1991

Chapter 15

A Computational Account of Children's Learning About Number Conservation

Tony Simon

Allen Newell

David Klahr

1. Introduction

This chapter focuses on the conceptual domain known in developmental psychology as "conservation of quantity". Acquisition of quantity conservation constitutes a fundamental part of cognitive development, and hundreds of studies have been published in which children's "naturally acquired" conservation concepts are assessed, or in which attempts are made to teach such concepts to children. We review some of the major regularities emerging from this literature, and then we describe one particular training study in detail. Following that, we describe Q-Soar, a computational model that simulates the acquisition of number-conservation knowledge.

The particular study that we examine is Gelman's (1982) training of three and four year old children, and we use Q-Soar to model the learning behavior exhibited by children in this study. This model is constrained by a unified theory of cognition, which is the Soar architecture (Lewis et al., 1990; Newell, 1990). The version of Q-Soar to be described in this chapter accounts only for the learning that occurs in Gelman's study. However, we will suggest that chunking, Soar's learning mechanism, which is the only learning process included in Q-Soar, could account for the acquisition of all conservation knowledge.

Our study of conservation learning and human development can be viewed as an instance of research in concept formation. The nature of knowledge acquisition in cognitive development is consistent with the concerns of work on incremental processing. In addition, the feedback that is naturally available during development rarely has the explicit nature that is assumed in supervised learning models (Fisher & Pazzani, Chapter 1, this volume). In fact, studies of conservation training (such as Gelman's) demonstrate that more explicit forms of feedback considerably accelerate the acquisition of conservation concepts. However, even under training conditions, the type of learning that is discussed here is in line with theory-driven approaches to concept formation (Fisher & Pazzani, Chapter 6, this volume), notably the construction of *ad hoc* concepts (Barsalou, 1983) and *explanation-based* methods (Mooney, this volume). We now turn our attention to conservation learning, and then to the particular system that we have developed to model it.

2. The Phenomenon of Conservation

A central tenet of Piagetian theory (Piaget, 1952, 1970) is that the acquisition of conservation knowledge is a crucial step in the child's development of mature conceptual capabilities. Piaget (1968, p. 978) defines conservation as follows:

> We call "conservation" (and this is generally accepted) the invariance of a characteristic despite *transformations* of the object or of a collection of objects possessing this characteristic. Concerning number, a collection of objects "conserves" its number when the shape or disposition of the collection is modified, or when it is partitioned into subsets.

Children's knowledge about the effects of transformations must be empirically derived in the first instance because all transformations have different effects on different physical dimensions of the transformed material. For example, whether or not the *pouring* transformation conserves quantity depends on what is poured and what is measured:

> If we pour a little sugar into red sugar water, we do not change temperature, amount, height, width, or redness, but we increase sweetness. If we add more of an identical concentration, we do not change temperature, redness or sweetness; however the

> amount increases, as does liquid height, but not width (in a rigid container). On the other hand, if we add water, we increase two extensive quantities (amount, liquid height), reduce two intensive quantities (redness, sweetness), and leave one unchanged (temperature) (Klahr, 1982, pp. 68–69).

Therefore, a central component of what must be learned, either in training studies or "naturally acquired" by the child outside the laboratory, are the linkages between transformational attributes and their dimensional effects as measured in a wide variety of contexts.

The centrality of conservation concepts to most theories of cognitive development has produced a vast database of empirical results. Nevertheless, a computational model that can account for the regularities has yet to be fully specified. There have been both structural and processing accounts of the knowledge *used* by a child who "has" conservation, as well as global characterizations of the acquisition of that knowledge, such as Piaget's assimilation and accommodation processes, Klahr and Wallace's (1976) time-line processing, and Halford's (1982) levels of cognitive systems. However, neither these nor any other accounts have completely stated a set of operations and their interaction with a specified learning mechanism and shown this to produce the pattern of behavior observed in children acquiring conservation knowledge.

Q-SOAR is a model of the acquisition of conservation knowledge designed to meet several desiderata for computational models of developmental phenomena:

1. Such models should be based on a principled cognitive architecture, rather than as a set of arbitrary and ad hoc mechanisms. For Q-SOAR, the architecture is SOAR, to be described in Section 5.

2. Computational models should be constrained by the general regularities in the relevant empirical literature. There are a number of such regularities, i.e., findings that are consistently reported and for which there is little or no disconfirming evidence.

    (a) Young children in the three to four year age range can consistently obtain accurate specific quantitative values for small sets of objects (up to four) (cf. Fuson, 1988).

    (b) Young children in this age range cannot consistently obtain accurate specific quantitative values for large sets of objects (more than five or six) (cf. Fuson, 1988).

(c) Children from three years of age are able to correctly produce a cardinal value for three or four items in under one second by the process of subitizing[1] (Campbell, Cooper, & Blevins-Knabe, 1988; Chi & Klahr, 1975; Svenson & Sjoberg, 1983).

(d) Subitizing has a limit of four items in these children (Campbell et al., 1988; Chi & Klahr, 1975; Svenson & Sjoberg, 1983).

(e) Children who have not fully acquired conservation knowledge can still correctly answer conservation questions when they can obtain a specific quantitative value for the objects concerned (cf. Siegler, 1981).

(f) Children who have not fully acquired conservation knowledge do not correctly answer the conservation question when they cannot obtain a specific quantitative value for objects concerned (cf. Halford & Boyle, 1985).

(g) Young children can respond correctly and provide explanations for tests of conservation of quantity when discrete materials are used (such as counters or cookies) before they can do this on tests involving materials with continuous properties, such as columns of water or areas of space (cf. Gelman & Baillargeon, 1983).

3. Computational models should generate the same behavior as do the children in the specific domain being modeled. More specifically, they should compute an approximation of subjects' final knowledge states, given an approximation of initial states and external inputs like those imposed by experimental and/or natural conditions.

Although more than 20 years have passed since Klahr and Wallace (1970) proposed an information-processing approach to cognitive development, as yet there are no computational models of any major developmental transitions that satisfy all of these criteria. The Klahr and Wallace work on the development of quantitative concepts (Klahr, 1973, 1984; Klahr & Wallace, 1973, 1976) consists of verbal descriptions, flow charts, and production-system models of distinct performance levels in the general domain of quantitative reasoning, including subitizing, counting, estimation, class inclusion, transitive reasoning, and quantity conservation. However, with respect to transition processes, their most

1. Subitizing is a fast and accurate process of determining the numerosity of small sets of entities.

fully developed model (Wallace, Klahr, & Bluff, 1987) goes only as far as a partially specified architecture for supporting developmental transitions.

More recent computational accounts of developmental phenomena have been of two kinds:

1. One type of account is highly constrained by data from empirical studies of children's acquisition of knowledge in a domain, but the computational model itself is not constrained by any theoretical principles. Instead, it is based on pragmatic decisions about how to implement a set of assumed mechanisms (e.g., Siegler, 1991).

2. The other type of account is based on a very broad set of theoretical assumptions that are consistent with a wide range of specific implementations, such as the adaptive production system used by Halford et al. (1991) to model the acquisition of transitive inference or the connectionist model used by McClelland (1991) to model the acquisition of balance scale rules. Computational models of this type, while suggesting interesting learning mechanisms, tend to be relatively unconstrained by any particular empirical results on children's knowledge acquisition.

The purpose of our approach is to formulate a model that is tightly constrained by both a general theory of the cognitive architecture and a specific set of empirical results. Q-SOAR's challenge is to demonstrate that it can model the learning reported in Gelman's (1982) training study, which we describe next. The further issue of explaining the general developmental regularities listed above is addressed in the final section of the chapter.

3. A Training Study

In constructing a simulation of training studies, we were faced with the choice of modeling either our own arbitrary view of the essential properties of a typical training situation, or one specific training situation chosen from the vast conservation training literature. The problem with the former choice is that there is no "typical" training study. Detailed examination of the literature on conservation training studies reveals that they vary along so many potentially relevant dimensions that it is nearly impossible to get agreement even on a prototypical training study, let alone a set of defining properties. For example, Field's (1987) review organizes a collection of 25 recent conservation training studies with preschoolers along nine dimensions and three theoretical orientations.[2] Without any principled basis on which to construct a typical study, we chose to simulate a specific training study with well-defined procedures and clear quantitative outcomes.

3.1 Gelman's Training Procedure

As noted, we chose a training study reported by Gelman (1982) in which three and four year olds were trained in a brief session using small collections of discrete objects ($N = 3 - 4$) in both equivalence (two rows of equal number) and inequivalence (two rows of unequal number) relations, and in which the transfer test included both small ($N = 4-5$) and large ($N = 8 - 10$) collections. Gelman trained one group and used two types of control groups. Children in the Experimental group were trained with two types of collections in counterbalanced order. Half the children were first shown an equivalence relation (two rows of four items each), and the other half were first shown an inequivalence relation (one row of four and one row of three). In both equivalence and inequivalence collections, the items were initially placed in one-to-one correspondence.

For each type of collection there were nine steps, as illustrated in Figure 1: (1) The display was presented in one-to-one correspondence and the child was instructed to count the number of items in one of the rows. (2) That row was covered by the experimenter and the child was asked, "How many are under my hands?" (3) The child was instructed to count the number of items in the other row. (4) That row was covered by the experimenter and the child was asked, "How many are under my hands?" (5) The child was asked to judge whether the two uncovered rows contained "the same number or a different number" of items. (6) While the child watched, the length of one of the rows was spread or compressed. (7) The experimenter pointed to the altered (or unaltered) row and asked, "Are there still N here?" (8) The experimenter pointed to the other row and asked the same question. (9) The child was asked

2. The procedural dimensions were design, pretest, training, materials, reinforcements, verbal rule instruction, post-test, justifications, and delayed post-test. The theoretical orientations were specific experience, cognitive readiness, and perceptual readiness. Space does not permit an elaboration of these "guiding models", as Field calls them, but it is clear that training studies vary widely along both procedural and theoretical dimensions.

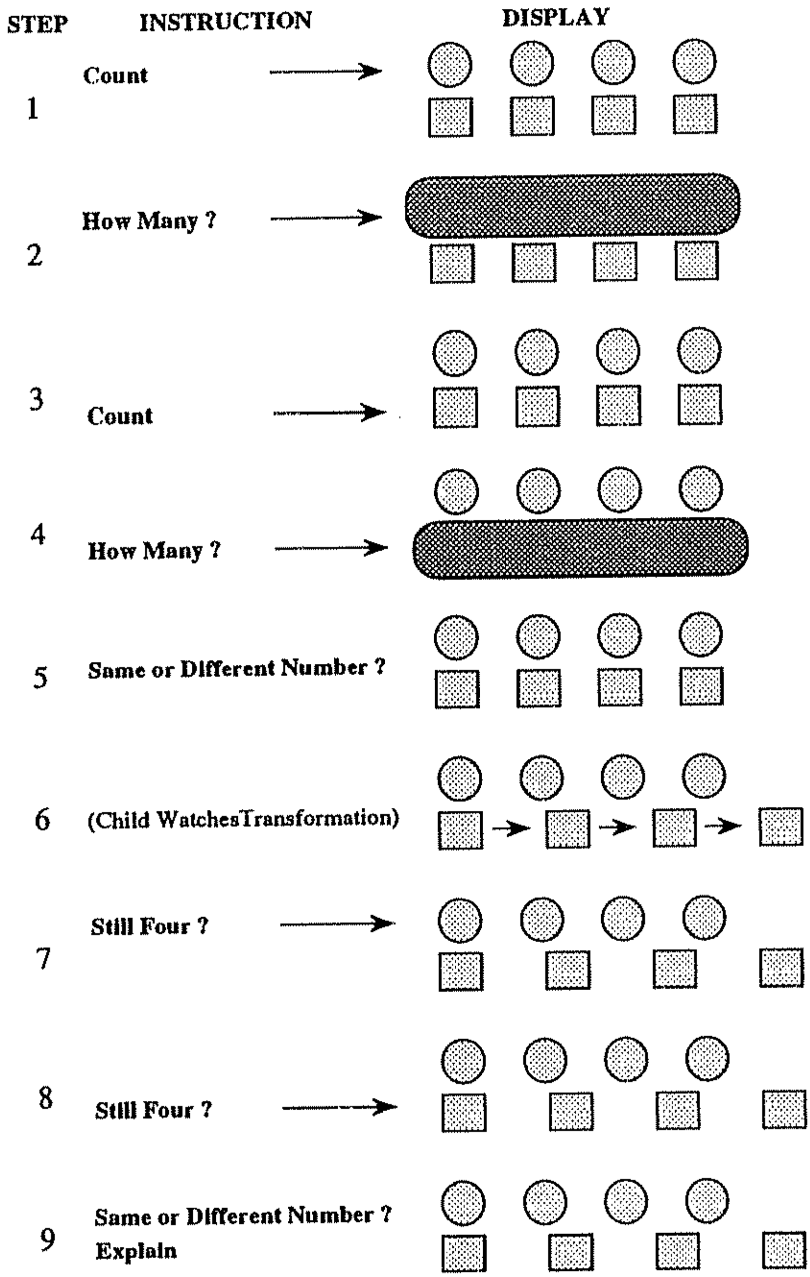

Figure 1. Graphical representation of the Experimental procedure.

whether the pair of rows had the same number or a different number of items, and to explain his/her judgment. All children answered the questions above correctly (except for one three year old who needed a slight extra prompt).

Gelman used two control groups. Children in the Cardinal-Once group were exposed to only one row (of three or four items). For that one row, they were exposed to steps 1–2 and 6–7 listed above. (Each row was altered four times to provide a comparable number of counting trials between the Experimental and control groups.) The other control group (No-Cardinal) simply counted single rows of three or four items, but the children in that group were not asked to "indicate the cardinal value rendered by the count".

3.2 Conservation Test

Immediately following the Experimental or other procedures, conservation tests were administered. Each child was given four different conservation tasks (large or small set size, and equal or unequal numbers of items in the two rows). Small sets included either four and five or five and five items, and large sets included either eight and ten or ten and ten items.

> The order of presentation of large and small set sizes was counterbalanced as was the order of conservation of equality and inequality tasks within a set-size range. The equal arrays were equal in length prior to the transformation and unequal in length after being transformed. The reverse was true for the nonequivalent arrays: before being transformed they were unequal in length and then equal in length after the transformation. The conservation trials were run in the standard way older children are tested, and *included requests for explanations*. Likewise, children were discouraged from counting. (Gelman, 1982, p. 213)

Because children were discouraged from counting, and because one or both of the rows had at least five items, and because children of this age do not count beyond three or four items very reliably (Fuson, 1988), it is likely that the equivalence (or non-equivalence) of both large and small arrays was established by one-to-one correspondence.

Table 1. Proportion of correct conservation judgments (over all four judgments and all subjects) in each condition, derived from Table 1 in Gelman (1982).

AGE & SET SIZE	EXP'L (N=32)	CARD'L-ONCE (N=24)	NO-CARD'L (N=16)
3's ON SMALL SET SIZE	71	11	13
4's ON SMALL SET SIZE	75	58	16
3's ON LARGE SET SIZE	72	8	0
4's ON LARGE SET SIZE	65	34	13
OVERALL THREES	71	9	6
OVERALL FOURS	70	46	15

3.3 Results

For both the large and small sets, there was almost no difference in the equal and unequal set sizes, so those results will be collapsed in the following discussion. Table 1 shows the overall proportion of correct judgments on conservation tasks. The effect of condition is striking: overall, the Experimental groups passed about 70% of the conservation trials, compared to passing rates from 0% to 15% for the "untrained" (No-Cardinal) groups. Trained threes and fours did equally well on large and small sets.

The interesting difference between the threes and fours occurred in the Cardinal-Once groups. For threes, Cardinal-Once training had no effect, but for fours it had a substantial effect when tested on both small and large set sizes. This is important, because the children in the Cardinal-Once group were trained only in identity conservation (transforming a single row), rather than equivalence conservation (transforming one of a pair of rows).[3] That is, they were never trained to notice that the relation between two different collections remained the same under a perceptual transformation, nor could they use one-to-one correspon-

3. See Klahr (1984) for a full discussion of the difference between identity conservation (IC) and equivalence conservation (EC). Note that Klahr's account of the acquisition of conservation rules is presented entirely in terms of the simple IC situation.

Table 2. Reference to transformation or one-to-one correspondence of Gelman's explanation categories.

	TRANSFORMATION	ONE-TO-ONE
IRRELEVANT TRANSFORMATION	×	
ADDITION/SUBTRACTION	×	
INITIAL EQUALITY/INEQUALITY	×	
ONE-TO-ONE CORRESPONDENCE		×

dence to reason about the effects of the transformations to which they were exposed. Instead, they could only learn that the initial and final cardinal value of a collection remained unchanged under certain kinds of transformations (i.e., spreading and compressing). Apparently, many four year olds, though few three year olds, were able to learn about transformations without the further help of one-to-one correspondence.

Gelman's categorization of children's explanations for their correct responses are presented in Table 2 in terms of (our interpretation of) whether the explanation makes reference to the transformation or to one-to-one correspondence. Gelman gives examples of two categories: *Irrelevant transformation* ("They just moved.") and *Addition/Subtraction* ("You need another one to make them the same.") explanations. *Initial equality/inequality of number* explanations presumably stated that the original value still held, while the content of *One-to-One Correspondence* explanations is obvious. The majority of the children's explanations referred to the transformation, and there were many more of these transformationally referenced explanations than there were explanations in terms of one-to-one correspondence. More specifically, for the Experimental threes, Experimental fours, and Cardinal-Once fours, the proportion of transformationally referenced explanations was 61%, 81%, and 65%, respectively, while for the same groups, one-to-one correspondence was used for only 21%, 9%, and 15% of the explanations.

Three year olds did not benefit from Cardinal-Once training, and four year olds in the Experimental group benefited more than did their age-mates in the Cardinal-Once group. Rather than attribute these differences to the role of one-to-one correspondence, we note that subjects in the Experimental group, but not in the Cardinal-Once group, got repeated exposure to observations of the following transitive relation. When, for example, two rows of objects have the same number, and after spreading the transformed row has the same number as before, then the untouched row and the transformed row still have the same number of objects.

Note that it is not difficult for the child to compute the effect of the transformation. First, each set was counted before and after the transformation for the Experimental trials, rendering one-to-one matching redundant. Second, in Experimental trials, the pre- and post-transformation information is visible in the form of a transformed and an untransformed row after the trànsformation has taken place. However, in the Cardinal-Once trials, memory of the pre-transformation information is always required to compute the transformation's effect. All of the three year olds and some of the four year olds apparently needed this additional information (and reduction of processing) that was provided to the Experimental group.

4. A Theory of Number Conservation Knowledge

We can summarize our view of the important difference between the Experimental group and the Cardinal-Once group as follows. Subjects in the Experimental group were exposed to equivalence (or inequivalence) conservation trials in which they observed and encoded an initial quantitative relation between two collections, and then observed a quantity-preserving transformation on one of the collections. They then requantified both collections and noted that the relation had not changed. In contrast, subjects in the Cardinal-Once group, because they were dealing with only one collection, rather than two, were in an identity conservation situation. That is, they had to judge, after observing a spreading or compressing transformation, whether the quantity following the transformation was the same as the quantity preceding it; they could not simply requantify and compare the two rows.

In both situations, acquired knowledge stemmed primarily from the discovery that certain types of transformations have no effect on the numerosity of an object set, even though the transformations may affect other properties, like the spatial density or length of the set. This conclusion is independent of the number of objects in the set that was measured when the new knowledge was created. In other words, what was learned was a characterization of the quantity-preserving aspects of the transformation in question.

Q-SOAR was built to model this piece of knowledge acquisition. In order to do this, the system must be able to specify:

1. The knowledge state prior to the training (i.e., a non-conserving child);
2. The encoding of the collection(s) prior to transformation. This will include salient features such as number, length, and density, as well as other features that may ultimately be irrelevant for the task at hand;
3. The encoding of the relation between collections (for the Experimental group);
4. The encoding of the collection(s) following transformation;
5. The encoding of the physical aspects of the transformation (e.g., salient motion, how objects were moved, how many were moved, direction of movement); and
6. New knowledge acquired from repeated trials of the kind presented to both the Experimental group and the Cardinal-Once group.

The model will have two variants, and each variant will be exposed to the three kinds of stimulus presentations (corresponding to the Experimental, Cardinal-Once, and No-Cardinal groups): Q-SOAR-4 will model the four year olds, who learn from both the Experimental manipulations and the Cardinal-Once manipulations. Q-SOAR-3 will model the three year olds, who learn only from the Experimental condition.

This set of general hypotheses about the essential mechanisms involved in the child's acquisition of number conservation can be called Q Theory, to distinguish it from Q-SOAR, which conjoins Q Theory with the assumptions of a particular cognitive architecture (SOAR) to form a more complete operational theory. A full theory of conservation will ultimately contain assumptions about the nature of the environments in which development takes place. Indeed, it is the lack of justifiable assumptions that can be made about naturally occurring conservation experiences that forces us to focus entirely on training studies.

5. The Soar Architecture

This section describes the relevant aspects of the Soar architecture. More detailed accounts of it exist elsewhere (Laird, Newell, & Rosenbloom, 1987; Laird, Swedlow, Altmann, & Congdon, 1989; Newell, 1990). Besides being an operational architecture, Soar is also a theory of cognition that has been shown to explain a wide variety of psychological phenomena. No attempt can be made to describe that wider background here (cf. Lewis et al., 1990; Newell, 1990).

All tasks are formulated in Soar as search in *problem spaces*, where operators are applied to states in an attempt to attain a goal state. Problem spaces can be thought of as packages of knowledge about different tasks. The operators within a given space (and knowledge about constraints on legal states) define the problem solver's competence for a task. For example, a complete problem space for the Missionaries and Cannibals puzzle contains the necessary operators to carry out moves, knowledge about the goal state, and knowledge about legal and illegal moves. Problem solving proceeds sequentially by decisions that select *problem spaces*, *states*, and *operators*. This processing gathers knowledge from a long-term recognition memory that is implemented as a production system. This memory matches structures in working memory and retrieves knowledge that elaborates the existing state and suggests preferences for the next step to take.

If Soar cannot make a decision, an impasse occurs. In this case, Soar automatically generates a subgoal in which a new problem space can be used to find the required knowledge. A major reason that Soar exhibits the impasse and subgoal pattern is that not all of the knowledge required to carry out a task can be searched for within a single problem space. For example, should the goal arise in the Missionaries and Cannibals context to explain why the boat does not sink, there will be no knowledge in the problem space to implement that process. In response, an impasse will arise and in the resulting subgoal, Soar will select a problem space for solving such an explanatory problem, since this is a different task requiring different knowledge. Once that knowledge is found, subgoals are resolved and processing continues where it left off (Newell, 1990).

Soar has a single learning mechanism, called *chunking*, that learns new productions, or *chunks*, for resolved impasses. When similar situations are encountered, the knowledge generated by the previous subgoal processing is retrieved automatically so that the impasse is not recreated. The chunk will apply in a wider set of circumstances than the exact conditions under which it was created. This is because the chunking mechanism carries out an analysis that is a form of explanation-based learning (Mitchell, Keller, & Kedar-Cabelli, 1986; DeJong & Mooney, 1986; Rosenbloom & Laird, 1986; Mooney, this volume) to determine the critical features of the situation that led to the creation of the new knowledge. In future situations these act as cues to make the new knowledge available. The behavioral implication of chunking is that SOAR exhibits a shift from deliberate to automatic processing as the situations that it encounters become increasingly familiar. In other words, knowledge becomes compiled from search-based retrieval to recognition-based retrieval (Anderson, 1987; Rosenbloom & Newell, 1986).

6. The Acquisition of Number Conservation Knowledge

The knowledge and processes that enable Q-SOAR to acquire number conservation knowledge are implemented as a set of problem spaces that are depicted in Figure 2. The figure shows the problem spaces that are selected to carry out processing in response to given deficiencies in available knowledge; these deficiencies are stated as labels on the downward-pointing sides of the arrows. Once sufficient knowledge is returned (as depicted by the upward-pointing side of the arrows), the original processing can continue. The new knowledge becomes immediately accessible on later occasions in the form of chunks. The top panel depicts the knowledge required to interpret task instructions and to establish initial values before a transformation is applied. The lower panel depicts the knowledge involved in determining the quantitative effects of transformations.

The figure also distinguishes between task-motivated problem spaces (unshaded) and theory-motivated problem spaces (shaded). The unshaded spaces contain those operations that any task analysis of the training studies would deem to be necessary for its successful completion. These processes include the ability to understand instructions, create responses, and determine relative or absolute values for the objects used in training and testing. The shaded problem spaces contain operations that we, as theorists, assert are necessary to enable the cognitive architecture, SOAR, to achieve the behavior and learning that constitute the attainment of number conservation as shown by children in the three to four year old age range.

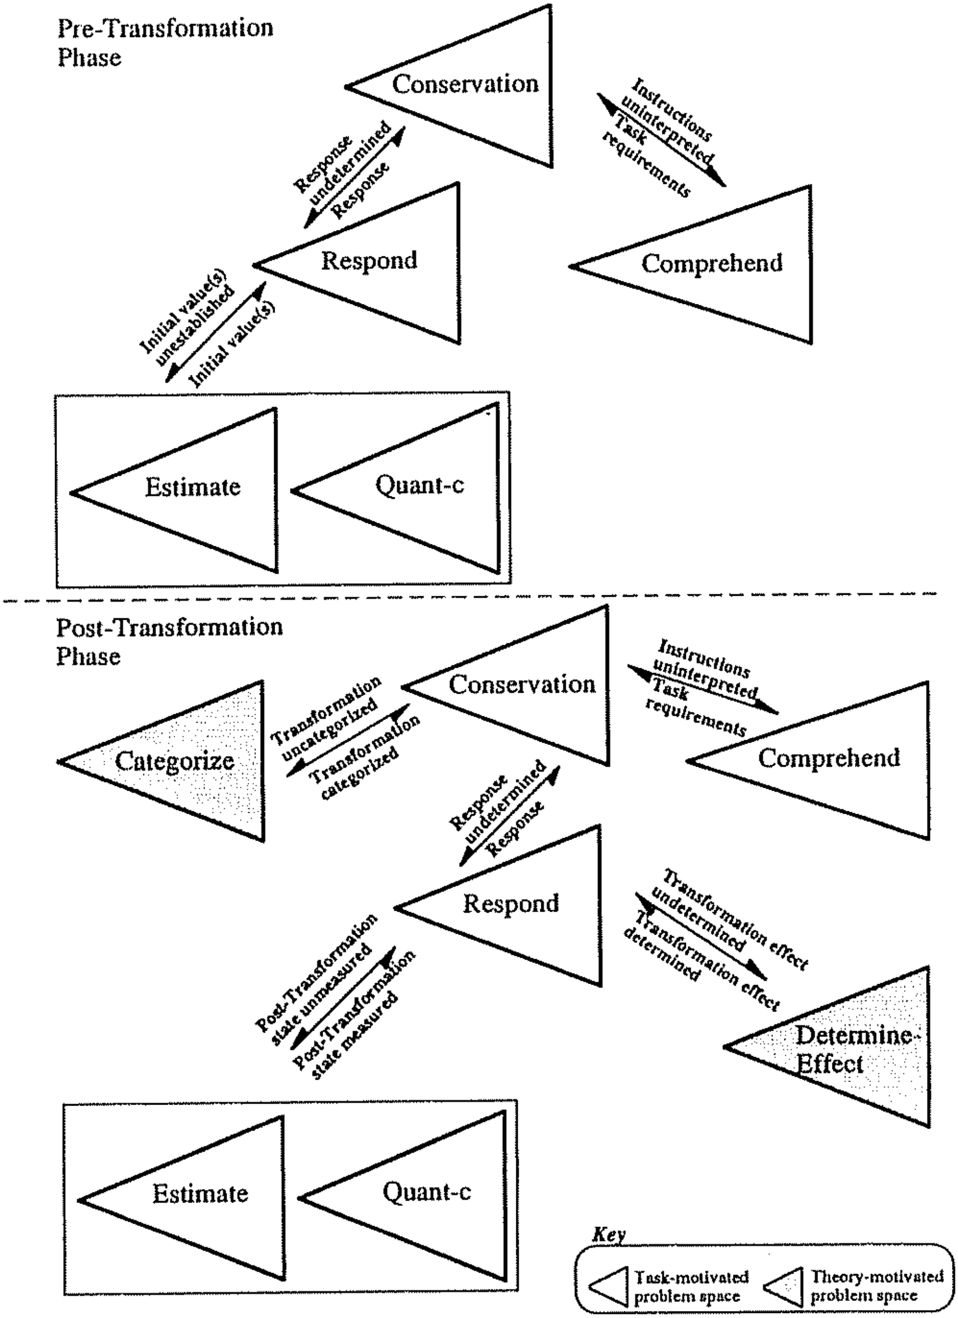

Figure 2. Problem spaces for the Q-SOAR model of conservation behavior.

Q-SOAR's design presumes that young children acquire number conservation knowledge by measurement and comparison of values to determine the effects of transformations on small collections of discrete objects. Having been shown a transformation to a set of objects, the child first categorizes the transformation and then initiates a conservation judgment about the transformation's effect. Ideally, categorization will identify the observed transformation as an instance of a larger class, with effects that are known to be associated (through chunking) with this class. If not, then pre- and post-transformation values created by measurement processes are compared to determine the effect of the transformation. The learning over this processing creates new knowledge about this kind of transformation, which will become available on future occurrences in similar contexts.[4] Now the transformation's effects can be stated without the need for any empirical processing. In other words, the necessity of the effects is recognized.

6.1 Theoretical Assumptions

The behavior exhibited by Q-SOAR is determined, in part, by five assumptions that we make about the knowledge and strategies used by three and four year old children in the context of conservation experiments. The five assumptions are:

1. *The numerical value of a row of objects will not change if the row is not transformed.* Gelman (1977) has shown in her "magic studies" that children as young as two years of age operate with the assumption that the numerical value of a set of objects will remain constant in the absence of any observable manipulation of the set.

2. *The numerical value of a row of objects is very likely to change if the row is visibly transformed by an experimenter.* Many experiments, notably those of Donaldson (1978), have shown that the act of an experimenter making explicit physical changes to a set of objects suggests to young children that some significant change will result from the action.

3. *Four year olds have the knowledge that they can verify whether assumption 2 is true or false where measurement is possible.* By mea-

4. The notion of similarity involved is the occurrence in the new situation of the same essential features used in the prior situation. There is no similarity metric involved.

suring the numerical value of the objects before and after the transformation, a determination can be made as to whether the action changed the number of the objects.

4. *Three year olds do not have the knowledge that a transformation's effects can be verified but they may be motivated to measure a transformation's effects if faced with two conflicting sources of evidence.* If one source of information suggests that the value of the objects has changed while another suggests that it has not, three year olds will attempt to determine the true effect of the transformation where measurement is possible.

5. *Three and four year olds have the capacity to store and recall pre- and post-transformation values but only four year olds do so systematically.* Four year olds have the knowledge that it is important to do this in order to determine the effect of a transformation. There is a sizable literature on young children's strategic use of their mnemonic capacity which indicates that, even if three year olds have the same capacity for remembering as older children and adults, they have little knowledge about how to exploit that capacity and thus their untutored memory performance is poor (Brown, Bransford, Ferrara, & Campione, 1983).

Q-SOAR's behavior in simulating the training study is described below in sections that correspond to the problem spaces involved. References will be made to the example procedure in Figure 1 so that the reader can keep track of both the current subproblem that Q-SOAR is attempting to solve and the arrangement of the objects concerned. The description will be presented using the Experimental procedure since it is a superset of the other two control conditions. The actual behavior that Q-SOAR produces during this procedure, and its subsequent behavior on one of the test conditions, is presented in the appendix to this chapter. (Recall that the Cardinal-Once group only experienced steps 1–2 and 6–7 and the No-Cardinal group experienced only step 1.)

6.2 Conservation Task Operations

To carry out a conservation task, five basic processes are required. These correspond to SOAR operators in the CONSERVATION problem space: COMPREHEND-INSTRUCTION, CATEGORIZE-TRANSFORMATION, DETERMINE-RESPONSE, RETURN-RESPONSE, and WAIT.

All operators process internal representations in working memory. These representations correspond to aspects of the external situation to which the model is attending. There is a focus of attention, which is determined partly by external sources, such as the experimenter asking questions or drawing attention to the experimental materials. It can also be determined by internal processing, such as attention to individual items during counting. The internal representations are in the form of *annotated models* (Lewis, Newell, & Polk, 1989), which are descriptions of attended aspects of the external situation expressed symbolically as objects with parts, properties, and relations to other objects. These models are manipulated by operators and augmented with knowledge from memory and from the external situation via perception.

Each step in the Experimental (and other) procedures is carried out in response to an instruction or request from the experimenter; we shall refer to these as *instructions*. When an instruction is perceived, its meaning must be represented as an annotated model. This model is constructed by a COMPREHEND-INSTRUCTION operator in the manner of an existing system, called NL-SOAR, which is implemented in the COMPREHEND problem space; Lewis, Newell, and Polk (1989) give a description of this process. COMPREHEND-INSTRUCTION operators produce a representation of a request called a *behavior-model object*. In step 1 of Figure 1, for example, the child is requested to count the row of circles. In this case, the COMPREHEND-INSTRUCTION operator would produce a behavior-model object representing the operation "measure", the argument "circles", and the dimension "number". Thus a behavior-model object is a child's representation to himself/herself of what behavior should be carried out to achieve the task. In general, the child is perfectly capable of behaving without such a plan, but in the case of an experiment, he/she must represent (and remember) the instruction to be carried out. These representations also play a role in mediating the speed of acquisition of conservation knowledge.

Once the instruction has been comprehended and represented, Q-SOAR must still produce a response. To do so, it selects the DETERMINE-RESPONSE operator. Q-SOAR implements steps 1 and 2 (and then steps 3 and 4) in Figure 1 with a single DETERMINE-RESPONSE operator, which is augmented with the instructions represented on the behavior-model object. If the response is not immediately available, then there will be an impasse and other problem spaces will be selected to compute the response. When the system has created a response that satisfies the instructions (such as the value "four" for steps 1 and 2), it selects the RETURN-RESPONSE operator to output the answer and then waits for the next instruction by selecting the WAIT operator.

Q-SOAR's first response to observing a transformation (step 6) is to categorize it. This is done before responding to post-transformation questions by selecting the CATEGORIZE-TRANSFORMATION operator. The categorization process is described in Section 6.6.

6.3 Response Determination

If the DETERMINE-RESPONSE operator in the CONSERVATION space cannot immediately produce a response to the counting instructions of steps 1 and 2, there will be an impasse. The RESPOND space will always be selected when the DETERMINE-RESPONSE operator impasses, since it contains the three operators that are required to create responses in the conservation task: MEASURE, COMPARE, and RECALL. In the case where steps 1 and 2 have been comprehended as requiring a measurement, the MEASUREMENT operator is selected. Depending on the number of objects and their representation (see Section 6.5.1), either the QUANT-C or ESTIMATE space is selected to carry out this measurement. That measurement is returned as the result of the MEASURE operator and, in turn, it is also returned as the result of the top-level DETERMINE-RESPONSE operator.

Once Q-SOAR has created and returned a measurement for the circles and squares, it perceives and comprehends the question "same or different number?" in step 5. The resulting DETERMINE-RESPONSE operator will be augmented with the instruction "compare", the arguments "circles" and "squares", and the dimension "number". Since this is the first time the instruction has been encountered, there will be no immediately available response and, in the resulting subgoal, the COMPARE operator will be selected in the RESPOND space. This operator tests whether the values for the measurements of the two rows match, since only a "same" or "different" response is required. In the current example, that processing will be carried out in the QUANT-C space if a value for that comparison is not immediately available.

Once the comparison has been created and returned, and the transformation has been observed and categorized (step 6 and lower panel of Figure 2), the experimenter asks, "Are there still N objects?" for each row (where N is the number of objects in the row). These instructions are treated in a similar way to steps 2 and 4, by selecting a DETERMINE-RESPONSE operator. However, the operator is now augmented with the operation "recall". If no answer to the question is immediately available, that operator will impasse and the RECALL operator will be selected in the RESPOND space. The implementation of this operator differs when it is applied to rows that have been transformed. The alternatives depend on the model variant (Q-SOAR-3 or Q-SOAR-4) and training condition (Experimental or Cardinal-Once), as discussed in Section 6.4. Responses to the "Are there still N?" question with respect to an untransformed row (e.g., step 7 in Figure 1) are dealt with in the following way. Recall that assumption 1 (in Section 6.1) stated that the numerical value of a row of objects will not change if the row is not transformed. This means that the value for the row in question is assumed to be the same before and after the transformation of the other row. Q-SOAR-4 is able to recall the pre-transformation value and return it as the answer to this question. This is not the case for Q-SOAR-3 (see assumption 5 in Section 6.1). However, even without retrieving the correct value with the RECALL operator, Q-SOAR-3 can answer the question correctly by requantifying the objects in question. This can be done by subitizing since no row has more than four objects.

6.4 Effect Determination

The way that Q-SOAR responds to the "Are there still N?" question for transformed rows is not only to produce a post-transformation value but also to determine the effect of the transformation. The DETERMINE-EFFECT problem space is selected if the RECALL operator impasses, because no effect of the transformation is immediately available. In order to learn about the effect of the transformation, the system must compare the pre- and post-transformation values for the row. Also, some role must be attributed to the transformation for the value of that comparative judgment. This can be as simple as identifying it as the action that created the post-transformation array. In other words, this process creates the knowledge to answer the implicit question, "What change did the transformation make to the number of objects?" This new knowledge states that, whenever such a transformation is applied, the relation between the pre- and post-transformation values that have just been computed will hold for the dimension in question. For example, the response to step 9 in Figure 1 is that a spreading transformation causes no change because it produces an identical value to the one that existed before it was applied, namely "four squares" in both cases.

6.4.1 EFFECT DETERMINATION IN Q-SOAR-4

Q-SOAR-4 determines the effect of transformations in the same way for both the Experimental and Cardinal-Once conditions. Although assumption 2 in Section 6.1 stated that children believe that the value of a row will change when it is transformed, assumption 3 stated that four year olds have the knowledge that this can be verified when it is possible to measure the materials before and after a transformation is applied. This is the case in both of these conditions because of the small number of discrete objects. Thus, in all cases, Q-SOAR-4 makes a pre- and post-transformation value comparison to determine the effects of observed transformations.

6.4.2 EFFECT DETERMINATION IN Q-SOAR-3

Q-SOAR-3 behaves differently in the Experimental and Cardinal-Once conditions. In the Experimental condition its behavior is like that of Q-SOAR-4, due to assumption 4. This assumption stated that three year olds do not possess knowledge about verifying a transformation's effect but may be induced to do so if they are faced with two conflicting sources of evidence. This is always the case with the Experimental condition. Before transformations, the two rows are in one-to-one correspondence so equal trials have equal-length rows and unequal trials have unequal-length rows. In other words, perceptual information and quantitative information are not in conflict. However, "transformations on unequal trials yielded rows of the same length; equal trials involved rows of different lengths" (Gelman, 1982, p. 212). In other words, after transformation, quantitative information, which was available via subitizing, and perceptual information were in conflict: unequal rows were the same length and equal rows were different lengths. This conflict leads Q-SOAR-3 to recall the pre-transformation value it measured and check it against the post-transformation value of the row.

In the Cardinal-Once condition, no such conflict exists. There is a single row of objects which, when transformed, takes on a new visual appearance. There is nothing in the visual array to suggest that the assumed change in its numerical value should be doubted. Thus Q-SOAR-3 makes no attempt to compare pre- and post-transformation values of rows. As with the untransformed row, it answers the question "Are there still N?" by requantifying. Since no comparison is made to the original value, no learning can take place as to whether the transformation has had any effect on the numerical value of the row.

6.4.3 EFFECT DETERMINATION OPERATIONS

There are two operators in the DETERMINE-EFFECT space. The RECALL operator recalls a pre-transformation value for comparison to the post-transformation value. The DETERMINE-EFFECT operator matches pre- and post-transformation values for the transformed row as described above. The process is a simple match that tests whether the values are the same or not. A requested determination of the magnitude or direction of the change will require accessing quantification knowledge. The result of this match is returned by the DETERMINE-EFFECT operator as the effect of the transformation, and the basis of its determination (e.g., that pre- and post-transformation values matched) constitutes an explanation. This will be the result of the RECALL operator in the RESPOND space and, ultimately, the DETERMINE-RESPONSE operator in the CONSERVATION space. The chunks that are built when this new knowledge is returned enable immediate retrieval of the effect of the current transformation. These chunks will fire in response to selection of the DETERMINE-RESPONSE operator, letting Q-SOAR immediately return the effect and explanation of the transformation. This demonstrates the shift in conservation performance from empirical examination of materials to direct explanation of the transformation's effects.

6.5 Quantification and Estimation

In the preceding sections we showed how the acquisition of number conservation knowledge in Q-SOAR is founded on empirical processing, whose results are then used by the effect determination process. In this section we examine the quantification and estimation abilities available to Q-SOAR that implement the measurement and comparison processes.

6.5.1 QUANTIFICATION

The primary measurement capability that is possessed by young children is quantification. This quantification subsystem, which we call *Quant-C*, for "quantification in conservation", includes only capabilities that produce cardinal values for small sets of entities. Thus it includes subitizing and counting of small sets.

In cases where values are to be determined in terms of number, the QUANT-C problem space may be selected to implement the measurement. Selection of the space depends on two factors. The first is whether the conservation property represented on the MEASURE operator suggests the use of Quant-C processes (e.g., in the case of discrete objects but not liquid). The second is the suitability of the representation for the application of operators in the QUANT-C problem space.

Even in cases where the conservation property suggests Quant-C processes for determining equivalence, the problem solver may still be unable to use it. In order to select the QUANT-C space, the representation of objects to be measured must be in the form of symbols representing discrete objects that are in one-*onto*-one mapping (hereafter *onto*) with their external referents. We assume such representations are only possible for set sizes within the range that young children can subitize: a limit of four objects. Above this limit, a much looser one-*into*-one mapping (hereafter *into*) is used.

The process of subitizing in the QUANT-C space is not controlled by an operator: it is simply that of creating an *onto* representation of up to four external referents. This approach is based on the view that there is a special code for the representation of discrete quantities which is primitive to the architecture and which differs from the formal code used to communicate about numbers with words such as "three" or symbols like "3". We call this primitive representation the *basic quantitative code* (Simon & Newell, 1990) and assume that it provides the agent with an ability to represent quantity in a primitive form.

There are six operators in the QUANT-C space: ATTEND, INITIALIZE, FETCH-NEXT, COUNT-NEXT, COMPARE, and MEMORIZE. The ATTEND operator attends to the objects specified in the behavior-model object and sets up an *onto* representation. All of the other operators are involved only if counting and not subitizing is to be carried out. The INITIALIZE operator selects a mark for identifying objects to be counted, selects an initial word from the counting string to be used, and selects an initial object to be processed. The COUNT-NEXT operator assigns a selected count word to a marked object and, where cardinal responses (Fuson, 1988) are to be returned, assigns that label to the cardinality of the set. FETCH-NEXT obtains a next item to be counted, marks it, and obtains a next count word to be assigned. The COMPARE operator can be used to test either the relative similarity or difference of values created by MEASURE operators. The MEMORIZE operator carries out a deliberate act of memorization on the final response to a pre-transformation and post-transformation instruction, so that the results are stored in long-term memory and are available for the processes that determine the effect of the transformation.

None of these operations can be directly applied to *into* representations. However, one can count large collections of objects if perceptual and motor operations can be carried out to serially map individual items onto their external referents, thereby creating transitory *onto* representations for up to four items at a time. If this cannot be done or if a decision is made against doing so, the only recourse is to use estimation operations.

6.5.2 PERCEPTUAL ESTIMATION

Perceptual estimation in the children modeled by Q-SOAR is unidimensional — the relative number of two numerous rows of objects is determined either by length or by density, but not both. Siegler (1981) showed that a child's ability to integrate more than one dimension to solve a range of problems does not develop until around eight years of age. Thus, estimation in conservation settings is inaccurate, since one dimension is often inadequate for an accurate quantitative judgment.

The ESTIMATION problem space is selected to obtain values for materials under a number of conditions. As discussed in Section 6.5.1, Q-SOAR may be requested to create a relative quantity judgment where there are too many objects to create an *onto* representation. In this case the model uses perceptual estimation, in which the primary cue as to quantity is the length of the rows. A MATCH operator carries out a type of one-to-one matching called *end matching* (Klahr & Wallace, 1976), which tests whether the end items of each row are above or below the end items of the other row. If this is not the case, the longer row is assumed to be more numerous. A MEMORIZE operator stores the result of this processing, just as in the QUANT-C space.

6.6 Categorization

As mentioned earlier, Q-SOAR categorizes observed transformations.[5] This means that it identifies critical features that are common to individual transformations, such as that all spreading actions move things further apart irrespective of the objects in question. Chunks created from processing in the DETERMINE-EFFECT problem space associate the new effect with the category of the transformation, not to the specific instance. As a result, invariance effects will be cued by any new transformation that can be identified as a member of that category. This enables novel situations to cue knowledge acquired about other members of the same category. Thus we assume that all novices, especially young children, form concepts to facilitate plausible generalizations about novel instances. Chunking models this desirable behavior, as do some other methods of explanation-based learning. As noted by Mooney (this volume), categorization need not be limited to a single dimension, but it should be sensitive to current goals. For example, when confronted with studies of number conservation like Gelman's, Q-SOAR may "CATEGORIZE" spreading, compressing, piling, and distributing together since they have no numeric effect. In contrast, if the concern is with spatial density, then compressing and piling would constitute a category with the opposite effect of spreading and distributing. The imposition of conceptual cohesiveness by goals or effects is related to *ad hoc* categorization introduced by Barsalou (1983) and discussed by Fisher and Pazzani (Chapter 6, this volume).

Q-SOAR selects the CATEGORIZE problem space when there is a transformation represented on the state and the CATEGORIZE-TRANSFORMATION operator in the CONSERVATION space cannot retrieve a type for it. The categorization process identifies in the representation of the transformation a set of features that are predictive of a certain classification. It is implemented as a recognition task. If a new instance is not immediately recognized as a member of a known class, features are progressively abstracted out of the instance description until the instance is recognized as a known class member. If no class is retrieved, then a new one is formed using the set of features in the new instance. For example, when all the features common to all spreading transfor-

5. The current version of Q-SOAR does not fully implement this process. Instead, the structures that would be created by an existing system called AL-SOAR (Miller & Laird, 1990) are fed into working memory.

mations are present, and none that are indicative of some other type of action (like compressing) are represented, the transformation will be treated the same way that other spreading transformations would be in the current context. This strategy is straightforward and sufficient to model Gelman's study, though a more general theory of conservation development will likely require more flexible principles of feature-discriminating power, like those found in Chapters 1 through 5 of this volume.

6.7 Learning Conservation Knowledge in Q-SOAR

The above subsections presented the problem spaces and operators that comprise Q-SOAR. How then do these components combine to create the number conservation knowledge that is the result of the effective training procedures? The answer is that they are called upon to contribute knowledge as Q-SOAR experiences impasses during problem solving. These impasses arise dynamically from the particular task that Q-SOAR is working on and the knowledge it brings to bear on each task at a given time. Thus, the conservation knowledge that the system has depends of what problems it has tried to solve and what knowledge it had available when it tried to solve them.

For example, if the knowledge required to respond to the question about the relative quantity of the two rows of objects in the initial array is not available, an impasse will arise. Q-SOAR will have already quantified the two values, but no comparative value will exist. The resulting series of impasses ground out in the selection of the COMPARE operator in the QUANT-C space. The successful creation of that comparative value resolves the impasse and creates a new piece of information that is available for later instances of the same problem. This kind of processing is repeated for every impasse that the system encounters. Some of these chunks simply reduce the amount of search that Q-SOAR engages in on subsequent trials (such as chunks that implement the instruction comprehension operators). Other chunks, such as those that arise from compare operators in the QUANT-C space, not only reduce search but also directly contribute to the ultimate conservation judgments that the system makes.

The result of one such impasse is the chunk (hereafter the *conservation* chunk) marked in the appendix that produces the conservation response. Due to the explanation-based nature of chunking, some generalization will occur with respect to the applicability of the chunk. Specifically, only features that existed before the impasse arose can become conditions for chunks. This is to ensure that an impasse for the same problem will not recur. However, not all of the pre-existing features will become conditions; only those that are used to compute the result in the subgoal will be selected. This means that the chunked result will be retrieved in a wider set of circumstances than the one in which it was formed. However, it does not mean that Q-SOAR exhibits conserving responses after one trial. When simulating human cognition, SOAR builds chunks only for the results created from the lowest goal in a subgoal stack. This *bottom-up* chunking causes the architecture to exhibit a gradual progression from deliberate (search-based) to automatic (recognition-based) behavior.

In the case of Q-SOAR, the conservation chunk at first only implements the response to the DETERMINE-EFFECT operator because that was the operator that led to the final impasse from which the chunk was created. Only after a series of trials is there a single subgoal caused by the top-level DETERMINE-RESPONSE operator. Then the information in the original chunk becomes available to implement that operator and thus enable a recognitional response to the effect of an observed transformation, as in the case of a conserving child.

By acquiring conservation knowledge in this way, Q-SOAR does not create any single knowledge structure that represents a "conservation concept". Instead, it builds a series of chunks that, when appropriately cued, enable the system to exhibit number conservation. In other words, rather than learning concepts that define the features of conserving transformations, Q-SOAR acquires generalized knowledge about the effects of observed transformations that is cued by other, similar transformations in similar contexts. As already described, these pieces of knowledge are acquired incrementally as problems are solved by the system with different amounts of available knowledge. In this particular modeling study, Q-SOAR was led to acquire its conservation knowledge by the use of a training regime. However, this was not a supervised concept-learning situation in which pre-classified examples of concepts are presented for the system to learn. Indeed, Q-SOAR is never presented with the concept of conservation; it is merely asked to solve a series of problems that were experimentally demonstrated to result in the acquisition of conservation knowledge. These kinds of problems can be encountered and solved without supervision and, as can be seen in the next section, we claim that Q-SOAR should be capable of learning conservation knowledge without training but at a slower speed than demonstrated here.

7. Toward a Full Theory Of Conservation

Q-SOAR successfully models the acquisition of conservation knowledge attained by subjects in Gelman's training study in an implementation within SOAR's unified cognitive theory. In this final section, we describe the behavior of Q-SOAR before and after training. We also describe what we anticipate as the necessary steps toward a full theory of conservation in other domains.

Ultimately, one can evaluate an enterprise such as the one presented here in terms of Piaget's (1964) well-known criteria for "real" conservation:

> But when I am faced with these facts [that *learning* of structures seems to obey the same laws as the *natural development* of these structures], I always have three questions which I want to have answered before I am convinced.
>
> The first question is, "Is this learning lasting? What remains two weeks or a month later?" If a structure develops spontaneously, once it has reached a state of equilibrium, it is lasting, it will continue throughout the child's entire life. When you achieve the learning by external reinforcement, is the result lasting or not and what are the conditions for it to be lasting?
>
> The second question is, "How much generalization is possible?" ... When you have brought about some learning, you can always ask whether this is an isolated piece in the midst of the child's mental life, or if it is really a dynamic structure which can lead to generalizations.
>
> Then there is the third question, "In the case of each learning experience what was the operational level of the subject before the experience and what more complex structures has this learning succeeded in achieving?"

To these three questions, we add a fourth: How can subjects (and Q-SOAR) learn so rapidly from a brief training study, when untrained subjects take several years to acquire the same knowledge?

7.1 Durability and Robustness of Learning

With respect to Piaget's first question, Q-SOAR makes a specific theoretical claim: a chunk, once learned, is always available, and will be evoked whenever the context-specific information that was included in the original chunk is recognized and encoded. For the SOAR architecture, chunking is an automatic acquisition mechanism that is applied to all processing that takes place. Thus, by undertaking the processing that is induced by the externally driven training procedure, the learning of conservation knowledge will occur.

The empirical prediction associated with this claim is not straightforward. The general pattern of results with increasingly remote post-tests is that, for a while, performance declines as a function of intervening time between training and testing, but then performance improves as one would expect with the "natural" acquisition of conservation. At present we have no principled explanation of this in terms of chunking.

7.2 Generalization

The second question refers to the specificity of learning from experience. This is a well-established empirical fact and is predicted by the chunking mechanism (Laird, Rosenbloom, & Newell, 1986). In the context of Q-SOAR, chunking predicts little generalization from learning about certain transformations of discrete objects to other transformations of different materials (e.g., the pouring of water). Indeed, this is what one usually finds from conservation training studies: very little generalization to other kinds of quantity conservation.

Transfer from small to large number tasks is achieved by the generalization inherent in SOAR's chunking mechanism. The actual objects that are measured in determining a conservation judgment are not tested when it is retrieved from memory. There are tests for the kind of transformation and the conservation property (in this case, number) and these delimit the scope of transfer. If that were not so, Q-SOAR would predict unrealistically fast learning: to transformations of quantities that are not affected in the same way as the one measured.

In addition, transfer is also limited with respect to continuous quantities such as volumes of liquid. Acquiring knowledge about continuous quantity is not addressed by Q-SOAR. Nevertheless, having acquired conservation knowledge based on small number measurement, a problem solver must come to appreciate what is common to transformations like lengthening and the pouring of liquids. This requires that these transformations be represented as actions that neither add nor remove any of the materials that they manipulate.

Finally, we suggest that the problem solving that enables the identification of the common features of different transformations and materials can best be described as a discovery process. The learner operates with a set of expectations based on current knowledge. This will at some point create a violation of the expected effects of a transformation. The learner's task then is to generate an hypothesis of what caused that violation, to devise ways of testing that hypothesis, and to integrate the results either into new hypotheses or modified knowledge. Research on scientific reasoning (Klahr & Dunbar, 1988), instructionless learning (Shrager, 1987), and analogy (Gentner, 1983) provides good explanations of the nature of such processes. Mediating factors in the effectiveness of that problem solving are the selection and combination of features that are considered for inclusion in the analysis (Bransford, Stein, Shelton, & Owings, 1981).

7.3 Operational Level and Structural Change

With respect to Piaget's third question, Q-SOAR makes explicit statements about the "complex structures" arising from the training of conservation responses. These can be seen by examinations of Q-SOAR-3 and Q-SOAR-4 before and after training.

7.3.1 Q-SOAR-3 AND Q-SOAR-4 BEFORE TRAINING

Before experiencing the three conditions of Gelman's training study, both versions of Q-SOAR are able to execute all the steps of the three experimental conditions. The only difference between the two versions is that Q-SOAR-3 does not start out with the knowledge that the effects of transformations can be verified by comparing pre- and post-transformation values. Apart from this difference, both versions have all the capabilities described in Section 6.

However, since neither variant has undergone any training and has not learned about the effects of any transformations, both versions of Q-SOAR fail all of the conservation tests that Gelman used. Neither system can accurately measure the large number of objects in the tests to yield the correct comparative answers. They must use estimation, a process that results in the assertion that a longer row contains more objects than a shorter row. Finally, both untrained variants of Q-SOAR-3 and Q-SOAR-4 are unable to determine the effects of the transformations and so cannot state an explanation. Before training, then, both are true non-conservers. We will now examine their behavior after training. Since the No-Cardinal condition is not expected to induce any change in behavior, we will discuss only the results of the other two conditions.

7.3.2 Q-SOAR-3 AFTER CARDINAL-ONCE TRIALS

Without employing its memorization capability to recall and compare pre- and post-transformation values, Q-SOAR-3 cannot learn anything about the numerical effect of observed transformations. Thus, based on assumption 2 in Section 6.1, it always assumes that the value of the row changes. Since this is never the case in the experiment, Q-SOAR-3 is always wrong and it fails the conservation tests. As can be seen in Table 1, three year olds produced few correct responses.

7.3.3 Q-SOAR-3 AFTER EXPERIMENTAL TRIALS

As explained in Section 6.4, the conflicting information in Experimental trials after a transformation induces Q-SOAR-3 to recall and compare values in the same way as Q-SOAR-4. Thus, Q-SOAR-3 can construct a correct comparison and explanation from such trials. These can then be recalled later, enabling it to pass the conservation tests. The behavior of Q-SOAR-3 after the Experimental condition produces correct responses and explanations, and thus is consistent with the pattern of results in Table 1. It seems likely that the experience of this conflict and the resulting recall and comparison of values provide the means by which three year olds acquire the effect-verification knowledge that we have assumed to be available to four year olds and which we provided for Q-SOAR-4.

7.3.4 Q-SOAR-4 AFTER BOTH TRIALS

Having produced a quantitative response before a transformation (e.g., step 4 in Figure 1), Q-SOAR-4 selects the MEMORIZE operator to store the computed values in long-term memory. Then, in the DETERMINE-EFFECT problem space, it selects the RECALL operator to enable it to compare pre- and post-transformation values to determine the effect of the transformation and create an explanation. This knowledge can then be recalled in the tests, enabling Q-SOAR-4 to pass the conservation tests after experiencing both the Experimental and Cardinal-Once procedures. This pattern of results is also consistent with that in Table 1. The higher proportion of correct responses in the Experimental group may reflect the fact that not all four year old children had acquired the effect-verification knowledge that we assumed for Q-SOAR-4. Those that had not would be expected to perform less well in the Cardinal-Once condition, just as was the case for three year olds.

7.4 Learning Speed

In Section 4 we stated that Q Theory is designed to account for the natural development of conservation, while Q-SOAR simulates only conservation learning in a single training study. Therefore, we should explain how the same processes can learn very quickly under experimental situations and yet take a few years to reach the same point during natural development. Two obvious factors are the differences in exposure and the availability of feedback. Intensive exposure to important features and informative feedback are characteristic of training studies, but neither of these is the case in unsupervised everyday activity.

However, we suggest that the greatest influence on learning speed is what we shall call the goal versus encoding interaction. A learner may activate the goal of measuring the effects of transformations. Alternatively, that learner's processing may be in the service of some other goal, such as building towers out of blocks. Even if the measurement goal has been activated, the learner may not attend to a property of the transformed materials that will reveal any number-invariance knowledge, such as the spatial density of a pile of blocks. Only if the child simultaneously has the goal of measurement and the encoding of number as the feature to be measured will he/she acquire number conservation. Well-designed training studies, such as Gelman's, foster just such optimal conditions, and in Q-SOAR these aspects are explicit in the representation of comprehended instructions. Similar directiveness appears to be provided for the child in relatively natural mother-child interactions, as set up by Saxe, Gearhart, and Guberman (1984). We know of no evidence to suggest that the goal and property combination optimal for number conservation learning would be chosen by the child any more or less often than any other, although it is evident that children often set themselves the goal of counting things. Thus, three of the four types of opportunities for learning number conservation knowledge would not produce conservation learning in Q-SOAR.

8. Conclusion

In this chapter we presented Q-SOAR, a computational model of the acquisition of conservation knowledge as reported in a single experimental training study. This is the first such account to present a set of mechanisms, constrained by a unified theory of cognition, that can be shown to acquire conservation knowledge. Furthermore, we demonstrated not only that Q-SOAR can account for the rapid learning observed in the Gelman (1982) training study but also that, without modification, it may also be able to explain the slower, more opportunistic acquisition of invariance knowledge that is characteristic of a young child's everyday unsupervised learning experiences. Much remains to be done before we can claim that Q-SOAR gives a complete account of the acquisition of conservation knowledge. There exist many other training studies (Field, 1987) whose results should also be explicable by the mechanisms of Q Theory. The transfer to conservation of continuous quantity remains to be explained, and an account of "natural" conservation development is still an important goal. Nevertheless, we believe that the work reported in this chapter represents real progress in the creation of computational theories of conceptual development.

Acknowledgements

We thank Bob Siegler for comments and access to his experimental data and Rochel Gelman for further explication of her experimental procedures. In addition, the first author wishes to thank members of the SOAR group for invaluable help, Richard Young for comments, discussions, and support, and Doug Fisher for helpful suggestions on earlier drafts. This work was supported by Contract N00014-86-K-0678 from the Computer Science Division of the Office of Naval Research.

References

Anderson, J. R. (1987). Skill acquisition: Compiling weak method problem solutions. *Psychological Review*, *94*, 194–210.

Barsalou, L. W. (1983). Ad hoc categories. *Memory and Cognition*, *11*, 211–227.

Bransford, J. D., Stein, B. S., Shelton, T. S., & Owings, R. A. (1981). Cognition and adaptation: The importance of learning to learn. In J. Harvey (Ed.), *Cognition, social behavior, and the environment*. Hillsdale, NJ: Lawrence Erlbaum.

Brown, A. L., Bransford, J. D., Ferrara, R. A., & Campione, J. C. (1983). Learning, remembering and understanding. In J. H. Flavell & E. M. Markman (Eds.), *Handbook of child psychology: Cognitive development* (Vol. 3). New York: John Wiley.

Campbell, R. L., Cooper, R. G., & Blevins-Knabe, B. (1988). *The development of subitizing in preschool and early elementary school children* (Tech. Rep. No. 63254). Yorktown Heights, NY: IBM, T. J. Watson Research Center.

Chi, M. T. H., & Klahr, D. (1975). Span and rate of apprehension in children and adults. *Journal of Experimental Child Psychology*, *19*, 434–439.

DeJong, G., & Mooney, R. (1986). Explanation-based learning: An alternative view. *Machine Learning*, *1*, 145–176.

Donaldson, M. (1978). *Children's minds*. Glasgow, Scotland: Fontana.

Field, D. (1987). A review of preschool conservation training: An analysis of an analysis. *Developmental Review*, 7, 210–251.

Fuson, K. C. (1988). *Children's counting and concepts of number*. New York: Springer Verlag.

Gelman, R. (1977). How young children reason about small numbers. In N. J. Castellan, D. B. Pisoni, & G. R. Potts (Eds.), *Cognitive theory* (Vol. 2). Hillsdale, NJ: Lawrence Erlbaum.

Gelman, R. (1982). Accessing one-to-one correspondence: Still another paper about conservation. *British Journal of Psychology*, *73*, 209–220.

Gelman, R., & Baillargeon, R. (1983). A review of some Piagetian concepts. In J. H. Flavell & E. M. Markman (Eds.), *Handbook of child psychology: Cognitive development* (Vol. 3). New York: John Wiley.

Gentner, D. (1983). Structure mapping: A theoretical framework for analogy. *Cognitive Science*, *7*, 155–170.

Halford, G. S. (1982). *The development of thought*. Hillsdale, NJ: Lawrence Erlbaum.

Halford, G. S., & Boyle, F. M. (1985). Do young children understand conservation of number? *Child Development*, *56*, 165–176.

Halford, G. S., Smith, S., Mayberry, M., Stewart, J., & Dickson, J. C. (1991). *A computer simulation model of acquisition of transitive inference*. Paper presented at the Biennial Meeting of the Society for Research in Child Development, Seattle, WA.

Klahr, D. (1973). Quantification processes. In W. Chase (Ed.), *Visual information processing*. New York: Academic Press.

Klahr, D. (1982). Nonmonotone assessment of monotone development: An information processing analysis. In S. Strauss & R. Stavy (Eds.), *U-shaped behavioral growth*. New York: Academic Press.

Klahr, D. (1984). Transition processes in quantitative development. In R. J. Sternberg (Ed.), *Mechanisms of cognitive development*. New York: W. H. Freeman.

Klahr, D., & Dunbar, K. (1988). Dual space search during scientific reasoning. *Cognitive Science*, *12*, 1–48.

Klahr, D., & Wallace, J. G. (1970). An information processing analysis of some Piagetian experimental tasks. *Cognitive Psychology*, *1*, 358–387.

Klahr, D., & Wallace, J. G. (1973). The role of quantification operators in the development of the conservation of quantity. *Cognitive Psychology*, *4*, 301–327.

Klahr, D., & Wallace, J. G. (1976). *Cognitive development: An information processing view*. Hillsdale, NJ: Lawrence Erlbaum.

Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). SOAR: An architecture for general intelligence. *Artificial Intelligence*, *33*, 1–64.

Laird, J. E., Rosenbloom, P. S., & Newell, A. (1986). Chunking in SOAR: The anatomy of a general learning mechanism. *Machine learning*, *1*, 11–46.

Laird, J. E., Swedlow, K. R., Altmann, E. M., & Congdon, C. B. (1989). *SOAR 5 User's Manual* (Tech. Rep.). Ann Arbor: University of Michigan, Department of Electrical Engineering and Computer Science.

Lewis, R. L., Huffman, S. B., John, B. E., Laird, J. E., Lehman, J. F., Newell, A., Rosenbloom, P. S., Simon, T., & Tessler, S. G. (1990). SOAR as a unified theory of cognition. *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society* (pp. 1035–1042). Cambridge, MA: Lawrence Erlbaum.

Lewis, R. L., Newell, A., & Polk, T. A. (1989). Toward a SOAR theory of taking instructions for immediate reasoning tasks. *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society* (pp. 514–521). Ann Arbor, MI: Lawrence Erlbaum.

McClelland, J. L. (1991). *Connectionist models of developmental change*. Paper presented at Biennial Meeting of the Society for Research in Child Development, Seattle, WA.

Miller, C. S., & Laird, J. E. (1990). *A simple, symbolic model for associative learning and retrieval*. Unpublished manuscript, Artificial Intelligence Laboratory, University of Michigan, Ann Arbor.

Mitchell, T.M., Keller, R.M., & Kedar-Cabelli, S.T. (1986). Explanation-based generalization: A unifying view. *Machine Learning*, *1*, 47–80.

Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.

Piaget, J. (1952). *The child's conception of number*. New York: W. W. Norton.

Piaget, J. (1964). Development and learning. In R. E. Ripple & V. N. Rockastle (Eds.), *Piaget rediscovered*. Ithaca, NY: Cornell University Press.

Piaget, J. (1968). Quantification, conservatism and nativism. *Science*, *162*, 976–979.

Piaget, J. (1970). *Structuralism*. New York: Basic Books.

Rosenbloom, P. S., & Laird, J. E. (1986). Mapping explanation-based generalization into SOAR. *Proceedings of the Fifth National Conference on Artificial Intelligence* (pp. 561–567). Philadelphia, PA: Morgan Kaufmann.

Saxe, G. B., Gearhart, M., & Guberman, S. R. (1984). The social organization of early number development. In B. Rogoff & J. V. Wertsch (Eds.), *Children's learning in the zone of proximal development*. San Francisco, CA: Jossey-Bass.

Shrager, J. (1987). Theory change via view application. *Machine Learning*, *2*, 1–30.

Siegler, R. S. (1981). Developmental sequences within and between concepts. *Monographs of the Society for Research in Child Development*, *46*, 1–74.

Siegler, R. S. (1991). *Variation and selection as cognitive transition mechanisms*. Paper presented at the Biennial Meeting of the Society for Research in Child Development, Seattle, WA.

Siegler, R. S. (1989). Mechanisms of cognitive development. *Annual Review of Psychology*, *40*, 353–379.

Simon, T., & Newell, A. (1990). *Subitizing is faster than you think — If you don't ask 'How Many?'* Unpublished manuscript, Department of Psychology, Carnegie Mellon University, Pittsburgh, PA.

Svenson, O., & Sjoberg, K. (1983). Speeds of subitizing and counting process in different age groups. *Journal of Genetic Psychology*, *142*, 203–211.

Wallace, J. G., Klahr, D., & Bluff, K. (1987). A self-modifying production system model of cognitive development. In D. Klahr, P. Langley, & R. Neches (Eds.), *Production system models of learning and development*. Cambridge, MA: MIT Press.

Appendix: Sample Q-SOAR Run

Here we illustrate how the problem spaces described in Section 6 generate behavior when Q-SOAR is presented with a task. The following trace is an abstracted version of the steps presented in Figure 1, which show Q-SOAR in operation for the first time. The second trace shows the model's successful performance on a conservation test.

The traces retain only the critical information, showing the problem spaces (denoted by P) and operators (denoted by O) that are selected in response to the impasses that arise. An impasse is shown by processing in a subgoal (G) being indented under the operator that produced the impasse. When an impasse is resolved, processing continues at the highest level at which an operator can be selected. The operators are augmented with the instruction that led to their initiation or by the objects on which they are focused.

External arrays and instructions are depicted to the right of the trace in lower case and Q-SOAR's output is given in the center in upper case. The trace is marked with '**' at the points in the run where the key conservation chunk is acquired and where it is evoked. Chunks are created continually throughout the run (one or more when returning from each impasse), but these are not shown.

ABSTRACTED RUN OF Q-SOAR DURING ITS FIRST OPERATION

```
                                                        () () () ()
                                                        [] [] [] []
P: (CONSERVATION)
                                        How many circles?
O: (COMPREHEND-INSTRUCTIONS)
O: (DETERMINE-RESPONSE)
==>G: (OPERATOR NO-CHANGE)
   P: (RESPOND)
   O: (MEASURE)
   ==>G: (OPERATOR NO-CHANGE)
      P: (QUANT-C)
      O: ((CIRCLE) ATTEND)
      O: (INITIALIZE)
      O: ((CIRCLE) COUNT-NEXT) Counting item: ONE
      O: ((CIRCLE) FETCH-NEXT)
      O: ((CIRCLE) COUNT-NEXT) Counting item: TWO
      O: ((CIRCLE) FETCH-NEXT)
      O: ((CIRCLE) COUNT-NEXT) Counting item: THREE
      O: ((CIRCLE) FETCH-NEXT)
      O: ((CIRCLE) COUNT-NEXT) Counting item: FOUR
      O: (MEMORIZE)
O: (RETURN-RESPONSE)                    Answer FOUR
                                        How many squares?
O: (COMPREHEND-INSTRUCTIONS)
O: (DETERMINE-RESPONSE)
==>G: (OPERATOR NO-CHANGE)
   P: (RESPOND)
   O: (MEASURE)
   ==>G: (OPERATOR NO-CHANGE)
      P: (QUANT-C)
      O: ((SQUARE) ATTEND)
      O: (INITIALIZE)
      O: ((SQUARE) COUNT-NEXT) Counting item: ONE
      O: ((SQUARE) FETCH-NEXT)
      O: ((SQUARE) COUNT-NEXT) Counting item: TWO
      O: ((SQUARE) FETCH-NEXT)
      O: ((SQUARE) COUNT-NEXT) Counting item: THREE
      O: ((SQUARE) FETCH-NEXT)
```

```
      O: ((SQUARE) COUNT-NEXT) Counting item: FOUR
      O: (MEMORIZE)
O: (RETURN-RESPONSE)           Answer FOUR
                               Same or different number?
O: (COMPREHEND-INSTRUCTIONS)
O: (DETERMINE-RESPONSE)
==>G: (OPERATOR NO-CHANGE)
   P: (RESPOND)
   O: (COMPARE)
   ==>G: (OPERATOR NO-CHANGE)
      P: (QUANT-C)
      O: (COMPARE)
O: (RETURN-RESPONSE)           Answer SAME
                                             () () () ()
                                             []   []    []    []
                               Still four circles?
O: (CATEGORIZE-TRANSFORMATION)
O: (COMPREHEND-INSTRUCTIONS)
O: (DETERMINE-RESPONSE)
==>G: (OPERATOR NO-CHANGE)
   P: (RESPOND)
   O: (RECALL)
O: (RETURN-RESPONSE)           Answer FOUR
                               Still four squares?
O: (COMPREHEND-INSTRUCTIONS)
O: (DETERMINE-RESPONSE)
==>G: (OPERATOR NO-CHANGE)
   P: (RESPOND)
   O: (RECALL)
   ==>G: (OPERATOR NO-CHANGE)
      P: (DET-EFFECT)
      O: (RECALL)
      O: (DETERMINE-EFFECT)**
O: (RETURN-RESPONSE)           Answer FOUR
                               Same or different number?
O: (COMPREHEND-INSTRUCTIONS)
O: (DETERMINE-RESPONSE)
==>G: (OPERATOR NO-CHANGE)
   P: (RESPOND)
   O: (COMPARE)
O: (RETURN-RESPONSE)           Answer SAME
      VALUES MATCH BEFORE AND AFTER THIS TRANSFORMATION.

End -- Explicit Halt
```