DOCUMENT RESUME

ED 237 326 SE 043 410

AUTHOR	Greeno, James G.
TITLE	Conceptual Entities.
INSTITUTION	Pittsburgh Univ., Pa. Learning Research and Development Center.
SPONS AGENCY	National Inst. of Education (ED), Washington, DC.
REPORT NO	LRDC-1983/14
PUB DATE	83
NOTE	29p.
PUB TYPE	Reports - Research/Technical (143) -- Viewpoints (120)
EDRS PRICE	MF01/PC02 Plus Postage.
DESCRIPTORS	*Cognitive Processes; Educational Research; Elementary Secondary Education; *Learning Theories; *Mathematics Education; *Mathematics Instruction; Physics; *Problem Solving; Proof (Mathematics); Statistics
IDENTIFIERS	*Mathematics Education Research

ABSTRACT

Discussed is one "quite general" attribute that can differentiate problem representations: the kinds of entities that are included -- the cognitive objects that the system can reason about in a relatively direct way, and that are included continuously in the representation. The ontology of a domain is significant for four reasons. First, ontology is a significant factor in forming analogies between domains, described in terms of two examples involving problem-solving procedures between domains: geometric proofs and subtraction procedures. Second, these entities provide arguments on which general reasoning procedures can operate directly; this is explored through physics problems; distance, time, and velocity; and sound transmission. That conceptual entities can enable more efficient computation is presented in terms of monster problems, isomorphic to the Tower of Hanoi problem. Finally, the ontology of a problem domain has important effects on goal definition and planning, illustrated by studies of binomial probability. (MNS)

5E

ED237326

LEARNING RESEARCH AND DEVELOPMENT CENTER

1983/14

CONCEPTUAL ENTITIES

JAMES G. GREENO

University of Pittsburgh

CONCEPTUAL ENTITIES

James G. Greeno

Learning Research and Development Center

University of Pittsburgh

1983

In D. Gentner & A. Stevens (Eds.), Mental models. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1983. Reprinted by permission.

The research reported herein was supported by the Learning Research and Development Center, supported in part as a research and development center by funds from the National Institute of Education (NIE), Department of Education. The opinions expressed do not necessarily reflect the position or policy of NIE and no official endorsement should be inferred.

10 Conceptual Entities

James G. Greeno
University of Pittsburgh

Representations of a problem can differ in several ways. In this essay, I discuss one quite general attribute that can differentiate problem representations: the kinds of entities that are included. By *the entities in a representation,* I refer to the cognitive objects that the system can reason about in a relatively direct way, and that are included continuously in the representation.

A system reasons directly about an object if it has procedures that take the object as an argument. In this regard, entities can be distinguished from attributes and relations, which have to be retrieved or computed using the entities as cues or arguments.

Continous inclusion is often achieved by creating an entity in the initial interpretation of a situation, and revising it whenever the situation is changed. Inclusion in the initial representation is not required; entities can be created in the course of working on a problem as well. The important feature is that an entity is maintained once it is created; this distinguishes entities from intermediate results that are removed from the representation after they have been used.

It seems appropriate to use the term *ontology* to refer to the entities that are available for representing problem situations. Therefore, by the *ontology of a domain* (for a representational system), I refer to a characterization of terms used in describing situations and problems in the domain. The ontology of the domain says which terms can refer to entities, and which only refer to attributes or relations.

I hypothesize that the ontology of a domain is significant for four reasons. The first hypothesis is that ontology is a significant factor in forming analogies between domains. An analogy is a mapping between objects and relations in two domains. If the domains are represented with entities that have relations that are similar the analogy might be found easily, but if either domain's representation lacks those entities, the analogy might be difficult or impossible to find.

A corollary of the first hypothesis is that an analogy can be used in facilitating the acquisition of representational knowledge in a domain. If an instructional goal is the learning of a representation that includes a specified set of conceptual entities, then that may be facilitated by providing an analogy with a domain for which a natural representation includes entities that correspond to those that are to be acquired in the target domain.

The second hypothesis is that ontology determines the kinds of information that are available for reasoning using general methods. It seems reasonable to suppose that human problem solvers have some very general reasoning procedures that can be used when appropriate information is available. Examples include reasoning about combinations of quantities that are related as parts and wholes, or comparisons of quantities in ordered sets. The ontology of a domain determines the kinds of information that will be available in the representation, and therefore will be available for use in general reasoning methods.

Third, the ontology of a domain has an obvious consequence for computational efficiency. Ontology determines which kinds of information will be available directly whenever they are needed, and which kinds of information will have to be computed. It clearly is an advantage to keep those items of information available that will be needed frequently, and this is achieved by creating entities corresponding to those items of information.

The fourth hypothesis is an extension of the third. It seems likely that ontology should be a significant factor in planning. A reasonable conjecture is that procedures of planning operate primarily on the entities that are formed in the initial representation of a problem. Thus, representational knowledge that includes an appropriate set of conceptual entities should enable a problem solver to evaluate problem information and choose among alternative goals and plans efficiently.

The fourth hypothesis applies especially to problem solving in domains where formulas are used to solve problems presented in text, such as physics problems and word problems in mathematics. Problem solving should be facilitated if representational knowledge that is applied to problem texts forms conceptual entities that correspond directly to variables in formulas. One way for this to occur would be for knowledge of formulas to include schemata that can be instantiated on the information in problem texts. Schemata that enable an integrated representation of problem information will facilitate judgments about the sufficiency and consistency of problem information and choice of problem goals.

In the remainder of this chapter, I discuss examples in which empirical findings are interpretable in terms of these four hypotheses about conceptual entities in problem solving.

I. ANALOGIES BETWEEN DOMAINS

I will discuss two ... mapping of problem-solving procedures between domains. The first example ... from high-school geometry, and provides an analysis of knowledge ... context of one domain of problems that can provide ... to another domain of problems. The second example is ... arithmetic, and provides an analysis of instruction that uses ... analogy between procedures in two domains in order to facilitate acquisition ... understanding of multidigit subtraction.

Geometry Proofs

The analysis ... first was concerned with an issue in the psychology of learning ... Wertheimer (1945/1959). The issue is whether when students ... their knowledge enables them to understand the problems or ... carry out rote, mechanical solutions.

An example ... Wertheimer discussed is in Fig. 10.1. Wertheimer contrasted two ways in which ... theorem of vertical angles can be proved. One method, which Wertheimer characterized as mechanical, uses an algebraic representation. Quantites in the problem, the sizes of angles, are translated into algebraic terms and a proof is derived using equations. The algebraic steps are indicated in solution (a) in Fig. 10.1.

The second method, which Wertheimer characterized as a solution with understanding, uses a geometric representation to a greater extent. The representation includes part-whole relations between angles, as indicated in solution (b) of Figure 1. The two whole structures, x with w and x with z, are equal because they are both angles formed by straight lines. Furthermore, they share a common part, x. The proof rests on the principle that if the same thing is removed from two equal quantities, then the remainders are equal.

The solution that Wertheimer preferred uses a representation that includes geometric entities that are not included in the more algebraic solution. In the more geometric solution, the straight-line angles are entities; that is, they are cognitive objects whose relationships are used in the solution. The only geometric entities that are needed in the algebraic solution are the labeled angles w, x, and z.

In geometry courses in high school, problems about angles, like the vertical angles theorem, are preceded by instruction in solving problems about line segments. A model was developed that simulated learning from three example problems about line segments. The model has been discussed previously, in another context (Anderson, Greeno, Kline, & Neves, 1981). The example problems are shown in Fig. 10.2. Note that the third problem has the same structure as the theorem of vertical angles, but is about lengths of line segments rather than

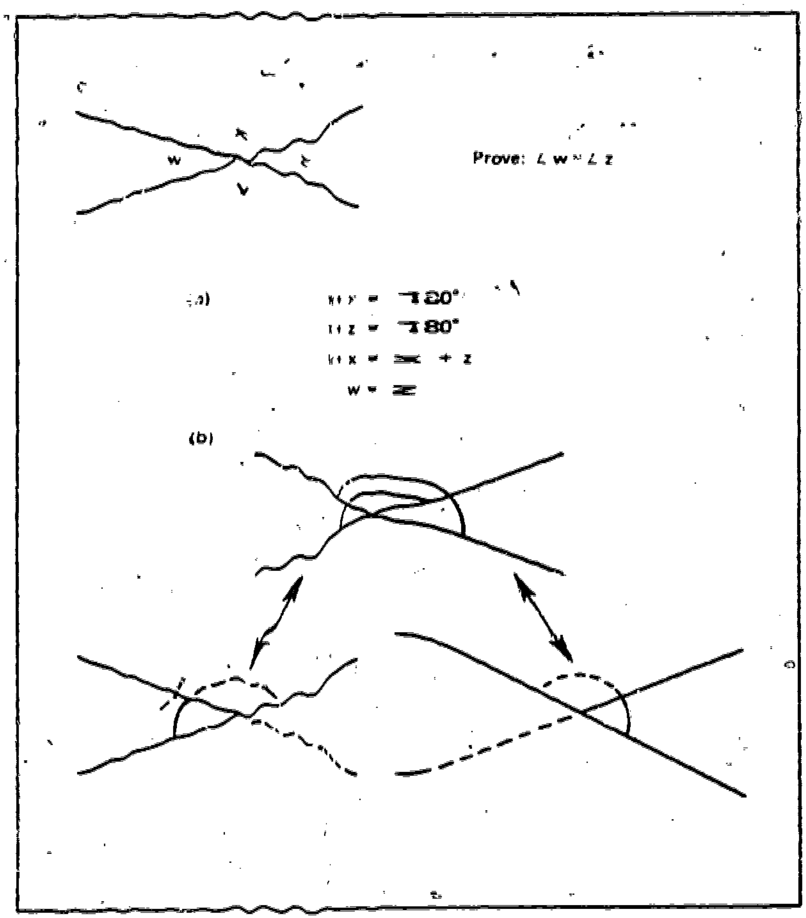

Fig. 10.1. The vertical angles problem with two solutions, from Wertheimer (1945/1959).

sizes of angles. The theoretical goal was to develop a hypothesis about knowledge structures that could be acquired in learning to solve the problems in Fig. 10.2 that would provide a basis for transfer to the vertical-angles problem.

Two simulations of learning were implemented. In one version, called stimulus-response learning, new problem-solving procedures were acquired by associating actions from the example problems with a representation of the problem situations in which the actions occurred. The knowledge acquired in this simulation was very limited in its applicability; however, if mechanisms of stimulus generalization and discrimination like those discussed by Anderson et al. (1981)

were provided, they probably would give a fairly accurate simulation of the knowledge that many students acquire from examples like these.

The second version, called meaningful learning, simulated learning with structural understanding. In meaningful learning, new problem-solving procedures were associated with schematic knowledge about part-whole relationships. The model's initial knowledge included a schema for representing situations involving whole quantities made up of parts, and making inferences about one of the quantities when the others were given.

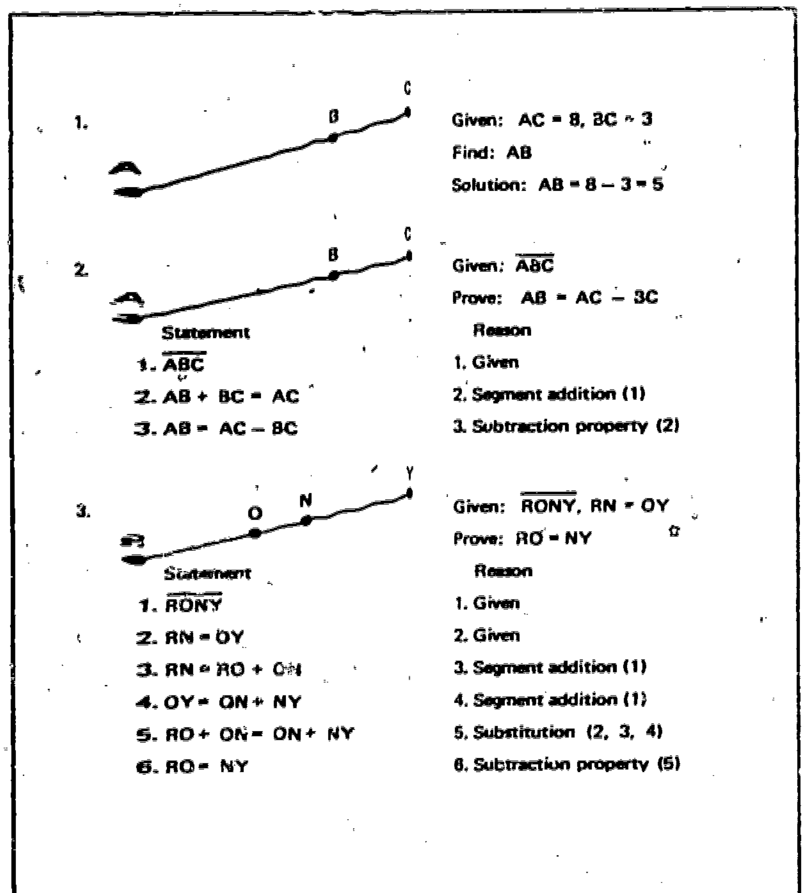

Fig. 10.2. Example problems used for simulations of learning.

From Problem 1, the meaningful-learning model acquired a production for applying its whole-parts schema in situations involving line segments. This knowledge enabled the model to represent problems about lengths of line segments in terms of their part-whole relations, and to use its general procedures for making quantitative inferences about parts and wholes in solving these problems.

From Problem 2, new problem-solving procedures were acquired, with actions of writing lines of proof corresponding to the steps in the example solutions. In meaningful learning, these were acquired as procedural attachments (in the sense of KRL, Bobrow & Winograd, 1977) associated with the whole-parts schema. The arguments of the acquired procedures are objects that occupy slots in the schema; for example, the procedure for writing a line with "Segment Addition" as the reason finds the segments that are the parts and the segment that is the whole, and writes "<part1> + <part2> = <whole>."

From Problem 3, the meaningful-learning model acquired a new schema, which it composed using its previously existing whole-parts schema. The new schema has two whole-parts structures as subschemata, with the provision that one of their parts is shared. The system had access to procedures attached to the subschemata; for example, the procedure for writing lines of proof stating that the whole is equal to the sum of the parts did not have to be acquired from Problem 3, since it was attached to the whole-parts schema previously.

The knowledge acquired in meaningful learning could provide a basis for transfer to problems about other kinds of objects, such as the vertical-angles problem in Fig. 10.1. There is evidence that some students acquire knowledge of that generality in studying problems like those in Fig. 10.2. In one study, six students were interviewed approximately once per week during the year that they were studying geometry. One interview included the problem shown in Fig. 10.3 and the vertical-angles problem. This interview was conducted just after the students had finished a unit on proof about line segments, which included Problem 2 and Problem 3 from Fig. 10.2 as example problems. The students had begun to study angles, and had learned some concepts such as supplementary angles and adjacent angles, but they had not yet done proofs about angles.

Three of the six students gave quite clear evidence in their protocols of conceptualizing the problem in Fig. 10.3 as a structure involving parts and wholes. Their protocols included comments such as "these are the same," and "I have to subtract," applied to appropriate quantities and combinations. Two of the students gave proofs that were conceptually sound, but that were technically incorrect. The errors made the proofs correspond more closely to the overlapping whole-parts structure than does a correct proof. The third of these students failed to prove Fig. 10.3, apparently because of weak knowledge of procedures.

The other three students did not show evidence for representing Fig. 10.3 as overlapping whole-parts structures. One student solved the problem easily using a theorem about supplementary angles. Another student worked out a proof that was technically correct, and appeared to involve applying a procedure for sub-

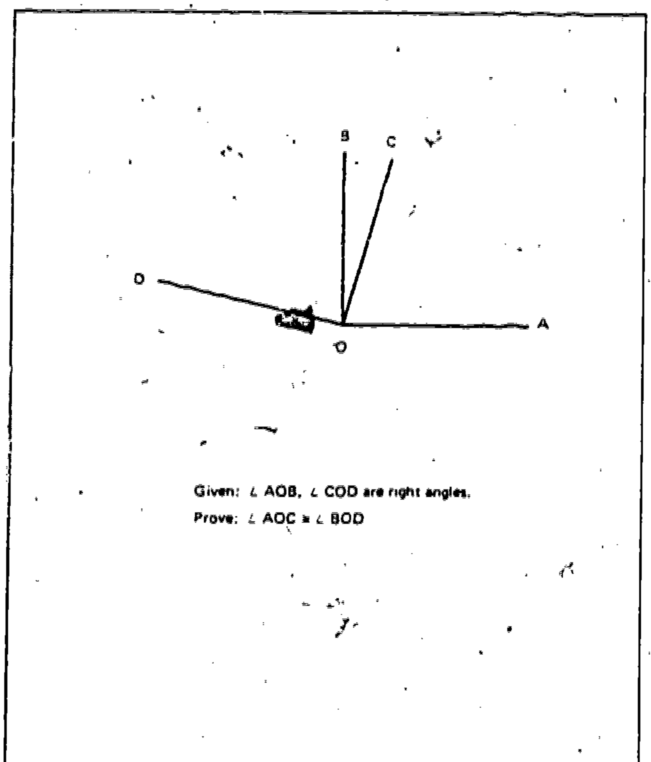

Fig. 10.3. Transfer problem given to students.

stitution in an equation. The sixth student was unable to make progress on Fig. 10.3, and in further questioning it seemed that this student had not learned how to solve the segment problems.

A similar variety of responses was obtained when the vertical-angles problem was presented. One of the students who solved Fig. 10.3 with the schema said, "This is the same problem again. You know something? I'm getting sort of tired of solving this problem." The student who appeared to apply the substitution procedure for Figure 3 failed to prove the vertical-angles theorem; this student got caught in a perceptual difficulty in the vertical-angles problem, where w and x are considered as a pair, and y and z are considered as the other pair.

The knowledge acquired in meaningful learning illustrates the role that conceptual entities can play in a problem representation. With the representational knowledge that enables line segments to be represented as parts and wholes, the model's general procedures for making inferences about parts and wholes can operate directly on the quantities presented in problem situations. This analysis also shows a way in which procedures that are acquired in one kind of problem situation can be applied in another kind of problem, if the procedures take arguments that are specified as the slots of a schema that can be applied to both problem domains.

Subtraction Procedure

The analysis of learning in geometry discussed earlier includes models that learn with and without understanding, but there is no analysis there of conditions that facilitate learning with understanding. In the domain of subtraction, we have analyzed a method of instruction that seems to make understanding likely. The method was developed by Resnick (in press); she calls it instruction by mapping. The instruction has been successful in correcting systematic errors in children's performance on subtraction problems. Children's explanations indicate that they also gain understanding of principles of place value in numeration and the subtraction procedure. We have developed a hypothetical analysis of learning that this instruction produces, in which representational knowledge of subtraction is acquired, including new conceptual entities.

The instructional method uses blocks to facilitate students' understanding of principles involved in addition and subtraction of multidigit numbers. Place-values of ones, tens, hundreds, and thousands are represented by blocks of different sizes and shapes. Representations of numbers are formed with the blocks, and procedures for addition and subtraction are defined. A correspondence can be formed between the procedures that use blocks and the procedures that use ordinary written numerals. For example, carrying and borrowing with numerals correspond to trading with blocks, where one block of a certain size is traded for ten blocks of the next smaller size. Use of blocks in the teaching of arithmetic is quite common. The distinctive feature of Resnick's instruction is that the correspondence between procedures in the two domains is spelled out in detail, and steps are taken to ensure that the student realizes which components of each procedure correspond to components of the other.

In Resnick's empirical research, the recipients of instruction have been children who needed remedial work on subtraction. The work has been done with fourth grade students who performed subtraction with bugs, according to Brown and Burton's (1978) analysis. Figure 10.4 shows two examples. The first problem is solved with a procedure called the smaller-from-larger bug; the answer in each column is found by subtracting the smaller from the larger digit in that column, regardless of which is on the top. The second and third problems illustrate another bug, called don't-decrement-zero. When borrowing is required and a zero is encountered, a *one* is added where it is needed, but nothing is decremented to compensate for that.

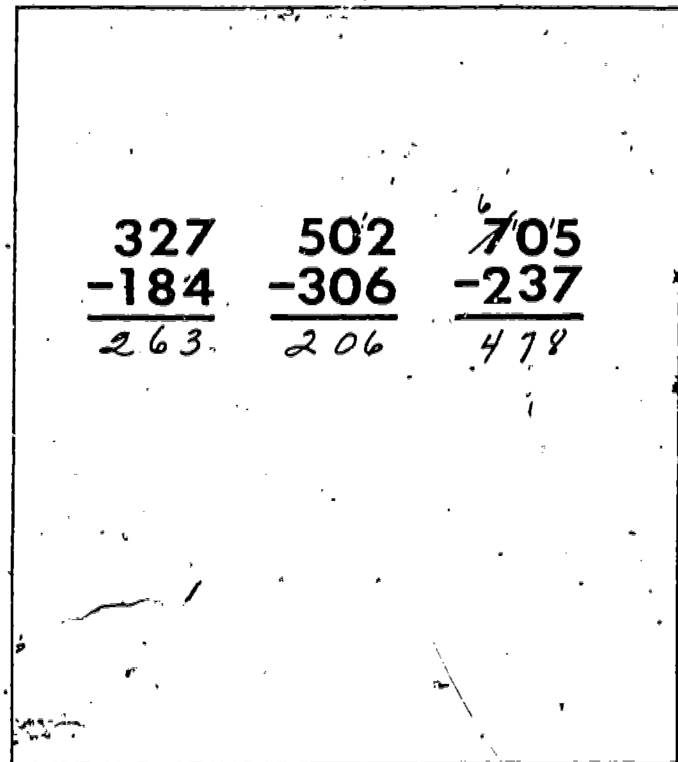

Fig. 10.4. Subtraction problems solved with buggy algorithms.

In Resnick's instruction, children are taught a procedure for subtracting with blocks. In this procedure, the top number in the subtraction problem is represented with blocks, and the number of blocks indicated by the bottom number of the problem is taken away, column by column. When there are too few blocks in one of the top-number piles, a block from the next pile to the left is traded for ten blocks of the size needed. If there are no blocks in the next pile to the left (corresponding to a zero in the top number) a block is taken from the next nonempty pile, traded for ten of the size to its right, one of those is traded for ten of the next smaller size, and so on, until the pile is reached where the extra blocks were needed.

After the child has learned to subtract with blocks, the correspondence between blocks and numerals is taught. For each action performed with blocks, a corresponding action is performed with the written numerals. An example is shown in Fig. 10.5. When a block is removed in borrowing, the corresponding

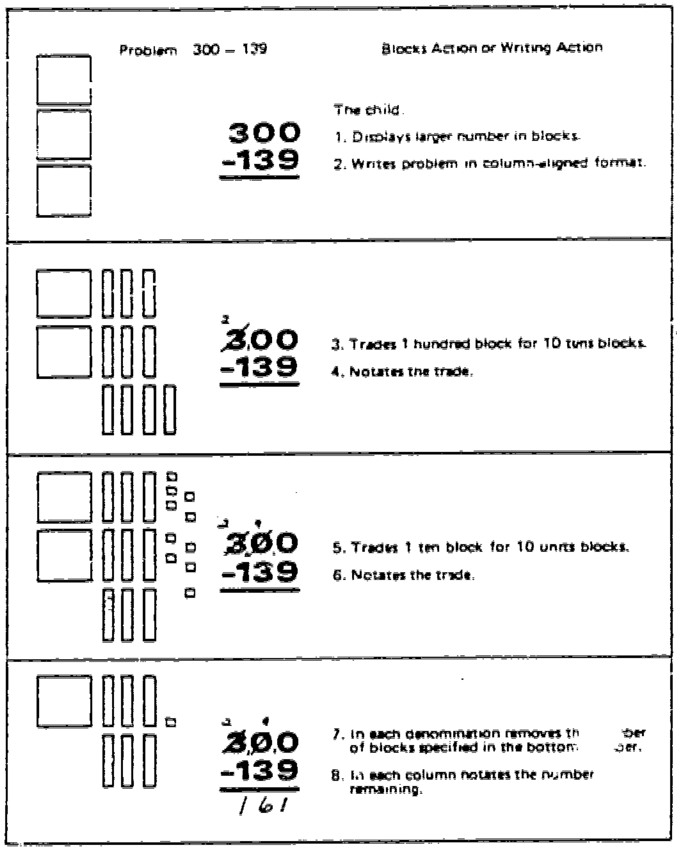

Fig. 10.5. An outline of mapping instruction for borrowing.

numeral is decremented. When ten blocks of the next size are put into the display, the digit for that column is increased by ten. When the number of blocks in a bottom digit are taken away from a pile, the remaining number of blocks is written as the answer for that column.

This instructional sequence can be quite effective. Resnick has recorded several successful cases in which children with bugs like those illustrated in Fig. 10.4 have learned to subtract correctly. Research on the instructional effectiveness of the method is continuing, but the data in hand are sufficient to establish that the instruction can provide effective remediation of subtraction bugs.

There also is evidence that children acquire a better understanding of general principles as a result of mapping instruction. This evidence is provided in part by explanations that children are able to give after the instruction. One child, whom we call Laura, started with the smaller-from-larger bug. She learned the correct procedure, and three weeks later she still remembered how to subtract correctly. She was asked whether she remembered how she used to subtract, and what the difference was. Her answer was, "I used to take the numbers apart. Now I keep them together, *and* take them apart." This remark seems to indicate that Laura came to understand an important principle: that the set of digits that are on a line collectively represent a single number.

Another wise explanation was given by a student who started with a bug involving borrowing when a zero is encountered. This student, whom we call Molly, learned to subtract correctly, and in a posttest solved the problem 403 − 275, correctly decrementing the *four*, replacing the *zero* with a *nine*, and placing a small *one* next to the *three* in the top number. She mentioned that she changed the *four* to a *three* "because I traded it for 10 *tens*." Then she was asked, "Do you know where the *nine* came from?" Molly answered, "It's 9 *tens* and the other *ten* is right here," pointing to the *one* near the *three*. Molly's remark seems to indicate that she appreciated the requirement of keeping the value of a number the same during borrowing.

In theoretical research in which I have collaborated with Lauren Resnick, Robert Neches, and James Rowland, we have tried to characterize the knowledge that is acquired in mapping instruction, and some of the learning processes that occur when students receive this instruction. We are working with two general ideas, one of which has been implemented as a simulation of learning, based on the protocol given by Molly. A simulation of the other idea is still being developed.

In both of these ideas, we assume that the effect of mapping is to elicit a generalization across the two procedures that are learned by the student. The generalization involves entitites that are abstractions over the domains in which the procedures are defined. In the case of blocks and numeral subtraction, the entities that are acquired in our simulation are quantitative concepts for which both the numerals and the blocks provide symbolic representations.

The main structures involved in the simulation are shown in Fig. 10.6. We assume that initially, the knowledge structure includes the whole-parts schema, including a procedure for adjusting the sizes of the parts while keeping the whole quantity constant. Instruction in the procedure with blocks has resulted in acquisition of a procedure called Trade, where a block of one size is removed and ten blocks of the next smaller size are put back in its place. The amounts that are taken away and put back are understood to be equal, since there is a ten-to-one ratio of the sizes of the blocks.

In mapping instruction, a procedure of borrowing is taught, and explicit connections are made between the components of Trade, and the components of Borrow; that is, Take-Away corresponds to Decrement, and Put-In corresponds to Add. We hypothesize that this correspondence influences the acquisition of Borrow, through the mediation of a third structure which we call Exchange.

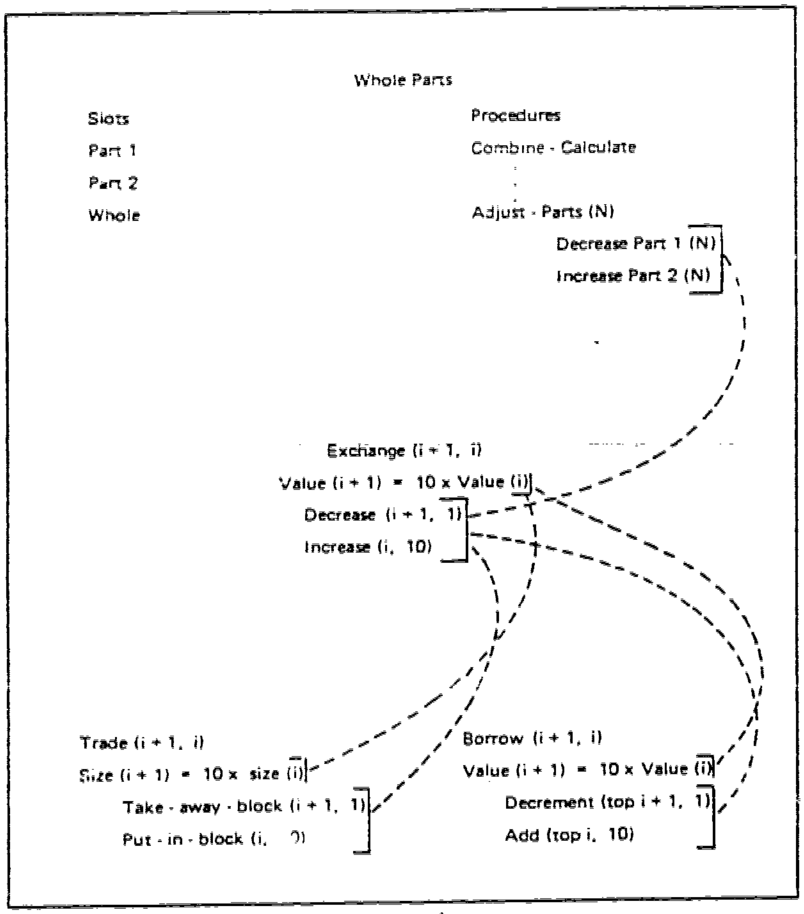

Fig. 10.6. Structures in simulation of learning from mapping instruction.

Exchange is a generalization across Trade and Borrow, and its components are propagated into the Borrow procedure. Decrease (i+1) and Increase (i) are generalizations of the surface-level actions Take-Away, Put-In, Decrement, and Add. The whole-parts schema provides a constraint that the amounts of increase and decrease should be equal. This is satisfied in Trade by the property of block size. We assume that a generalization of block size is included in Exchange as the property of Value, and that this is propagated into the Borrow procedure as a Value associated with the place of each digit.

The structures that our simulation acquires were designed to provide information of the kind needed for explanations like those given by Laura and Molly.

One important component is the concept of value, included in the Borrow procedure. This is an important general principle of numeration. Another important principle is that when borrowing occurs, the value of the number should remain the same. In our simulation, this principle is represented by the procedure's connection to the whole-parts scheme, and the constraint of its Adjust-Parts procedure. We provided our system with some primitive question-answering capability, and it can answer the question, "Where did the *nine* come from?" after it has borrowed through *zero* in a problem like 403 − 275. It finds the value of the block that it took away from the hundreds column, identifies the value of the nine *tens* as being part of the ten *tens* that it put back, and locates the other ten *ones* that it exchanged for one of the *tens*. Laura's answer about keeping the numbers together involves a more subtle use of information, which we have not simulated. However, we conjecture that the answer depends on conceptualizing the value of the numeral as a whole quantity, made up of parts corresponding to the values of the digits, and the concepts needed for this conceptualization are all included in our simulation.

The conceptual entities in this analysis are similar to those acquired in meaningful learning of geometry. In both cases, representations of problem situations include conceptual units that are interpreted as elements with part-whole relationships. In geometry, a conceptual entity represents a structure composed of two segments or angles that are combined in a whole segment or angle. In subtraction, there is a conceptual entity that represents the value corresponding to two adjacent digits, the sum of the values of the separate digits.

II. REASONING WITH GENERAL METHODS

The second function of conceptual entities that I propose is that they provide arguments on which general reasoning procedures can operate directly. In this section, I discuss findings that can be interpreted with this idea. First, analyses of processes in solving physics text problems suggest that experts' representations include entities that provide arguments to general procedures for reasoning about parts and wholes. Then, two experiments involving instruction provide further information about conditions that facilitate acquisition of representational knowledge that includes conceptual entities.

Physics Problems

In physics text problems, experienced problem solvers use representations in which forces, energies, momenta, and other abstractions are treated as entities. An example is in force diagrams, in which the collection of forces acting on an object in the problem is shown as a set of labeled arrows. The diagram shows various relations among these entities, such as opposition between pairs of forces acting in opposite directions. Chi, Feltovich, and Glaser (1981) have shown that abstract concepts such as conversion of momentum are salient for expert physicists when they are asked to classify problems into groups and when they are deciding on a method for solving a problem. McDermott and Larkin (1978) have simulated the process of forming representations based on abstract conceptual entities, such as forces.

I will discuss two specific examples in which representation using conceptual entities enable general reasoning procedures to be used. In both of these examples, the general procedures involve relationships between quantities that can be considered as parts of a whole. Tables 10.1 and 10.2 show partial protocols that were kindly made available by D. P. Simon and H. A. Simon. They were among the protocols obtained from a novice and an expert subject working on problems from a high school text (Simon & Simon, 1978). The problem for these protocols was the following: "An object is dropped from a balloon that is descending at a rate of four meters per sec. If it takes 10 sec for the object to reach the ground, how high was the balloon at the moment the object was dropped?"

In the novice's protocal, shown in Table 10.1, the process was one of search guided by a formula. Quantities in the problem text were interpreted as the values of variables. The subject applied some general constraints, such as a requirement that distances have positive values, but the protocol lacks evidence that velocities and accelerations functioned as conceptual entities.

In the expert's protocol, Table 10.2, there is a rather clear example of a conceptual entity, the "total additional velocity." The expert apparently represented the velocity that would be achieved at the end of a 10-sec fall as the sum of two components: the initial velocity, and the amount that would be added during the fall. The added amount can be found easily, because it is proportional to the time. Then the velocity at the end of the fall was found by combining its two components. The average velocity during the fall, needed to compute the distance, was found by averaging the initial and terminal velocities. Finally, the distance was found by multiplying the average velocity by the given duration.

A reasonable interpretation of this solution is that three general procedures for making quantitative inferences were used. One is a procedure for finding a whole quantity by adding its parts together. The second is a procedure that finds the average value of a quantity that undergoes linear change. The third is a procedure that finds the total amount of a quantity by multiplying its average rate during a time interval by the duration of the interval. All of these procedures correspond to physics formulas, but there is no evidence in the protocol that formulas were used in the solution. A plausible hypothesis is that the solution was obtained by forming representations of quantities that served as arguments for general inferential procedures. That inference seems particularly well justified in the case of the "total additional velocity," a quantity for which there is no specific variable in the formulas that are usually given.

Another example from physics is in the discussion in this volume by Larkin, regarding the loop-the-loop problem that deKleer (1975) discussed earlier.

TABLE 10.1
Novice Protocol

1. "An object dropped from a balloon descending at 4 meters per second."
2. 4 meters per second is v zero.
3. "lands on the ground 10 seconds later."
4. t equals 10 seconds.
5. "What was the altitude of the balloon at the moment the object was dropped?"
6. Now we want s equals v zero times the time plus one half of . . .
7. . . . a equals g equals in this case, minus 32. . . .
8. Oh, minus 9 point 8 meters a second.
9. It's descending at the rate of 4 meters per second.
10. One half g t squared.
11. that equals v zero,
12. which is 4,
13. times 10,
14. plus one half of minus 9 point 8,
15. equals minus 4 point 9 times.
16. . . . . Oh, we're going to come out with a minus number?
17. It was descending at 4 meters per second.
18. Oh, great.
19. "How high was the balloon?"
20. "An object dropped from a balloon descending at 4 meters per second"
21. "lands on the ground 4 seconds later."
22. It was already going. . . .
23. The initial velocity was 4 and not zero, that's it.
24. minus 4 point 9 times 100,
25. But this is its absolute . . . um . . .
26. We want its absolute value, don't we?
27. That equals 40 minus 49 hundred, that, obviously, . . .
28. . . . 4 . . . 4 hundred and 90 . . .
29. . . . 'cause it drops. . . .
30. Its initial velocity was 4., . . .
31. and starting from zero,
32. Now we've got something we really don't know how to handle.
33. Now we really don't know how to handle this.
34. Because it doesn't start from zero;
35. it started from 4 meters per second.
36. and the first second accelerates . . . so each one . . .
37. that initial velocity . . . starts at 4 and not zero.
38. So, I think it's 40 plus, because although it's a negative . . .
39. no, no, it's increasing.
40. Oh no, it's increasing, it's not slowing down.
41. Okay. So the distance equals 40 plus 4 hundred and 90
42. equals 5 hundred and 30 meters.
43. That's my answer.

TABLE 10.2
Expert Protocol

1.	"An object dropped from a balloon descending at 4 meters per second.
2.	lands on the ground 10 seconds later.
3.	What was the altitude of the balloon at the moment the object was dropped?"
4.	So it's already got a velocity of 4 meters per second
5.	and it accelerates at 9.8 meters per second per second
6.	so its final velocity 10 seconds later.
7.	well, let's say its total additional velocity 10 seconds later
8.	would be 98 meters per second per second
9.	and that . . . ah . . . plus the 4 that it had to start with
10.	would be 102 meters per second per second
11.	so its average velocity during that period
12.	would be 106 over 3 or 53 . . . ah . . . 53 meters per second
13.	and at 10 seconds that would mean it had dropped 530 meters.

Larkin notes that in place of the sequential envisionment procedure that deKleer described and analyzed, experts frequently represent the problem using the conservation of energy. In this representation, there is a quantity, the total energy, that remains constant. The total energy is made up of two components: the potential energy (associated with height) and the kinetic energy (associated with speed). As the ball moves downward, potential energy is converted to kinetic energy, which is then reconverted to potential energy as the ball moves up the other side. The requirement of the problem is satisfied if the amounts involved in the two phases are equal.

A reasonable interpretation of this solution includes another general inferential procedure involving additive combinations. If a whole quantity is constrained to be a constant, then one of its parts can be increased by a transfer from the other part. The use of a general procedure for inferring quantitative changes based on that principle in the loop-the-loop problem seems a reasonable conjecture.

Distance, Time, and Velocity

The interpretation that I proposed in the last section regarding expert problem solving in physics includes conceptual entities that are available as arguments for general methods of reasoning. A question that arises is how representational knowledge of that kind is acquired. Some suggestive findings were obtained in an instructional study conducted at Indiana University in 1967 (Greeno, 1976). The suggestion is that new conceptual entities can be acquired when procedures are learned that use those entities as arguments.

In the experiment, seventh-grade students were given instruction in solving problems about simple motion using the formula: distance = speed × time. Different groups received differing pretraining prior to the instruction. The pretraining that was effective included training in two kinds of procedures. One was observational: students were shown examples of simple linear motion and were given procedures for manipulating distance and velocity and for measuring distance and time. The other procedures were computational: students had practice in calculating one of the three quantities given the other two. Results of the study suggest that from these experiences students acquired representational knowledge in which distance, duration, and velocity were conceptual entities about which the students could reason in a direct, flexible manner.

The experiment took place in three consecutive daily sessions. In the first session a pretest was given. The second session was an instructional treatment that varied among groups of students. In the third session all of the students received some instruction in solving problems about motion and a posttest was given.

The instructional group of greatest interest was given experience with simple motion in a setup shown in Fig. 10.7. Model railroad tracks were marked at one-foot intervals. A timer, visible to the students, ran as an engine moved along the track. Velocity was variable from .5 to 3 feet per sec. A regulator was available to the students for one of the tracks.

In the instruction, a series of problems was presented to groups of four or five students. In each problem, two of the three quantities—distance, velocity, and duration—were given, and students calculated the third. When the unknown was distance or velocity, students performed the operations that determined the quantity, either by adjusting the transformer or by placing the photocell that stopped the timer. Each result was tested by running an engine. The correspondence between distance and time was noted as the engine moved along the track, a

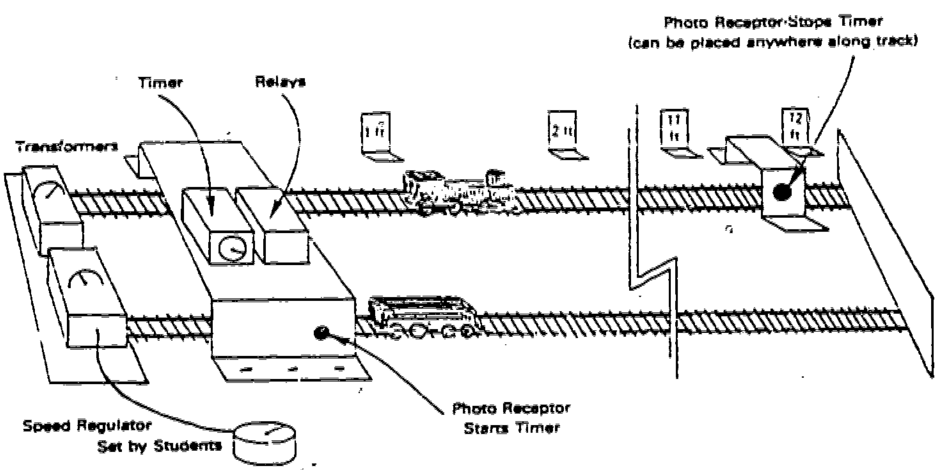

Fig. 10.7. Apparatus for distance-duration-velocity demonstrations (from Greeno, 1976).

record of results on all the problems was kept, and results of different combinations of quantities were discussed. A few problems with two engines moving simultaneously at different velocities were given at the end of the session.

The effect of this experience was compared with two other instructional groups and a control group. The other two instructional groups received experience of a more mathematical kind, involving the inverse relation of multiplication and division or use of ratios in solving problems. The fourth group went to a study hall.

The instruction that all students received on the third day was a straightforward presentation of the formula, distance = speed x time, with examples of its use in solving simple problems.

The tests that were given before and after instruction consisted of seven problems. Three were easy, requiring calculation of one of the three quantities from the other two, for example, "A man drove at a speed of 60 mile per hour for 4 hours. How far did he drive?" The other four problems were more complicated, requiring analysis of motions into components, either of durations or of distances. An example is, "The distance between Bloomington and Chicago is 240 miles, and there are two airline flights between the two cities. One flight is nonstop and takes 1½ hours. The other flight stops for ½ hour in Terre Haute, but also takes 1½ hours. How fast does each plane fly?" Pretest and posttest problems were variants of each other, involving different kinds of moving objects and different numbers.

The best posttest performance was given by the group with experience with the model trains. On the four complicated test problems, that group improved by an average of 1.21 problems between pretest and posttest, the control group improved by .57 problems, and the other instructional groups improved by .21 problems.

An interpretation that seems reasonable is that students who received experience with model trains acquired representational knowledge in which distance, velocity, and duration were conceptual entites. The complicated problems on which they excelled required combining parts of a trip. The students' ability to solve these problems suggest that their representations of quantities in problems were in a form that enabled them to be used by general reasoning procedures associated with a whole parts schema or other similar structures. A plausible conjecture is that entites may have resulted from the students' acquisition of observational and computational procedures that operated directly on the quantities of distance, duration, and velocity.

Sound Transmission

The last example I discuss in this section also involves an instructional experiment. This study was motivated by discussions of mental models as mechanisms of reasoning. In analyses such as Stevens and Collins' (1978) discussion of inferential reasoning about weather, knowledge about the detailed internal structure of processes enables individuals to generate conjectures about the behaviors of the processes in new conditions. In a study in which I collaborated with Gregg T Vesonder and Amy K. Majetic, we investigated the question whether instruction regarding the detailed causal structure of sound transmission would enhance students' ability to reason about properties of that process. A full report of this experiment is available in another study (Greeno, Vesonder, & Majetic, in prep.).

We designed two instructional units about transmission of sound. One was patterned after the usual textbook sequence, focusing on amplitude and frequency of sound waves. We refer to this instruction as a Steady-State unit, since it focused on temporal properties of sound waves at a single point in space: alternating compressions and rarefactions varying in amplitude and frequency. We gave a simpler discussion than is often used in texts. We made no attempt to discuss longitudinal waves, restricting our discussion to transverse waves consisting of alternating compressions are rarefactions. We also related the properties of waves to concrete phenomena, using a guitar to produce tones varying in loudness and pitch. The mechanism of transmission was discussed, mainly in the context of these properties. A Slinky toy was used to show transmission of a transverse wave, and a piece of plastic foam with dots painted on it was used to model compressions and rarefactions. Waves with varying amplitudes and frequencies were illustrated with both of these models and related to differing sounds made with a guitar.

We refer to the other instructional unit that we designed as a Transmission unit. It focused on the causal mechanism of sound transmission. The idea of a pulse was modeled using a row of dominoes and was reinforced using a tube covered on both ends with balloon rubber, so that pressing on one end caused the other end to bulge. A Slinky toy was used to show a pulse moving through a medium, and foam rubber with painted dots was used to model compression of molecules. Finally, a shallow round dish containing water was used to show that a pulse moving from the center is distributed over a greater area and therefore becomes weaker at any single point. After showing all these aspects of transmitting a pulse, we discussed sound waves as alternating increases and decreases in pressure caused by a vibrating source, and illustrated the effects of that with each of the models.

Our two instructional units can be considered as containing a common core of information, elaborated in different ways. The common information was about the components of sound transmission: the requirements of a source, a medium, and a detector, and some basic causal relations involving vibrations, compressions, and rarefactions. In the Steady-State unit, this information was elaborated by discussing attributes of sounds, identifying properties of pitch and loudness that vary between different sounds and relating these to variables of frequency and amplitude in the theoretical system of sound transmission.

In the Transmission unit, the basic information was elaborated by a more detailed discussion of the causal mechanism of sound, using the simpler case of a pulse to make the causal system easier to understand. This instruction was designed to teach the microstructure of the causal system. We anticipated that this might enable students with Transmission instruction to reason more successfully about situations involving transmission of sound than their counterparts, whose instruction focused more on attributes and less on the causal structure. This anticipation was not borne out in the results.

We tested our sixth-grade student subjects by asking a set of 12 questions. Their answers were tape recorded and transcribed, and we evaluated them using an analysis of propositions that would constitute correct knowledge and understanding. We were particularly interested in four questions that required inferences about sound transmission. One involved a simple application of knowledge that sound will be softer at a greater distance. A second question required the inference that sound will not be transmitted through a vacuum, but that it will be transmitted through water. The other two questions required conjectures about rates of transmission: one that sound could travel faster through one medium than another, and the other that one form of energy might travel faster than another.

To our surprise, scores on these inferential questions were not significantly different among students who had different units of instruction. Indeed, students who received either or both units did not differ from students in a control condition who received neither unit. The trend favored the students in the Steady-State condition, in opposition to our expectation of an advantage due to the Transmission unit.

This finding was reinforced by a more detailed analysis of evidence for knowledge of specific propositions. We divided propositions into four sets, judging whether each proposition was included explicitly in the Transmission unit, the Steady-State unit, both units, or neither unit. On propositions that were in both units, there was a nearly significant difference favoring the Steady-State unit. On propositions that were in only the Transmission unit, students with only Steady-State instruction did as well as students with Transmission instruction. This was not a symmetric finding: on propositions that were only in the Steady-State unit, Steady-State students were much better than Transmission students.

The students' responses to questions suggested that most of them learned about the requirements for sound transmission: a source, a medium, and a receptor. All except four of the 20 students correctly said that sound would not be transmitted through a vacuum when air was pumped out of a jar with a bell in it. Thirteen of the 20 students correctly said that sound would be transmitted if the jar were filled with water. The number of correct answers about either the vacuum or the water did not depend on the instruction that students received.

On the two questions requiring conjectures about velocities of transmission, correct answers were given by only six, and four of the 20 students, and there was no relationship between the answers and the instruction that students had received. Apparently the knowledge that they acquired about sound did not make contact with their general concepts about faster or slower motion. Several students gave answers indicating that the concepts of source, medium, and receptor were applied in answering the questions. One question asked why lightning is seen before thunder is heard: six students conjectured that lightning occurs earlier. The other question asked why a train is heard sooner if your ear is close to the railroad track: 15 students conjectured that the rail becomes a source of sound, being caused to vibrate by the wheels of the train.

The conclusion that we reach is that both of our instructional units probably led to acquisition of conceptual entities corresponding to the components of sound transmission: a source, a medium, and a receptor. This acquisition did not seem to be strengthened substantially by explanation of the detailed causal structure of the system. Of course, we may have chosen poor questions in trying to tap that knowledge. The main opportunity to show improved performance require conjectures about speed of transmission, a global property. The difficulty could have been in children's making contact between their knowledge of sound and their general knowledge about motions with differing speeds, rather than a lack of representational knowledge about sound. Even so, we are led to conclude that knowledge of the detailed causal structure of a mechanism may not be as useful an instructional target as knowledge of attributes that are directly relevant to question-answering and other target tasks.

III. COMPUTATIONAL EFFICIENCY

The hypothesis that appropriate conceptual entities can enable more efficient computation is probably obvious. I present a single example in which the point is illustrated with unusual clarity.

Monster Problems

An example in which alternative representations of problems have been analyzed in detail is a set of puzzles about monsters and globes that are isomorphs of the Tower of Hanoi problem, analyzed by Simon and Hayes (1976). The entities that are involved in this example are sets of objects, and the procedures for which the entities are arguments are operations on sets, such as finding the largest member of a set.

Simon and Hayes classified problems into two categories, called Transfer and Change problems, which differ in the way that applicability of operators depends on attributes and entities. The distinction was very significant empirically: Change problems were about twice as difficult as Transfer problems.

To illustrate the problem categories, consider two problems in which there are three monsters each holding a globe. The monsters and globes both vary in size:

the sizes are small, medium, and large. Initially, the small monster holds the large globe, the medium monster holds the small globe, and the large monsters holds the medium globe. The goal is a situation in which the size of each monster matches the size of the globe that it is holding.

In the Transfer problem, globes are moved from monster to monster. Only one globe can be moved at a time, a monster can only give away its largest globe, and the transferred globe must be larger than any the receiving monster is holding prior to the transfer.

In the Change problem the sizes of globes are changed by shrinking and expanding. To change a globe from its initial size to some terminal size, the monster holding the globe must be the largest monster currently holding a globe of its initial size, and no larger monster may be holding a globe of its terminal size.

To explain the greater difficulty of Change problems, Simon and Hayes suggested a plausible hypothesis about the representation of states and operators. In the representation of a state: (1) there is a list of the monsters; (2) each monster's size is an attribute; (3) a list of the globes held by each monster is a second attribute; and (4) each globe's size is an attribute of the globe. The operator for the Transfer problems has the form Move(GS, MS1, MS2), which means "Move the globe of size GS from the monster of size MS1 to the monster of size MS2." The operator for the Change problems has the form Change(MS, GS1, GS2), which means, "Change the globe held by the monster of size MS from its present size GS1 to size GS2."

The problems differ in a way that involves conceptual entities. The list of globes held by each monster is an entity in the representation; the lists are included in the initial representation of the problem, and are modified after each change in the problem state. These entities are used directly in the Transfer problems. To test whether move(GS, MS1, MS2) can be applied, the solver retrieves the lists of globes held by monsters MS1 and MS2 and determines whether globe size GS is the largest of both sets. The corresponding test in the change problems does not use entities in the representation, and requires construction of lists that are to be tested. Testing applicability of change(MS, GS1, GS2) involves retrieving the monsters holding globes of size GS1 and GS2, and testing whether monster size MS is the largest of both of these sets. The sets have to be constructed, since the lists of monsters holding globes of the three sizes are not entities in the representation.

Simon and Hayes' suggested explanation has not been confirmed empirically, and they are continuing their experimental research on the problem (H. A. Simon, personal communication). There probably are several factors that contribute to the difference in difficulty between the two kinds of problems. Even so, their hypothesis is plausible and provides an especially clear example of the importance of conceptual entities in problem representation.

IV. PLANNING

The final hypothesis considered in this essay is that the ontology of a problem domain has important effects on goal definition and planning. This point is illustrated by results of another set of instructional studies.

Binomial Probability

In the early 1970s, Richard Mayer, Dennis Egan, and I conducted a series of experiments (Egan & Greeno, 1972; Mayer, 1974; Mayer & Greeno, 1972; Mayer, Stiehl, & Greeno, 1975) in which we gave instruction in the formula for binomial probability:

$$P(R|N) = \binom{N}{R} p^R (1-p)^{N-R},$$

where N is a number of trails, R is a number of success outcomes, and p is the probability of success on each trial. The studies involved comparisons between alternative instructional conditions. In most of the experiments we compared two sequences of expository instruction. One sequence focused attention on calculation with the binomial formula. The other sequence emphasized meanings of concepts, providing definitions of variables in relation to general experience and giving explanations about how the concepts combine to form components of the formula. The conceptual instruction discussed outcomes of trials and sequences of trials with different outcomes, and defined the probability of R successes as the sum of probabilities of the different sequences that include R successes. We also compared expository learning that emphasized the formula with discovery learning, and obtained similar results to those we found with formula and conceptual emphases.

Our interpretation of these studies was that conceptual expository instruction and instruction by discovery led to knowledge that was more strongly connected to the students' general knowledge than the knowledge that was acquired in expository instruction that emphasized the formula. That still seems a correct interpretation, but a more specific hypothesis may be warranted. It seems likely that conceptual instruction and discovery learning may have facilitated formation of conceptual entities corresponding to the variables and that these were less likely to be acquired by students whose instruction emphasized calculation with the formula.

Several of the findings of our experiments are consistent with this interpretation. First, students with conceptual or discovery instruction were able to solve story problems nearly as easily as they could solve problems with information presented in terms of the variables of the formula, whereas for students with formula instruction story problems were considerably more difficult. This is consistent with the idea that conceptual entities facilitate interpretation of problem information in novel contexts.

Three further findings can be interpreted as indications that conceptual entities facilitate planning. First, some of the problems that we presented had inconsistent or incomplete information and hence were unsolvable. For example, one problem gave $R = 3$, $N = 2$, and $p = \frac{1}{2}$ and asked for $P(R|N)$. The information is inconsistent, because there cannot be more successes than trials. The students with conceptual instruction identified these as unsolvable problems more frequently than students with formula instruction. Students with conceptual instruction also were better at solving problems in which the probability of a specific sequence of outcomes was requested, rather than the probability of a number of success outcomes. We called a third kind of problem Luchins problems, because Luchins (1942) studied performance on similar problems extensively. These were problems in which the answer could be found by simple direct means, but if students tried to apply formulas they could be led into a complicated sequence of fruitless calculations. An example was the following: "You play a game five times in which the probability of winning each time is .17, and the probability of winning three games out of five is .32. What is the total number of successes plus the total number of failures?" Luchins problems were almost as easy as ordinary problems about binomial probability for students who had discovery learning, but they were much harder than ordinary problems for students with expository learning.

All three of these findings are consistent with the idea that a representation with conceptual entities corresponding to the variables enables a problem solver to reason directly about the quantities rather than simply through the medium of the formula. The conceptual instruction gave more emphasis to discussion of sequences of their outcomes and their properties. Thus, it seems likely that in conceptual instruction, students gained representational knowledge enabling them to interpret problems and questions in terms of individual sequences when that was appropriate. This would provide information that could be used directly to determine the problems were incoherent, to identify problem goals involving individual sequences rather than the quantity given by the binomial formula, and to find direct solution methods.

V. CONCLUSIONS

In this essay I have explored hypotheses about ways in which representational knowledge can influence problem solving. The discussion has been focused on effects of an aspect of representation that I have referred to as the ontology of a problem domain, the kinds of conceptual entities that are included in representations of problem situations. I have presented interpretations of several empirical findings and theoretical analyses that indicate four ways in which ontology can influence problem solving: by facilitating the formation of analogies between domains, by enabling use of general reasoning procedures, by providing efficiency, and by facilitating planning.

The idea of problem ontology raises significant issues relevant to instruction and the acquisition of cognitive skill. It seems important to design instruction so that students will acquire the conceptual entities that are needed for representing problems in the domain, as well as acquiring the procedures needed to make the calculations and inferences required for solving problems. Three studies described in this essay provided evidence of successful instruction that can be interpreted as acquisition of conceptual entities. In each of these the procedures that were taught were related to other information of various kinds. In mapping instruction for arithmetic, the procedure of multidigit subtraction with numerals was related to an analogous procedure of subtraction with place-value blocks. In instruction for solving problems about simple motion, the procedures for calculating answers were related to observational experience and procedures for manipulating and measuring values of the variables. And in instruction for solving problems using the binomial formula, the instruction that led to better understanding provided relationships between the computational formula and general concepts of trials, outcomes, and sequences. These findings suggest a general principle: perhaps the acquisition of cognitive entities is most effective when variables in procedures are related to other entities in cognitive structure. The kinds of relationships that can be useful in this way are clearly quite variable; on the other hand, we cannot expect everything to work, as evidenced by the results of our experiment on sound transmission. A detailed theory of learning will be required to characterize the favorable conditions specifically, but it seems reasonable to propose that the acquisition of the ontology of a domain is one of the significant issues to be addressed in our study of learning processes.

ACKNOWLEDGMENT

This research was supported by the Personnel and Training Research Programs, Office of Naval Research, under Contract Number N00014-79-C-0215. Contract Authority Identification Number NR 667-430.

REFERENCES

Anderson, J. R., Greeno, J. G., Kline, P. J., & Neves, D. M. Acquisition of problem-solving skill. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition.* Hillsdale, N.J.: Lawrence Erlbaum Associates, 1981.

Bobrow, D. G., & Winograd, T. An overview of KRL, a knowledge representation language. *Cognitive Science,* 1977, *1,* 3–46.

Brown, J. S., & Burton, R. R. Diagnostic models for procedural bugs in basic mathematical skills. *Cognitive Science,* 1978, *2,* 155–192.

Chi, M. T. H., Feltovich, P., & Glaser, R. Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 1981, *5*, 121–152.

deKleer, J. *Qualitative and quantitative knowledge in classical mechanics* (Tech. Rep. AI-TR-352). Cambridge, Mass.: Artificial Intelligence Laboratory, Massachusetts Institute of Technology, December 1975.

Egan, D. E., & Greeno, J. G. Acquiring cognitive structure by discovery and rule learning. *Journal of Educational Psychology*, 1972, *64*, 86–97.

Greeno, J. G. Some preliminary experiments on structural learning. In J. M. Scandura (Ed.), *Structural learning II. Issues and approaches*. New York: Gordon and Breach Science Publishers, Inc., 1976

Greeno, J. G., Vesonder, G. T., & Majetic, A. K. *Conveying knowledge about a process*. In preparation.

Luchins, A. S. Mechanization in problem solving: The effect of Einsrellung. *Psychological Monographs*, 1942, *54* (6, Whole No. 248).

Mayer, R. E. Acquisition processes and resilience under varying testing conditions for structurally different problem-solving procedures. *Journal of Educational Psychology*, 1974, *66*, 644–656.

Mayer, R. E., & Greeno, J. G. Structural differences between learning outcomes produced by different instructional methods. *Journal of Educational Psychology*, 1972, *63*, 165–173.

Mayer, R. E., Stiehl, C. C., & Greeno, J. G. Acquisition of understanding and skill in relation to subjects' preparation and meaningfulness of instruction. *Journal of Education Psychology*, 1975, *67*, 331–350.

McDermott, J., & Larkin, J. H. Representing textbook physics problems. In *Proceedings of the Second National Conference. Canadian Society for Computational Studies of Intelligence*. Toronto, Canada: University of Toronto, 1978.

Resnick, L. B. A developmental theory of number understanding. In H. P. Ginsburg (Ed.), *Children's Knowledge of arithmetic*, in press.

Simon, D. P., & Simon, H. A. Individual differences in solving physics problems. In R. S. Siegler (Ed.), *Children's thinking: What develops?* Hillsdale, N.J.: Lawrence Erlbaum Associates, 1978.

Simon, H. A., & Hayes, J. R. The understanding process: Problem isomorphs. *Cognitive Psychology*, 1976, *8*, 165–190.

Stevens, A. L., & Collins, A. *Multiple conceptual models of a complex system*. (BBN Report No. 3923). Cambridge, Mass.: Bolt Beranek and Newman, Inc., August 1978.

Wertheimer, M. *Productive thinking* (Enlarged ed.). N.Y.: Harper & Row, 1959, (originally published 1945).