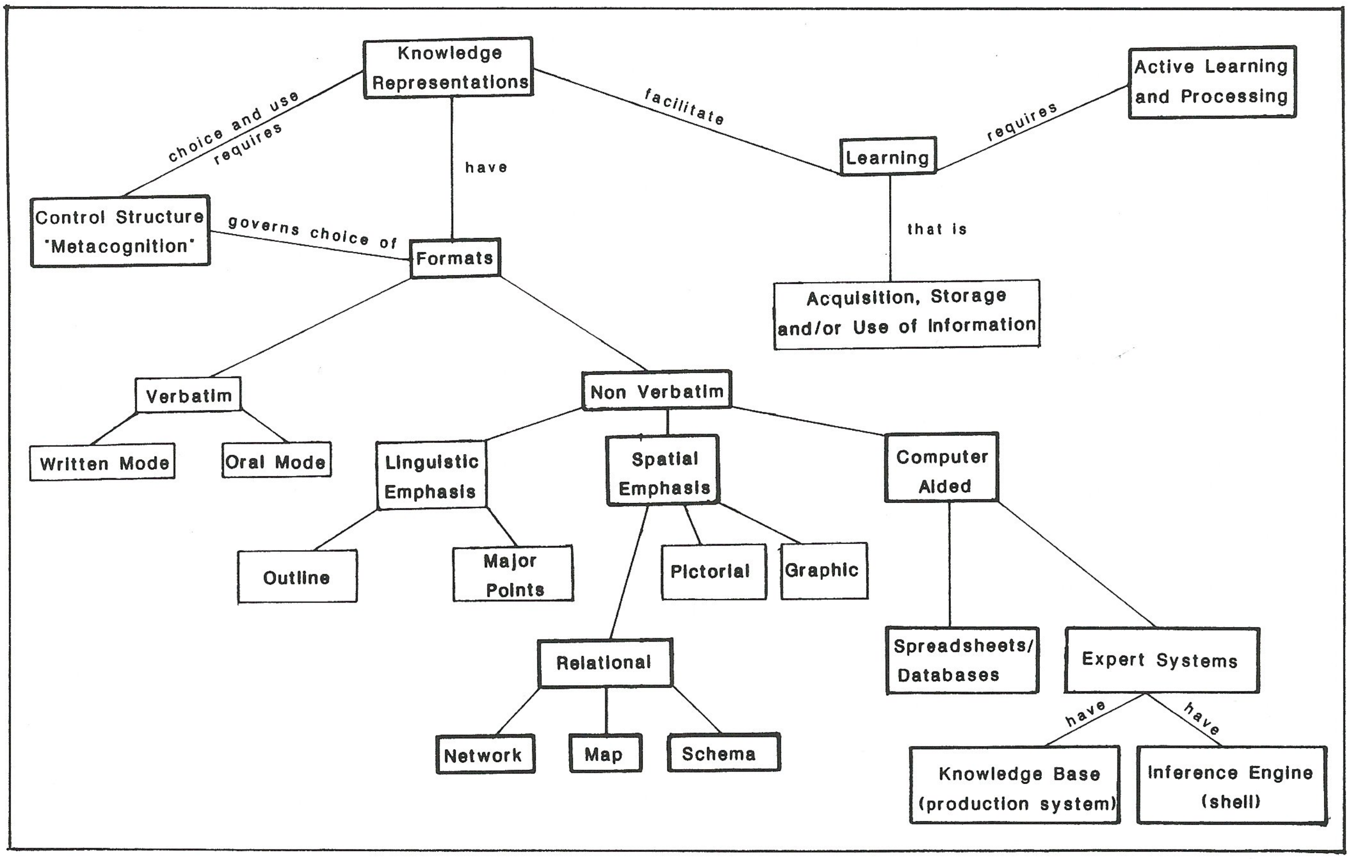

FIGURE 1. CONCEPT MAP FOR 'CONSTRUCTING KNOWLEDGE BASES'

out spatially, and (3) filling in relationships among the concepts. A set of named links can be used to code the relationships between ideas. The networking process emphasizes the identification and representation of (1) hierarchies (type-part), (2) chains (lines of reasoning-temporal orderings-causal sequences), and (3) clusters (characteristics-definitions-analogies). Link types and structure types used in the networking technique described by Holley and Dansereau[7] are shown in Figure 2. Assessments of networking have shown that students using this strategy perform significantly better on text processing tasks than do students using their own methods.

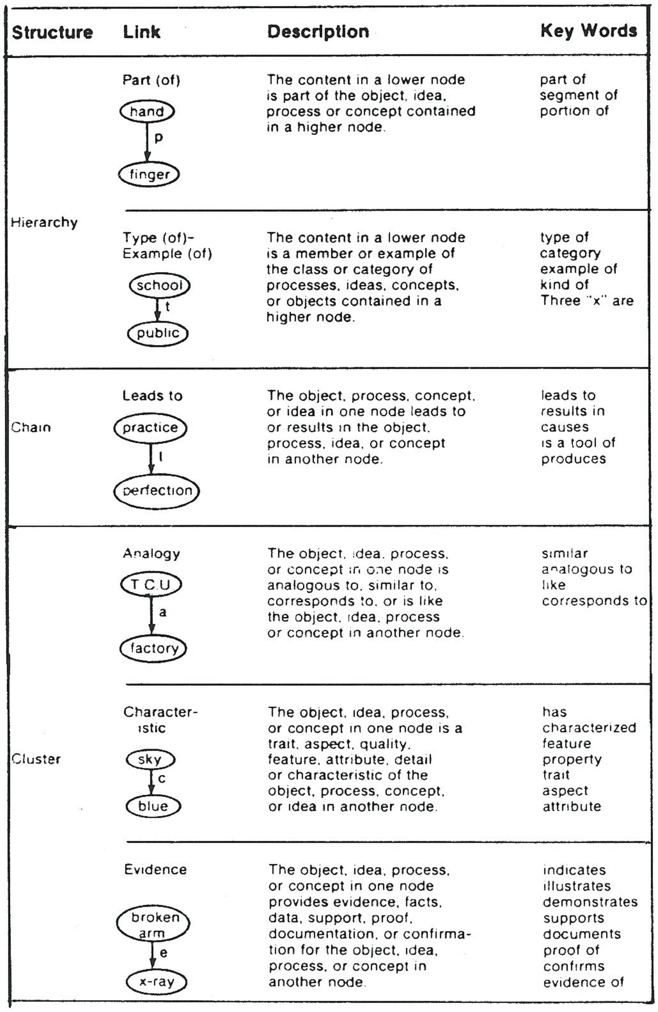

Structure	Link	Description	Key Words
Hierarchy	Part (of) hand p finger	The content in a lower node is part of the object, idea, process or concept contained in a higher node.	part of segment of portion of
	Type (of)- Example (of) school t public	The content in a lower node is a member or example of the class or category of processes, ideas, concepts, or objects contained in a higher node.	type of category example of kind of Three "x" are
Chain	Leads to practice l perfection	The object, process, concept, or idea in one node leads to or results in the object, process, idea, or concept in another node.	leads to results in causes is a tool of produces
Cluster	Analogy T.C.U. a factory	The object, idea, process, or concept in one node is analogous to, similar to, corresponds to, or is like the object, idea, process or concept in another node.	similar analogous to like corresponds to
	Characteristic sky c blue	The object, idea, process, or concept in one node is a trait, aspect, quality, feature, attribute, detail or characteristic of the object, process, concept, or idea in another node.	has characterized feature property trait aspect attribute
	Evidence broken arm e x-ray	The object, idea, process, or concept in one node provides evidence, facts, data, support, proof, documentation, or confirmation for the object, idea, process, or concept in another node.	indicates illustrates demonstrates supports documents proof of confirms evidence of

FIGURE 2. LINK TYPES AND STRUCTURE TYPES EMPLOYED WITH THE NETWORKING TECHNIQUE[7]

Schematizing

Schematizing involves the labeling (and, where appropriate, the clustering) of concepts and the depiction of relationships between concepts by lines that are annotated to reflect seven types of relationships.[8] The seven relationships are (1) similarity, (2) interaction, (3) denial of similarity, (4) denial of a static relation (5) denial of a dynamic relation, (6) negative influence, and (7) positive influence.

Comparison and Outcomes

Outlines and concept maps differ in three main ways:[5] (1) Concept maps show key concepts and propositions in very explicit and concise language whereas outlines usually intermix instructional examples, concepts, and propositions in a matrix that my be hierarchical, but fails to show the superordinate-subordinate relationship between key concepts and propositions; (2) Concept maps are concise and show the key ideational relationship in a simple visual representation; and (3) Concept maps visually emphasize both hierarchical relationships between concepts and propositions and cross links between sets of concepts and propositions.

The principal differences between mapping and networking appear to be that the former strategy (1) emphasizes local organization rather than abstraction of an overall framework or schema, and (2) employs spatial representation of relationships rather than labeled relationships.

Schematizing is similar to networking, but different from mapping in that it uses annotated lines to depict relationships between concepts and emphasizes the extraction of an overall framework or macrostructure. It differs from networking in the types of relations depicted, the methods of annotation used, and the organizational structure of the resulting diagrams.

Although there is not a great deal of empirical support yet for these spatial learning strategies, what is available shows that mapping facilitates delayed recall of narrative prose, networking results in students preforming significantly better on text processing tasks, and that schematizing is an effective processing aid in the context of a general study-skills course.[7]

Computer Aided Representation

Expert Systems

A small expert system shell has been an indispensable part of the way in which we have used knowledge bases as a teaching tool. The shell, described in Starfield, Adams and Bleloch,[9] is written in Pascal and runs on a personal computer with 128K memory. The knowledge bases constructed by the students are stored in text files, according to a simple and flexible format. These files are input data for the shell, which will read, parse, check and interpret the text. The shell then permits users to interact with the knowledge base.

The knowledge base itself is divided into three parts. First there is a set of numbered decisions. The second part consists of a series of questions which aim to solicit the information necessary to select an appropriate decision; associated with each question is a limited number of answers. The third part consists of a set of production rules. Each rule has the format

IF < condition > THEN < decision >

where the condition is a Boolean expression relating to answers to the questions and perhaps also to decisions.

In a typical project, described in Starfield, et al.,[2] students built a knowledge base to select the appropriate explosive for breaking rock. The decisions in that case consisted of a list of different explosives, such as

dec 1 dynamite

dec 2 aluminized slurry, etc.

The questions related to the rock type, the available drill holes and other conditions. A typical question might be

question 3 'Is the drill-hole'

answer 1 'wet?'
2 'dry?'.

An example of a production rule might be

IF (q2 ans3) and (not dec4) and (q5 ans1 or q6 ans2) THEN dec2.

In the interaction between a user and the expert system shell, the shell will ask questions, record and remember the user's responses, and systematically test out the rules until it finds a rule that is valid. It will then print out the appropriate decision. An additional, and from the instructional point of view, important feature is that the user can at any stage ask the shell "why?". The knowledge base may contain reasons (text strings) that are associated with each question and each rule. These should explain, briefly, the reasoning behind the questions and rules respectively. When a user asks "why?" he is first given the reason associated with the current question. If he asks "why?" again he is then given the reason associated with the rule that the shell is currently testing.

We have found the following sequence to be effective in the classroom:

1. First we introduce the concept of an expert system and knowledge base, explain the structure of the production rules, and demonstrate a small system on the computer.
2. We then divide the class into groups of 2 or 3 and ask them to suggest topics that would lend themselves to this kind of approach. The ensuing discussion highlights the differences between topics that are suitable and those that are not.
3. Each group is then requried to construct a knowledge base as a homework assignment over a period of one or two weeks. The students are told to pay particular care to the explanation facility and are required to implement and demonstrate their work, using the shell. This allows faculty and fellow students a chance to critique the assignments.

Our experience is that students adapt very quickly to this formal structure, learn to exploit it, and get a very real sense of achievement when they implement their work. The structure forces them to approach their problem in a pragmatic and purposeful manner. It guides them into thought processes that they may not have previously encountered in a structured environment and teaches them to think explicity in ways that will be essential to them in their professional careers.

Databases and Spreadsheets

The expert system described above is an example of 'build/run' software. It enables students to construct a knowledge base, operate on it, and examine the results. Standard database and spreadsheet programs facilitate the same operations. Students can compare, merge, and test information in a database; in addition, a spreadsheet allows them to introduce and simulate mathematical relationships.

Spreadsheets and data bases can also be thought of as build/run software. A successful spreadsheet appplication must be carefully constructed of knowledge and rules (i.e. values and formulae) in order that one might later observe rippling effect of column by column output of the calculated values. Similarly, a data base must be constructed before various sorting and merging functions can be invoked. These programs provide a means to represent and manipulate knowledge.

Spreadsheet programs are used increasingly outside the traditional business domain. For example, spreadsheets have been used by students in mine network analysis to simulate air flow, wherein known values of air pressure at selected network junctions (entered as real values) and physical laws (entered as formulae) expressing the interdependence of network junctions and network branches.

Instructional Use of Knowledge Representations

The strategies described above are techniques for externalizing concepts and propositions. However, learning the meaning of a piece of knowledge requires dialog, exchange, sharing, and sometimes compromise. Meanings can be shared, discussed, negotiated, and agreed upon. When spatial learning strategies are used in groups of two or three students, it can serve a useful social function and also lead to lively classroom discussion. Cooperative learning groups, an active learning technique described by Smith, Johnson and Johnson[10] and Smith, Wassyng and Starfield,[11] has shown similar positive outcomes. Preparing to teach or tutor another, whether or not any teaching is actually done, results in greater achievement and liking of the subject and other classmembers.[12,13] Cognitive psychology researchers have shown that the two principal contributors to the development of expertise are talking with peers and preparing to teach.[14,15]

Nickerson[16] has argued that understanding is an active process and that representations play an important role in understanding. Carrol and Mack[17] describe an active learning approach to metaphor. They show that metaphor is critical to understanding text processing software and, furthermore, outline two general theories of metaphor--operational and structural. Their operational and structural theories are roughly equivalent to our discussions of expert systems and spatial learning strategies, respectively.

The learner must engage in active analysis of the structure in order to construct a spatial representation. Spatial strategies may be effective not because they provide an image, but rather because by constructing a graphic representation, the learner carries out activities such as analysis, encoding, and organization that are themselves effective regardless of whether or not they result in a spatial representation.

In addition to training and encouraging students to construct knowledge bases to assist in learning, instructors could incorporate a variety of forms of knowledge bases in their lectures or handout materials. Day[18] described modifying lecture materials to incorporate some of these ideas.

Finally, the construction of knowledge bases could be incorported in evaluation procedures. Novak and Gowin,[5] for example, have prepared a scoring procedure for concept mapping, shown in Figure 3. We have required students to build the knowledge representation for expert systems in exams.

Many recognize microcomputers as a revolution in education, however not to supplant traditional educational processes, rather to supplement them by allowing students to experiment with many different situations and to 'instruct' the machine rather than be 'instructed'.[19]

Programs which enable one to modify data and quickly recalculate are excellent tools for sensitivity analysis and encourage a deeper understanding of the behavior of the system being studied. Once a knowledge representation is constructed, one is free tö operate on it with the tools of the particular package using a "WHAT IF?" approach.

Scoring Model

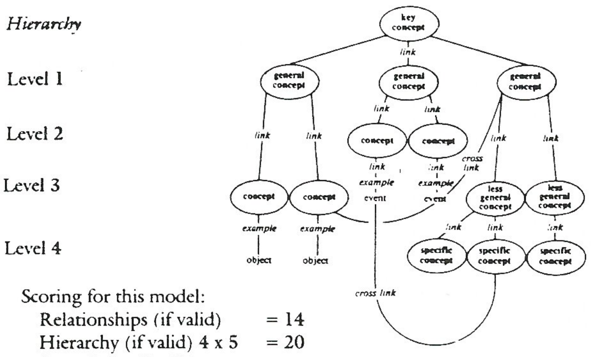

Scoring for this model:
Relationships (if valid) = 14
Hierarchy (if valid) 4 x 5 = 20
Cross links (if valid and significant) 10 x 2 = 20
Examples (if valid) 4 x 1 = 4
58 points total

1. *Propositions.* Is the meaning relationship between two concepts indicated by the connecting line and linking word(s)? Is the relationship valid? For each meaningful, valid proposition shown, score 1 point. (See scoring model below.)
2. *Hierarchy.* Does the map show hierarchy? Is each subordinate concept more specific and less general than the concept drawn above it (in the context of the material being mapped)? Score 5 points for each valid level of the hierarchy.
3. *Cross links.* Does the map show meaningful connections between one segment of the concept hierarchy and another segment? Is the relationship shown significant and valid? Score 10 points for each cross link that is both valid and significant and 2 points for each cross link that is valid but does not illustrate a synthesis between sets of related concepts or propositions. Cross links can indicate creative ability and special care should be given to identifying and rewarding its expression. Unique or creative cross links might receive special recognition, or extra points.
4. *Examples:* Specific events or objects that are valid instances of those designated by the concept label can be scored 1 point each. (These are not circled because they are not concepts.)
5. In addition, a criterion concept map may be constructed, and scored, for the material to be mapped, and the student scores divided by the criterion map score to give a percentage for comparison. (Some students may do better than the criterion and receive more than 100% on this basis).

FIGURE 3. SCORING MODEL AND CRITERIA FOR CONCEPT MAPS[5]

A spreadsheet format includes imbedded formulae expressing relationships between bits or cells of data (knowledge). One can edit and update values in the knowledge base and quickly recalculate. (e.g. "What if I double the energy dissipation?")

Sorting and extracting on a body of knowledge can be quite instructive by exhibiting complex inter-relations which might not be obvious or easily obtainable. (e.g. "What if I sort and extract all elements which possess characteristic A, but not characteristic B?").

Metacognition

The selection of an appropriate knowledge representation technique is one of the most difficult tasks facing students and teachers. To effectively select and use the knowledge representation techniques described one must employ a control structure. This control structure is commonly called metacognition. Flavell's[20] generally accepted definition of metacognition is as follows:

> "Metacognition refers to one's knowledge concerning one's own cognitive processes and products or anything related to them, e.g., the learning-relevant properties of information or data. . . Metacognition refers, among other things, to the active monitoring and consequent regulation and orchestration of these processes in relation to the cognitive objects on which they bear, usually in the service of some concrete goal or objective." (p. 232).

Garofalo and Lester[21] have outlined the major aspects of metacognition as it pertains to instruction in mathematics. The two main attributes of their model are knowledge of cognition (person, task and strategy) and regulation of cognition.

Since there are numerous strategies available to each learner, the approach that is needed is to help the student choose appropriate strategies for each learning task. Bruner[22] claims that "We would do well to equip learners with a menu of their possibilities and, in the course of their education, to arm them with procedures and sensibilities that would make it possible for them to use the menu wisely."

McKeachie[23] points out in a critique of spatial learning strategies that he has difficulty justifying the amount of time needed to teach students how to implement these strategies. In contrast, our experience is that it takes very little time for students to learn how to operate within the formalism of a rule-based knowledge system or to exploit the facilities provided by a database or spreadsheet package. One might speculate that if spatial learning strategies were automated, they would become a far more effective teaching tool. They would then share the advantages of the computer aided representations discussed here, namely, that they provide a framework for the student, an intellectual 'jungle jim' which encourages them (1) to review (information or data) and select from it, (2) to think about what is important and what is secondary, (3) to hypothesize inter-connections, consequences and relationships, (4) to test those hypotheses and explore sensitivity, and (5) to communicate what they learn.

Perry and Downs[24] stress that new technology and the need to learn quickly and effectively, requires that learners become deliberately instrumental in their own learning. They present evidence to support the idea that learning strategies can be introduced to, and assimilated by the learner and be used in subsequent learning tasks.

Conclusion

The development of higher cognitive skills that enable students to be independent learners and independent, creative problem-solving users of their knowledge is a very important goal for educators. It is not a new concern as evidenced by the following 1915 quote from Judd:[25]

> "Most science textbooks can be criticized by drawing attention to the fact. . . .that these books are chiefly concerned with the statements of results. Usually the most general results are put near the beginning of the textbook. A textbook in physics begins by telling about molecules and the constitution of matter or by giving some of the most compactly formulated statements about the principles of mechanics. . . .The degree of enthusiasm of the ordinary student for these introductions which he gets in the textbooks is very slight indeed. . . .The student, confronted by these verbal additions to his experience, gets into the habit of thinking of science as verbal additions to experience, and he faithfully learns the words and keeps them in store against the time when the teacher demands them (p. 334)."

Providing students with an active learning environment where they can get involved with the material to be learned in a mutually supportive situation with other people and providing them with tools such as the ones described here will contribute to meaningful learning.

Even if learning, thinking, and problem-solving strategies, whether general or specific, are shown to exist, it might not be possible to teach them directly. Perhaps they must spontaneously emerge as consequence of substantial experience.[25] Brown[26] has suggested, in fact, that metacognition--conscious awareness of and control of cognitive processes--emerges only as knowledge and skills in a particular domain become quite well developed. At the very least, it should be possible to select and design experience to result in a more rapid and complete emergence of such skills. A key to the success of developing students' skill at using these strategies is for faculty to incorporate them in their handouts, exercises, lectures, assignments and exams.

Experts, such as faculty, have difficulty explaining their expertise, due to what Paul Johnson has called "the paradox of expertise." Experts have very well developed knowledge structures as evidenced by their performance. The question is whether or not the difficulty in explaining is embedded in human psychology or whether it is because they have never been explicitly taught how to formally structure knowledge.

References

1. Norman, D.A. (1980). Cognitive Engineering and Education. In D.T. Tuma & F. Reif (Eds.) Problem Solving and Education: Issues in Teaching and Research. (pp. 97-107), Hillsdale, NJ: Erlbaum.

2. Starfield, A.M., Butala, K.L., England, M.M. and Smith, K.A. (1983). Mastering Engineering Concepts by Building an Expert System. Engineering Education, 74, 104-107.

3. Johnson, P.E. (1983). What Kind of Expert Should a System Be? The Journal of Medicine and Philosophy, 8, 77-97.

4. Weinstein, C.E. and Mayer, R.E. (In press). The Teaching of Learning Strategies. In M.C. Wittrock (Ed), Handbook of Research on Teaching (3rd ed.), New York: Macmillan.

5. Novak, J.D. and Gowin, D.B., (1984). Learning How to Learn, Cambridge: Cambridge U Press.

6. Armbruster, B.B. and Anderson, T.H. (1984). Mapping: Representing Informative Text Diagrammatically. In C.D. Holley & D.F. Dansereau (Eds.), Spatial Learning Strategies (pp. 189-208), New York: Academic Press.

7. Holley, C.D. and Dansereau, D.F., (1984). Networking: The Technique and the Empirical Evidence. In C.D. Holley & D.F. Dansereau (Eds.), Spatial Learning Strategies (pp 81-106), New York: Academic Press.

8. Mirande, M.J.A. (1984). Schematizing: Technique and Application. In C.D. Holley & D.F. Dansereau (Eds.), Spatial Learning-Strategies (pp 149-160), New York: Academic Press.

9. Starfield, A.M., Adams, S. and Blelock, A. (1985). A Small Expert System Shell and its Applications. Proceedings Fourth Annual Phoenix Conference, IEEE, 262-267.

10. Smith, K.A., Johnson, D.W. and Johnson, R.T. (1981). Structuring Learning Goals to Meet the Goals of Engineering Education. Engineering Education, 72(3), 221-226.

11. Smith, K.A., Wassyng, A. and Starfield, A.M. (1983). Development of a Systematic Problem Solving Course: An Alternative to the Use of Case Studies. In L.P. Grayson and J.M. Biedenbach (Eds.), Proceedings Thirteenth Annual Frontiers in Education Conference, Worcester, MA, Washington: IEEE/ASEE, 42-46.

12. Goodlad, S. (1979). Learning by Teaching: An Introduction to Tutoring. London: Community Service Volunteers.

13. Annis, L.F. (1983). The Process and Effects of Peer Tutoring. Human Learning, 2, 39-47.

14. Champagne, A.B., Klopfer, L.E., and Gunstone, R.F. (1982). Cognitive Research and the Design of Science Instruction. Educational-Psychologist, 17(1), 31-35.

15. Stewart, J. (1985). Cognitive Science and Science Education. European Journal of Science Education, 7, 1-17.

16. Nickerson, R.S. (1985). Understanding Understanding. American Journal of Education, 93, 201-239.

17. Carrol, J. M. & Mack, R. L. (1985) Metaphor, Computing Systems, and Active Learning. International Journal of Man-Machine Studies, 22, 39-57.

18. Day, R.S. (1980). Teaching from Notes: Some Cognitive Consequences. In W.F. McKeachie (Ed.), Learning, Cognition, and College Teaching (pp. 95-112), New Directions for-Teaching and Learning 2.

19. Wassyng, A. & Sharp, S. (1985) Effective Teaching in Engineering using Computer Aided Learning. Proceedings of the North Midwest Section Meeting, ASEE, 49-60.

20. Flavell, J.H. (1976). Metacognitive Aspects of Problem Solving. In L. Resnick (Ed.), The Nature of Intelligence, (pp. 231-236), Hillsdale, NJ: Erlbaum.

21. Garofalo, J. and Lester, F.K. (1985). Metacognition, Cognitive Monitoring, and Mathematical Performance. Journal for Research on Mathematics Education, 16, 163-176.

22. Bruner, J. (1985). Models of the Learner. Education Researcher, 14(6), 5-8.

23. McKeachie, W.J. (1984). Spatial Strategies: Critique and Educational Implications. In C.D. Holley and D.F. Dansereau (Eds.), Spatial Learning Strategies (pp.301-312), New York: Academic Press.

24. Perry P. and Downs, S. (1985). Skills, Strategies and Ways of Learning: Can we Help People Learn How to Learn? Programmed Learning and Educational Technology, 22, 177-182.

25. Judd, C.H. (1915). Psychology of High School Subjects. New York: Wiley.

26. Chipman, S.F. and Segal, J.W. (1985). Higher Cognitive Goals for Education. In S.F. Chipman, J.W. Segal and R. Glaser (Eds.), Thinking and Learning Skills (V.1): Relating Research to Instruction (pp. 1-20). Hillsdale, NJ: Erlbaum.

27. Brown, A.L. (1980). Metacognitive Development and Reading. In R.J. Spiro, B.C. Bruce, W.F. Brewer (Eds.), Theoretical Issues in Reading Comprehension. Hillsdale, NJ: Erlbaum.

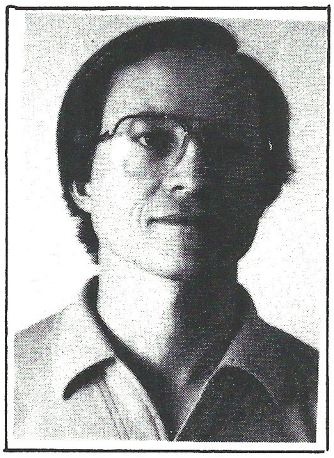

Karl A. Smith

Karl A. Smith is an Assistant Professor of Metallurgical Engineering at the University of Minnesota. He teaches undergraduate courses in mineral processing, extractive metallurgy and systems analysis, and conducts a seminar on environmental issues. His research is in the areas of mineral processing, extractive metallurgy, expert systems, and engineering education. He has extensive experience using cooperative learning procedures and structured controversy.

Karl is a memeber of AIME, ASEE, AERA, AAAS, Phi Kappa Phi, and Alpha Sigma Mu. His education consists of a B.S. and M.S. degree in Metallurgical Engineering from Michigan Technological University and a Ph.D. degree in Educational Psychology from the University of Minnesota.

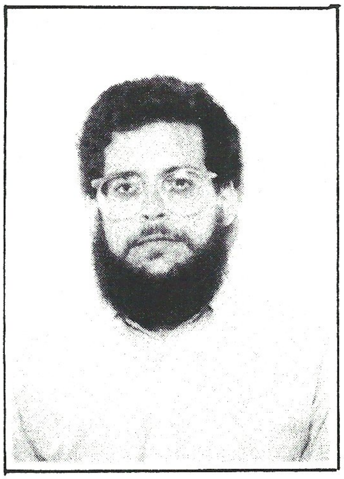

Bob Macneal

Bob Macneal is an undergraduate in Geo-Engineering at the University of Minnesota with a keen interest in microcomputer applications in traditional numerically based programs, as well as descriptively oriented decision support software. Bob has been involoved in several software development projects for the department of Civil and Mineral Engineering (Minnesota) as an undergraduate research assistant. In addition, he has been an undergraduate teaching assistant for a Computer Aided Instruction (CAI) pilot program in Introductory Soil Mechanics. Recently he has been involved in authoring an automated knowledge acquisition program for the Starfield-Adams expert system shell. Bob has a BS in Film and Television Production from Montana State University and has received the Public Works Association Scholarship, NORCUS - DOE Summer Fellowship Award, and Mineral Industries Education Fund Scholarship.

Anthony M. Starfield

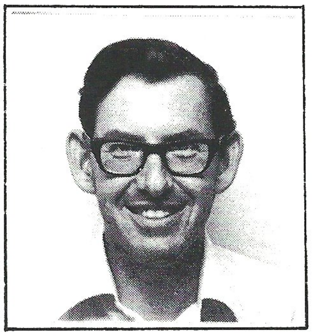

Professor Starfield received his Bachelors degree in Applied Mathematics and his Ph.D. degree in Mining Engineering from the University of Witwatersrand. He was Professor and IBM head of the Department of Applied Mathematics at the University of Witwatersrand and is currently a Professor in the Department of Civil and Mineral Engineering at the University of Minnesota. His research interests are in numerical methods and computational approaches for a variety of problems, the central theme being an orientation towards decision making or design. He has supervised 13 graduate theses and has authored computer methods in rock mechanics and ecological modeling. He is the co-author of two books on applied mathematics, and a third, on ecological modeling, is in press.