7

A Computational Theory of Children's Learning About Number Conservation

Tony Simon
Georgia Institute of Technology

David Klahr
Carnegie Mellon University

Among the child's most important conceptual acquisitions are the information processes that underlie quantification: the knowledge and skills related to the encoding and manipulation of quantitative information. Like the acquisition of another major conceptual system, language, the earliest developments occur long before children enter a classroom, whereas most of the formal aspects of numbers are learned in school. In this chapter, we suggest that the earliest bases for understanding numbers are founded upon innate or early developing perceptual capabilities of the young child. More specifically, we focus on the development of children's understanding of number conservation and the construction of the knowledge required by that transition.

Since Piaget, conservation has been thought to be the first major conceptual advance in children's numerical development. It is not the kind of knowledge that results from what young children learn in school, and yet it appears to be a foundation that must be securely in place before many other concepts of number and arithmetic can be acquired. Although conservation of number is one of the most heavily researched phenomena in cognitive development, a satisfactory theoretical account of its acquisition remains to be formulated. The first computational models of different *states* of conservation knowledge were proposed nearly 20 years ago (Klahr & Wallace, 1976). More recently, Klahr, (1984) sketched a flowchart type model of the *transition processes* but he did not implement it as a computational model. In this chapter, we extend those earlier ideas and integrate them with Soar (Newell, 1990)—a unified theory of cognition—to present a *computation model of conservation acquisition*. The model was first proposed by Simon, Newell and Klahr (1991), and we elaborate that account in this chapter.

Our computational model demonstrates that conservation learning derives from the ability to make accurate measurements and to use them to evaluate the numerical effects of transformations on collections of objects. Our account is the logical inverse of the one presented by Piaget. The difference between our position and the Piagetian one concerns the developmental roles of two central conceptual attainments in the development of quantification abilities. These are *conservation knowledge* (understanding the behavior of quantities under transformation) and *measurement skills* (creating quantitative values for bodies of material).

The Piagetian view (Piaget, Inhelder and Szeminska, 1960) is that conservation is a logical prerequisite to the ability to measure. Piaget reasoned that, without an understanding of the essential nature of quantity, measurements in terms of those quantities would mean nothing and would be of no practical use. The opposing view is that measurement is the necessary precursor of conservation (Klahr & Wallace, 1976; Miller, 1984). Measurement is the empirical tool used to gather information about whether or not some dimension of a transformed entity has remained quantitatively invariant. Miller stated that "practical measurement procedures appear not to be late-developing concomitants of a more general understanding of quantity. Instead, the measurement procedures of children embody their most sophisticated understanding of the domain in question. The limitations of these procedures constitute significant limits on children's understanding of quantity" (p. 221).

Such measurement is not always possible. The limitations Miller spoke of determine what children can learn about quantity. They are responsible for the pattern in the development of conservation. Number, or discrete quantity, conservation is acquired first. Also, preconservers can reason successfully about transformations of small discrete quantities but not about large ones (Cowan, 1979; Fuson, Secada, & Hall 1983; Siegler, 1981). Conservation of continuous quantities such as length, area, and volume is acquired a year or two later (Siegler, 1981).

One type of limitation is on *processes*, that is, to what kinds of things measurement procedures can be applied. As Piaget et al. (1960) stated, "to measure is to take out of a whole, one . . . unit, and to transpose this unit on the remainder of the whole." Thus, any material to be measured must afford the measurer some unit that can be used in that process. This characteristic is not present in continuous quantities. Beakers of water or pieces of string do not exhibit any evident subunits. Only the employment of special tools such as rulers or measuring cylinders (and the knowledge of how to use them) can create subunits that can be used for quantification. On the other hand, discrete quantities are defined by collections of individual subunits of the quantity as a whole. No special tools are needed because quantification abilities are present to some extent in even the youngest children. Young children appear to be particularly sensitive to the fact that unitary objects, and not subparts of those objects, have a special status. In learning language, they choose that level for the assignment of novel word labels (Markman, 1990), and in quantification of collections, they choose that level for the assignment of units (Shipley & Shepperson, 1990) Thus, discrete quantities are clearly easier to identify, and thus, to measure.

A second type of limitation is in the abilities of children who attempt to use measurement procedures. The children that need to carry out measurements to determine quantitative invariance are those below the age of 5. However, their quantification skills are not well developed. They are efficient at *subitizing*: a fast, accurate perceptual quantification mechanism (Chi & Klahr, 1975; Svenson & Sjoberg, 1983). Subitizing, however, has a limit of about four objects (Atkinson, Campbell & Francis, 1976; Simon, Cabrera, & Kliegl, 1993). Young children's counting is only reliable for collections of about the same size (Fuson, 1988).

Therefore, the measurement-before-conservation view predicts the learning events that enable the acquisition of quantitative invariance knowledge. It follows that, if measurement is needed to be able to reason about quantity, learning can occur only when the effects of transformations of small collections of objects are evaluated. These quantities have to be discrete because young children are not capable of creating consistent subunits from continuous quantities. Gelman (1977) showed that 1-year-olds can reason about some transformations when the number of objects involved is very small. The discrete quantity requirement was supported by Piaget et al.'s (1960) and Miller's (1984) findings that, given the task of dividing up an object such as a cookie into equal parts, young children created many arbitrarily sized subunits. These are unsuitable for quantification because counting them fails to produce accurate absolute measures for single entities, or to produce relative measures of multiple entities.

Miller (1989) further demonstrated the interaction between the use of measurement procedures and the acquisition of quantitative knowledge. He tested 3- to 10-year-olds on a modified equivalence-conservation task, where the effect of transforming one of a pair of quantities must be established. A variety of transformations were applied to different materials to test number, length, and area conservation. Miller demonstrated that the effects of transformations are easy to determine when measurement procedures provide good cues to the actual quantity, and vice versa. For example, spreading out a row of objects has no effect if number is the conservation dimension in question, but it does affect length conservation. Therefore, the appropriate measurement tool for this transformation would be enumeration in a number task. However, it would produce no information useful for evaluating length conservation. Thus, Miller predicted that performance would be best where transformations were relevant to the domain: The effects of spreading a row of objects would be easily evaluated in the number task, but the effects of changing their size would not. The results were as predicted, showing that the acquisition of quantitative knowledge depends on the selection and application of appropriate measurement procedures.

Our theory follows that of Klahr (1984) in stating that it is measurement of collections of discrete objects that provides information upon which knowledge about quantitative invariance is built. Conservation knowledge is acquired in situations where invariance can be empirically verified. In other words, learning events occur when the materials allow children to use their measurement capabilities to obtain a numerical measurement for a collection of objects before and after it has been transformed. The two measurements can then be compared and the result attributed to the transformation as its effect. If the results are identical, the quantity is unchanged and the transformation is deemed to have a nonquantitative effect for the dimension in question—it *conserves number*. If some difference is detected, the transformation is found to be nonconserving. Such differences can be simply detected by means of discriminations based on subitizing. With sufficient domain knowledge, the direction and magnitude of the change can also be determined. Thus, we conclude that the initial learning experiences for invariance knowledge are based on measurements of small collections of discrete objects within the subitizing range. Such a view is consistent with Starkey's (1992) conclusion that simple numerical abstraction competence supports numerical reasoning before the emergence of the mature counting skill.

In the following sections of this chapter, we present the various components of our theory of conservation learning and the accompanying computational model. We begin in by examining the phenomenon of conservation in order to establish the learning task that our model will account for. Then, we discuss the particular training study that we used as the vehicle for our demonstration of conservation learning, and present our theory of number conservation learning in more detail. Next, we briefly overview Soar, the computational medium within which our model is constructed, and then present a detailed account of Q-Soar, our computational model of number conservation learning. Finally, we extend the account of conservation beyond the behaviors directly modeled by Q-Soar and draw conclusions.

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) defined 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.

As we stated, 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 by being naturally acquired by the child outside the laboratory, are the linkages between transformational attributes and their dimensional effects as measured in a variety of contexts.

The centrality of conservation concepts to most theories of cognitive development produced a vast database of empirical results. Nevertheless, a computational model that can account for the regularities has yet to be fully specified. There are 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 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 criteria for computational models of developmental phenomena:

1. Such models should be based on a principled cognitive architecture, rather than a set of arbitrary and ad hoc mechanisms. For Q-Soar, the architecture is Soar, to be described in a later section.
2. Computational models should be constrained by the general regularities in the relevant empirical literature. There are a number of such regularities, that is, findings that are consistently reported and for which there is little or no disconfirming evidence. The critical regularities for the construction of Q-Soar are later discussed in detail.
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, 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) consisted 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 fully developed model (Wallace, Klahr, & Bluff, 1987) went 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 broad set of theoretical assumptions that are consistent with a range of specific implementations. Examples include the adaptive production system used by Halford et al. (chap 3., this volume) to model the acquisition of transitive inference, the connectionist model used by McClelland (1989) to model the acquisition of balance scale rules and the concept-formation system used by Jones and

VanLehn (in press) to model the strategy changes in young children's arithmetic. Computational models of this type, although they suggest interesting learning mechanisms, tend to be relatively unconstrained by either any particular empirical results on children's knowledge acquisition or by illustrations that their central constructs are critical in accounting for other aspects of human cognition.

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, to be described next. The further issue of explaining general developmental regularities is addressed in the final section of the chapter.

A TRAINING STUDY

Simon and Halford (chap. 1, this volume) discuss, simulating the activities of children engaged in a training study is one way to ensure that the model and the child engage in very similar activities while acquiring the knowledge involved in a transition of interest. In constructing our 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 impossible to get agreement on even a prototypical training study, let alone a set of defining properties. For example, Field's (1987) review organized a collection of 25 recent conservation training studies with preschoolers along nine dimensions and three theoretical orientations.[1] 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.

Gelman's Training Procedure

As noted, we chose a training study reported by Gelman (1982) in which 3- and 4-year-olds were trained in a brief session using small collections of

[1]The procedural dimensions were design, pretest, training, materials, reinforcements, verbal rule instruction, posttest, justifications, and delayed posttest. The theoretical orientations were specific experience, cognitive readiness, and perceptual readiness. Space does not permit an elaboration of these *guiding models*, as Field (1987) called them, but training studies vary widely along both procedural and theoretical dimensions.

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 Fig. 7.1:

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 whether the pair of rows had the same number or a different number of items, and to explain his or her judgment.

All children answered the questions correctly (except for one 3-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. (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."

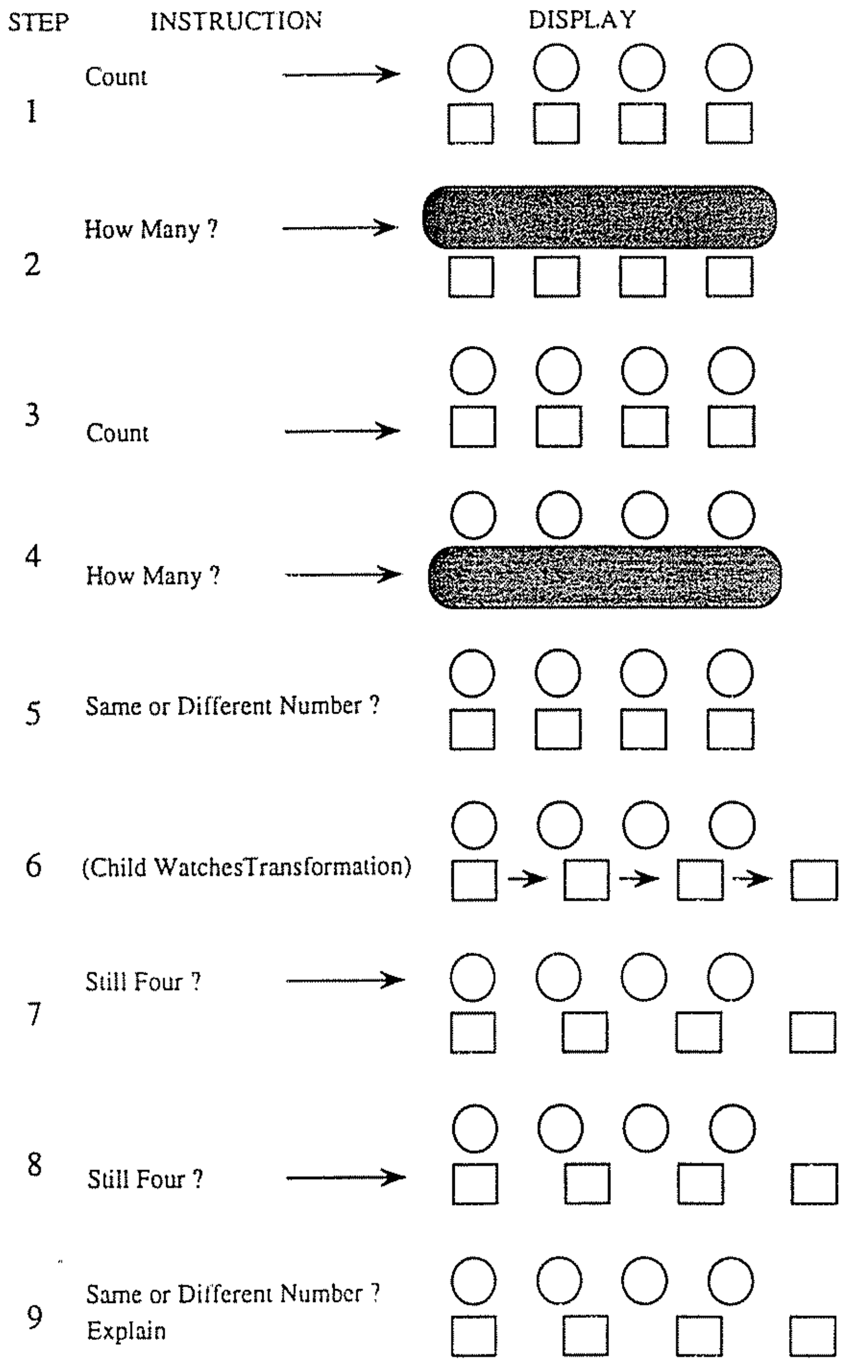

FIG. 7.1. Graphical representation of the experimental procedure

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 4 and 5 or 5 and 5 items, and large sets included either 8 and 10 or 10 and 10 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 reliably (Fuson, 1988), it is likely that the *pre-transformation* equivalence (or nonequivalence) of both large and small arrays was established by one-to-one correspondence.

Results

For both the large and small sets, there was almost no difference in the equal and unequal set sizes, so those results are collapsed in the following discussion. Table 7.1 shows the overall proportion of correct judgments on conservation tasks. The effect of condition is striking: Overall, the exper-

TABLE 7.1.
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	*Experimental (n = 21)*	*Cardinal-Once (n = 24)*	*No-Cardinal (n = 16)*
3s on small set size	71	11	13
4s on small set size	75	58	16
3s on large set size	72	8	0
4s on large set size	65	34	13
Overall 3s	71	9	6
Overall 4s	70	46	15

imental 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).[2] 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 correspondence 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 4-year-olds, though few 3-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 7.2 in terms of (our interpretation of) whether the explanation makes reference to the transformation or to one-to-one correspondence. Gelman gave examples of two categories: *irrelevant transformation* (*They just moved*) and *addition/subtraction* (*You need another one to make them the same*) explanations. Initial equality or inequality of number explanations presumably stated that the original value still held, whereas the content of one-to-one correspondence explanations is obvious. The majority of the children's explanations referred to the transformation and, as can be seen in Table 7.2, there were 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, whereas 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 4-year-olds in the experimental group benefited more than did their agemates 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, received repeated exposure to

[2]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 7.2
Reference to Transformation or One-To-One Correspondence of Gelman's Explanation Categories

	Transformation	*One-to-One*
Irrelevant transformation	X	
Addition/subtraction	X	
Initial equality/inequality	X	
One-to-one correspondence		X

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.

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 posttransformation information is visible in the form of a transformed and an untransformed row after the transformation has taken place. However, in the cardinal-once trials, memory of the pretransformation information is always required to compute the transformation's effect. All of the 3-year-olds and some of the 4-year-olds apparently needed this additional information (and reduction of processing) that was provided to the experimental group.

A THEORY OF NUMBER CONSERVATION KNOWLEDGE

We 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 or not 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 nonconserving child).
2. The encoding of the collection(s) prior to transformation. This includes salient features such as number, length, and density, as well as other features that may 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).
6. New knowledge acquired from repeated trials of the kind presented to both groups.

The model will have two variants, and each 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 4-year-olds, who learn from both the experimental manipulations and the cardinal-once manipulations. Q-Soar–3 will model the 3-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 (see Simon & Halford, chap. 1, this volume). 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.

THE SOAR ARCHITECTURE

This section describes the relevant aspects of the Soar architecture. More detailed accounts 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 explains a variety of psychological phenomena. We make no attempt 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 and 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, because 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 automatically retrieved 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 (DeJong & Mooney, 1986; Mitchell, Keller, & Kedar-Cabelli, 1986; Mooney, 1991; Rosenbloom & Laird, 1986) 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 it encounters become increasingly familiar. In other words, knowledge becomes compiled from search-based retrieval to recognition-based retrieval (Anderson, 1987; Rosenbloom & Newell, 1986).

Q-SOAR'S 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 Fig. 7.2. 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 3- to 4-year-old age range.

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 posttransformation 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, thatwill become available on future occurrences in similar contexts.[3] 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.

[3] 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.

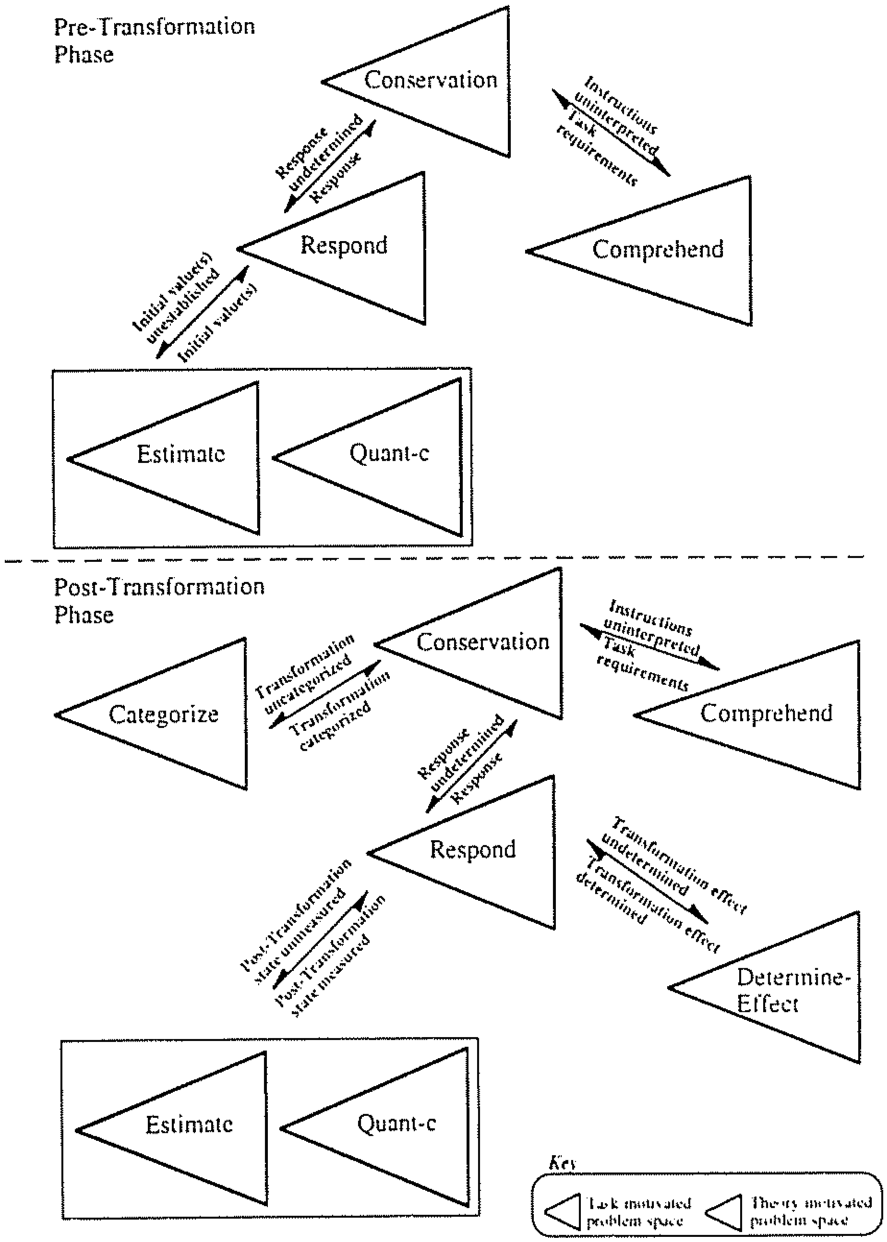

FIG 7.2. Problem spaces for the Q-Soar model of conservation behavior.

Regularities and Theoretical Assumptions

The behavior exhibited by Q-Soar is constrained by a set of regularities in the literature on early numerical abstraction and reasoning. The processes that Q-Soar employs to generate that behavior are determined, in part, by five assumptions we make about the knowledge and strategies used by 3- and 4-year-old children in the context of conservation experiments. These regularities and theoretical assumptions are discussed in detail and are integrated in Table 7.3. The top half of the table relates the regularities (R1 to 6) to 3- and 4-year-olds' abilities and tendencies to quantify, and to their abilities to correctly answer conservation questions about sets of different sizes. The lower half of the table maps the competencies provided for our Q-Soar models of 3- and 4-year-olds onto those abilities and it shows the role played by the theoretical assumptions we made.

Six critical regularities define the requirements for the behavior of Q-Soar:

R1. Young children in the 3- to 4-year age range can consistently obtain accurate specific quantitative values for small sets of objects (up to three or four; cf. Fuson, 1988; Starkey, 1992).

R2. Young children can execute a counting procedure for larger numbers when requested and can accurately monitor the counts of others. However, children below the age of 4 do not spontaneously use counting as an effective method for solving a range of quantitative problems, including comparing the relative numerosity of two sets of objects (Briars & Siegler, 1984; Fuson, 1988; Sophian, 1987; Starkey, 1992; Wynn, 1990).

TABLE 7.3
Regularities and Assumptions Underlying the Construction of Q-Soar

	Accurate Quantification?		*Count to Quantify or Compare Set Sizes?*		*Solve Conservation of Discontinuous Quantity?*	
Age	$n < 4$	$4 < n \leq 10$	*Spontaneously*	*On Demand*	$n < 4$	$4 < n \leq 10$
3	yes *R1 R3*	no *R2*	no *R2*	yes *R2*	yes *R4 R6*	no *R5*
4	yes *R1 R3*	no *R2*	yes *R2*	yes *R2*	yes *R4 R6*	no *R5*
			Verify Outcome Quantitatively?			
			Spontaneously	*Via Conflict*		
Q-Soar–3	yes	no	no *A4*	yes *A4*	yes *A1–A5*	no
Q-Soar–4	yes	no	yes *A3*	yes *A3 A4*	yes *A1-A5*	no

R3. children from an early age are able to accurately and rapidly enumerate three or four items by the process of subitizing (Chi & Klahr, 1975; Svenson & Sjoberg, 1983; van Loosbroek & Smitsman, 1990).

R4. 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).

R5. 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).

R6. 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).

There are five assumptions underlying the construction of Q-Soar. The first two concern the typical expectations held by young children about the behavior of sets of objects under transformation in the context of conservation experiments. The remaining three are our modeling assumptions about how these first two are implemented in Q-Soar in terms of knowledge and processes:

A1. Young children assume that the numerical value of a row of objects will not change if the row is not transformed. Gelman (1977) showed in her "magic studies" that children as young as 2 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.

A2. Young children assume that the numerical value of a row of objects is likely to change if the row is visibly transformed by an experimenter. Many experiments, notably those of Donaldson (1978), showed 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.

A3. Four-year-olds know that they can use measurement to verify whether Assumption 2 is true or false. By measuring the numerical value of the objects before and after the transformation, a determination can be made as to whether or not the action changed the number of the objects. Sophian (1987) showed that by this age, children are beginning to spontaneously use counting to solve quantitative problems. However, all 4-year-olds can verify outcomes using subitizing where set sizes are small enough.

A4. Three-year-olds do not typically attempt to verify the effects of a transformation, but they may be motivated to create a measurement if faced with two conflicting sources of evidence. If one source of information suggests that the value of the objects has changed whereas another suggests that it has not, 3-year-olds will attempt to determine the true effect of the transformation where measurement is possible. Sophian (1987) showed little spontaneous use of counting in relevant tasks for children of this age.

A5. Three- and 4-year-olds have the capacity to store and recall pre- and posttransformation values, but only 4-year-olds systematically do so. 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 that indicates that, even if 3-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 in sections that correspond to the problem spaces involved. References are made to the example procedure in Fig. 7.1, 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 is presented using the experimental procedure because 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. (Recall that the cardinal-once group only experienced Steps 1–2 and 6–7 and the no-cardinal group experienced only Step 1.)

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, 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)—descriptions of attended aspects of the external situation symbolically expressed 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 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 et al. (1989) described this process. COMPREHEND-INSTRUCTION operators produce a representation of a request called a *behavior-model object*. In Step 1 of Fig. 7.1, 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 or 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 or 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 Fig. 7.1 with a single DETERMINE-RESPONSE operator, that 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 posttransformation questions by selecting the CATEGORIZE-TRANSFORMATION operator. The categorization process is described in a later section.

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, because 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, 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*. Because 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 or not the values for the measurements of the two rows match, because 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 Fig. 7.2), 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 the next section. Responses to the "Are there still *N*?" question with respect to an untransformed row (e.g., Step 7 in Fig. 7.1) are dealt with in the following way. Recall that Assumption 1 stated that the numerical value of a row of objects will not change if the row is not transformed. 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 pretransformation value and return it as the answer to this question. This is not the case for Q-Soar–3 (see Assumption A5). However, even without retrieving the correct value with the RECALL operator, Q-Soar–3 can correctly answer the question by requantifying the objects. This can be done by because because no row has more than four objects.

Effect Determination

The way that Q-Soar responds to the "Are there still *N*?" question for transformed rows is not only to produce a posttransformation value, but 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 posttransformation 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 posttransformation 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 posttransformation values that have just been computed will hold for the dimension in question. For example, the response to Step 9 in Fig. 7.1 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.

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 A2 stated that children believe that the value of a row will change when it is transformed, Assumption A3 stated that 4-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 posttransformation value comparison to determine the effects of observed transformations.

Effect Dertermination 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 A4. This assumption stated that 3-year-olds do not readily engage in 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). 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 pretransformation value it measured and check it against the posttransformation value of the row.

In the cardinal-once condition, no such conflict exists. There is a single row of objects that, 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 posttransformation values of rows. As with the untransformed row, it answers the question "Are there still *N*?" by requantifying. Because no comparison is made to the original value, no learning can take place regarding whether or not the transformation has had any effect on the numerical value of the row.

Effect Determination Operations. There are two operators in the DETERMINE-EFFECT space. The RECALL operator recalls a pretransformation value for comparison to the posttransformation value. The DETERMINE-EFFECT operator matches pre- and posttransformation values for the transformed row as previously described. The process is a simple match that tests whether or not the values are the same. 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 posttransformation 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.

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.

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 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 that is primitive to the architecture and that 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* (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 pretransformation and posttransformation 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.

Preceptual 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 8 years of age. Thus, estimation in conservation settings is inaccurate, because 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. 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), that 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.

Categorization

As mentioned earlier, Q-Soar categorizes observed transformations.[4] 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

[4] 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

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 (1991), 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 because 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).

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 transformations 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.

Learning Conservation Knowledge in Q-Soar

The preceeding 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 on 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 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 preexisting 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 preclassified examples of concepts are presented for the system to learn. 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.

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.

One can evaluate an enterprise such as that 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?

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 posttests is that, for awhile, 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.

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: 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. In other words, this is a problem of representation change. The child must move from domain-specific characterizations of the effects of transformation classes in terms of discrete number to representations where the effect is separated from the dimension that it impacts. This would allow commonalities between transformations that have a "more" or a "same" effect to be noticed, irrespective of what it is that they are affecting. In their account of the development of analogical reasoning, Gentner, Ratterman, Markman, and Kotovsky (chap. 6,this volume) present a transition mechanism that is concerned with precisely this sort of representation change.

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 is to generate a 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) provided 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).

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.

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 posttransformation values. Apart from this difference, both versions have all the described capabilities.

However, because neither variant has undergone any training or 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 nonconservers. We now examine their behavior after training. Because the no-cardinal condition is not expected to induce any change in behavior, we discuss only the results of the other two conditions.

Q-Soar–3 After Cardinal-Once Trials. Without employing its memorization capability to recall and compare pre- and posttransformation values, Q-Soar–3 cannot learn anything about the numerical effect of observed transformations. Thus, based on Assumption A2, it always assumes that the value of the row changes. Because 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 7.1, 3-year-olds produced few correct responses.

Q-Soar–3 After Experimental Trials. As explained, the conflicting information in experimental trials after a transformation induces Q-Soar–3 to recall and compare values in the same way as does 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 is thus consistent with the pattern of results in Table 7.1. It seems likely that the experience of this conflict and the resulting recall and comparison of values provide the means by which 3-year-olds acquire the effect-verification knowledge we have assumed to be available to 4-year-olds and that we provided for Q-Soar-4.

Q-Soar–4 After Both Trials. Having produced a quantitative response before a transformation (e.g., Step 4 in Fig. 7.1), 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 posttransformation 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 7.1. The higher proportion of correct responses in the experimental group may reflect the fact that not all 4-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 3-year-olds.

Learning Speed

We stated that Q Theory is designed to account for the natural development of conservation, whereas Q-Soar simulates only conservation learning in a single training study. Therefore, we should explain how the same processes can learn 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 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 or 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.

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. The central concept in our theory is that conservation learning is premised on young children's ability to make and use measurements. These measurements are used to make conservation *judgments*—evaluations of the quantitative effects of observed transformations. Therefore, the first kinds of conservation processing that children carry out are empirical. Young children's measurement capabilities are limited to small, discrete quantities and so the first kind of quantity for which conservation judgments can be made is number. The results of these number conservation judgments are turned, by a learning mechanism, into new conservation *knowledge*. Due to the nature of the learning mechanism, that new knowledge applies to more cases than just the one it was constructed from. When similar new number conservation problems are attempted, the new knowledge is immediately retrieved and no effect-determination is required.

Thus, we demonstrated a developmental shift where the child moves from empirically determining the effect of transformations via measurement, to making direct inferences about the necessity of conservation based on prior knowledge. This reverses the logical relationship between measurement and conservation that existed in Piaget's theory and now makes measurement a prerequisite for conservation learning. We also identified transformations involving small discrete collections as the learning events that children use to acquire number conservation knowledge. This opposes Piaget's view that conservation is a domain-independent principle that children acquire, by demonstrating that it arises from and initially applies only to domain-specific experiences with transformations relating to number. Furthermore, we demonstrated that Soar's chunking mechanism is sufficient to account for significant developmental transitions, such as the acquisition of number conservation knowledge. This challenges the Piagetian view that developmental change mechanisms are distinct from simple learning mechanisms. Chunking in Soar began as a model of practice effects in human learning and has since been extended to a wide range of cognitive phenomena (Lewis et al., 1990).

Finally, we showed that not only can Q-Soar account for the rapid learning observed in the Gelman (1982) training study, but also, without modification, it may 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 progress in the creation of computational theories of conceptual development.

APPENDIX: SAMPLE Q-SOAR RUN

Here we illustrate how the problem spaces generate behavior when Q-Soar is presented with a task. The following trace is an abstracted version of the steps presented in Fig. 7. 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

```
                                                      0 0 0 0
                                                      [] [] [] []
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
                                                  0000
                                                  [] [] [] []
                                                  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
```

The following rule paraphrases the chunk that Q-Soar learns at the point marked ** in this run and that applies at the point marked ** in the following trace. Pattern-match variables are preceded by question marks.

If Goal ?G1 has State ?S1 and Operator ?O1,
  and State ?S1 has Transformation ?T1 marked on it,
  and Transformation ?T1 is Spreading,
  and Operator ?O1 is Determine-Response,
Then mark State ?S1 with Effect ?E1 of Transformation ?T1,
  where ?E1 states that ?T1 has the effect NONE on the property
  Number, because pre- and posttransformation numerical values matched

Q-SOAR RUN ON A CONSERVATION TEST AFTER LEARNING

```
                                                  00000000
                                                  [][][][][][][][]
P: (CONSERVATION)
                                                  Same or different number?
O: (COMPREHEND-INSTRUCTIONS)
O: (DETERMINE-RESPONSE)
==>G:(OPERATOR NO-CHANGE)
  P: (RESPOND)
  O: (COMPARE)
  ==>G:(OPERATOR NO-CHANGE)
    P:                     (ESTIMATE)
    O:                     (MATCH-1-TO-1)
                           One-to-one end-match: SAME
    O: (MEMORIZE)
O: (RETURN-RESPONSE)       Answer SAME
                                                  00000000
                                                  [][][][][][][][]
                                                  Same or different number?
                                                  Explain
O: (CATEGORIZE-TRANSFORM)
O: (COMPREHEND-INSTRUCTIONS)
O: (DETERMINE-RESPONSE)**
O: (RETURN-RESPONSE)       Answer SAME
      VALUES MATCH BEFORE AND AFTER THIS TRANSFORMATION
End—Explicit Halt
```

ACKNOWLEDGMENTS

This chapter is a minor revision of work completed in collaboration with the late Allen Newell Simon, Newell, & Klahr, 1991; Simon, Klahr, & Newell, 1992). We acknowledge the profound impact of Allen's intellectual vision on this project, and we are grateful to him for inviting us to explore with him one of the "frontiers" of his unified theory of cognition (Newell, 1990).

We are sure this would have been a better chapter had Allen Newell lived long enough to make a sustained contribution to it.

We thank Robert Siegler for his comments on earlier drafts, as well as for granting us access to his experimental data, and Rochel Gelman for further explication of her experimental procedures. Finally, the first author wishes to thank members of the Soar group for invaluable help, discussions, and support. This work was funded in part 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.

Atkinson, J., Campbell, F. W., & Francis, M. R. (1976). The magic number 4 plus or minus 0: A new look at visual numerosity judgements. *Perception, 5,* 327–334.

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

Briars, D., & Siegler, R. S. (1984). A featural analysis of preschooler's counting knowledge *Developmental Psychology, 20,* 607—618.

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 Associates.

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 pp. 77–166). New York: Wiley.

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

Cowan, R. (1979). Performance in number conservation tasks as a function of the number of items. *British Journal of Psychology, 70,* 77–81.

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.

Fuson, K. C., Secada, W. G., & Hall, J. W. (1983) Matching, counting and conservation of numerical equivalence. *Child Development, 54,* 91–97.

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, pp. 219–283). 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, pp. 168–230). New York: 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 Associates.

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

Jones, R. M., & VanLehn, K. (1994) Acquisition of children's addition strategies: A model of impasse-free, knowledge-level learning. *Machine Learning, 16*, 11–30.

Klahr, D. (1973). Quantification processes. In W. Chase (Ed.), *Visual information processing* (pp. 3–34). 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: 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 Associates.

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). Hillsdale, NJ: Lawrence Erlbaum Associates.

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

McClelland, J. L. (1989). Parallel distributed processing: Implications for cognition and development. In R.G. Morris (Ed.), *Parallel distributed processing: Implications for psychology and neurobiology* (pp. 8–45). Oxford, England: Clarendon Press.

Markman, E. M. (1990) Constraints children place on word meanings. *Cognitive Science, 14*, 57–78.

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

Miller, K. F. (1984). Child as measurer of all things: Measurement procedures and the development of quantitative concepts. In. C. Sophian (Ed.), *The origin of cognitive skills* (pp. 193–228). Hillsdale, NJ: Lawrence Erlbaum Associates.

Miller, K. F. (1989). Measurement as a tool for thought: The role of measurement procedures in children's understanding of quantitative invariance. *Developmental Psychology, 25*, 589–600.

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

Mooney, R. (1991) Explanation-based learning as concept formation. In D. Fisher & M. Pazzani (Eds.), *Concept formation: Knowledge and experience in unsupervised learning* (pp. 174–206). San Mateo, CA: Morgan Kaufmann.

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.

Piaget, J., Inhelder, B., & Szeminska, A. (1960). *The child's concept of geometry*. London: Routledge & Kegan Paul.

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). San Mateo, CA: Morgan Kaufmann.

Rosenbloom, P. S., & Newell, A. (1986). The chunking of goal hierarchies: A generalized theory of practice. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artifical intelligence approach,*)Vol. 2, pp. 247–288). Palo Alto, CA: Tioga.

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* (pp. 19–30). San Francisco, CA: Jossey-Bass.

Shipley, E. F. & Shepperson, B. (1990) Countable entities: Developmental changes. *Cognition, 34,* 109–136.

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

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

Simon, T., Cabrera, A., & Kliegl, R. (1993). A new approach to the study of subitizing as distinct enumeration processing. *Proceedings of the 15th Annual Meeting of the Cognitive Science Society* (pp. 929–934). Hillsdale, NJ: Lawrence Erlbaum Associates.

Simon, T., Klahr, D., & Newell, A. (1992). The role of measurement in the construction of conservation knowledge. *Proceedings of the 14th Annual Meeting of the Cognitive Science Society* (pp. 66–71). Hillsdale, NJ: Lawrence Erlbaum Associates.

Simon, T., Newell, A., & Klahr, D. (1991). A computational account of children's learning about number conservation. In D. H. Fisher, M. J. Pazzani & P. Langley (Eds.), *Concept formation: Knowledge and experience in unsupervised earning* (pp. 423–462). San Mateo, CA: Morgan Kaufmann.

Sophian, C. (1987). Early developments in children's use of counting to solve quantitative problems. *Cognition and Instruction, 4,* 61–90.

Starkey, P. (1992). The early development of numerical reasoning. *Cognition, 43,* 93–126.

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

van Loosbroek, E., & Smitsman, A. W. (1990). Visual perception of numerosity in infancy. *Developmental Psychology, 26,* 916–922.

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* (pp. 359–436). Cambridge, MA: MIT Press.

Wynn, K. (1990). Children's understanding of counting. *Cognition, 36,* 155–193.