Counting Strategies and Semantic Analysis as Applied to Class Inclusion

ALEXANDER WILKINSON

University of Michigan

This study examined strategic and semantic aspects of the answers given by preschool children to class inclusion problems. The Piagetian logical model for class inclusion was contrasted with an alternative, problem processing model in three experiments. A major component of the alternative model is an enumeration strategy which is advantageous for learning reliable counting skills. The counting strategy was found to explain the inclusion errors of young children better than did the logic of the task. It was also found that young children understand the semantics of inclusion but are unable to coordinate their semantic knowledge with their counting strategy. Methodologically, one of the experiments suggested a fruitful extension of task analysis (Simon, 1969) to experimental design.

The class inclusion problem occupies a central place in the Piagetian theory of cognitive development (Piaget, 1970). In this problem, a child must compare the numerosity of a part or subclass with that of its superordinate whole or supraclass (e.g., more dogs vs. more animals). When making these comparisons, young children commonly, but mistakenly, name the included subclass as more numerous. Genevan psychologists, in their more recent studies, have used this problem to examine children's competence in logical reasoning (Inhelder & Piaget, 1964; Inhelder & Sinclair, 1969). This view of class inclusion as a logical problem contrasts somewhat with Piaget's earlier (1952) analysis, in which children's performance on class inclusion was compared with their ability to conserve number. The present study returned to the earlier view by using the inclusion paradigm to investigate the development of enumeration skills. These skills were also examined for their interaction with semantic processes that mediate the resolution of verbally communicated problems.

PIAGET'S LOGICAL MODEL

Piaget's model of inclusion performance is formalized as the logical operation:

$$(A + A' = B) \longleftrightarrow (B - A' = A).$$

Here the letters indicate classes, the equations within parentheses indicate relations among classes, and the double arrow indicates that two classification schemes are reversible or interchangeable. For example, the logically equivalent schemes might be (the dogs plus the cats equal the animals) on the left side of the operation, and (the animals less the cats equal the dogs) on the right side. Inclusion errors are said to arise from the absence of a fully reversible operation. Once young children have decomposed the whole into parts by applying the operation from left to right, they are not free to recompose the whole by applying the reverse operation from right to left. They are thus unable to compare the numerosity of a decomposed part with that of its recomposed whole. (For elaboration, see Flavell, 1963, pp. 172–176, 190–191; Piaget, 1970.)

Taken literally, this model implies that inclusion problems having the same logical structure should all elicit the same pattern of errors from young children. Empirically, however, there is wide variation in performance on problems having this basic structure (Klahr & Wallace, 1972a). Admittedly, even Inhelder and Piaget (1964) do not interpret the model so strictly, but neither do they offer a qualifying amendment to their logical equation. Furthermore, no detailed description has yet been given of the psychological processes that accomplish the logical reversal. These deficiencies suggest the need for an alternative model.

A PROBLEM PROCESSING MODEL

The model to be proposed here emphasizes two aspects of problem processing: semantic analysis and the use of goal-directed strategies.

Semantic Analysis

As used in this model, *semantic analysis* indicates processes that translate strictly grammatical analysis into a characterization of the relevant problem-space. For example, given the verbally posed problem *Are there more dogs or more animals?* a strictly grammatical analysis would reveal, among other things, that *more* modifies both *dogs* and *animals*. Of course, such grammatical comprehension would hardly be sufficient to solve the problem. There must also be a semantic analysis to interpret the phrase *more dogs or more animals* as requesting a quantitative comparison of two classes. In addition, it would be necessary to determine from context the intended reference of *dogs* and *animals*. These nouns might refer in one context to concrete classes shown in a picture, while they could refer in another context to abstract classes mentioned in some immediately preceding conversation.

When viewed in this way, the result of semantic analysis may be said to be twofold. A *start-state* is defined which identifies the referential and contextual sources of relevant information from which a solution may be

extracted. In addition, an *end-state* is defined which specifies the goal that must be reached for the problem to be solved. For the example given above, the start-state might be that the relevant dogs and animals were those shown in a picture. The end-state would be the achievement of two quantifications that need only be precise enough to determine which of the two classes is more numerous. In the experiments reported below, problems were examined which had similar start-states but different end-states, or the reverse.

This view of semantic analysis differs from that used by other investigators of class inclusion. For both Hayes (1972) and Markman (1973), semantic analysis was concerned with the assignment of meanings to individual key words in the inclusion question. Here, in contrast, semantic analysis refers to the more integrative meaning assigned to the verbal statement of the inclusion problem as a whole. Moreover, this holistic meaning is characterized as the conjoint specification of a start-state and an end-state that are suitable for manipulation by a problem-solving system. A similar definition of semantic analysis has been proposed, and successfully implemented on a computer, by Winograd (1972).

Goal-Directed Strategies

Problem-solving strategies find a connecting path between the start-state and end-state that have been defined by semantic analysis. Strategies for finding goal-directed paths of this type are often used in models of artificial intelligence (Nilsson, 1971) and in research which simulates the problem-solving of children and adults by means of production systems (Klahr, 1973; Klahr & Wallace, 1972a; Newell & Simon, 1972). The present model, like a production system, assumes that a strategy either is stored in toto in memory or is assembled ad hoc from subroutines that are stored in memory. The relevance of strategies for class inclusion is that even when young children have correctly analyzed the semantics of the problem, their limited repertoires of enumeration techniques may only permit them to find false paths to the desired goal. The experiments reported below attempted to separate the semantic from the strategic components of inclusion performance.

Essential Features of the Model

Two counting strategies are assumed to form a developmental sequence. Strategy I, the earlier of the two, forbids double-counting. This constraint is assumed to derive from young children's discovery that when they are attempting to determine how many members are in a class, they must exercise care to count each member once, but only once. For this reason, the child who is just learning to count should find it advantageous to use a

counting strategy that prevents any particular item from being counted twice. But a similar strategy, when employed for an inclusion problem, would lead the child to an incorrect solution. It would prohibit items counted as subclass members from being counted again as supraclass members. The postulated constraint against double-counting serves the same explanatory purpose as Piaget's concept of irreversibility. However, in the problem processing model the child is viewed not as being deficient in logical capacity, but as preferring a counting strategy that is adaptive for learning to count but maladaptive for class inclusion.

In actual use, a counting strategy might be applied to items appearing in a picture. Consider, for the sake of simplicity, a picture in which several dogs each exhibit a common pattern P_1 (dog-like) and several cats each exhibit a different pattern P_2 (cat-like). Suppose that a child uses Strategy I to count the single class of animals in this picture. In this case, the child would first count the dogs, then the cats. The counting of animals would thus be reduced to two subproblems corresponding to the patterns P_1 and P_2.

This process of problem-reduction is an important feature of the model. The path from start-state to end-state is stepwise, proceeding by the successive reduction of higher and more complex goals to simpler goals.

To achieve the reduction of a counting problem to subproblems, a child is assumed to use a SCAN operator. Briefly, SCAN is a recursive procedure by which an array of items is first scanned or inspected, then conceptually subdivided into mutually exclusive patterns. The term *pattern* is used in a very broad sense to indicate any organizational scheme, including appearance (e.g., color), relative location (e.g., spatial grouping), and even direction (e.g., left-to-right ordering over a specified range). Wohlwill (1968) and more recently Isen, Riley, Tucker, and Trabasso (1975) have reported that salient visual items and salient relations among items exercise an important influence on how children respond to inclusion questions. In view of these findings, it is assumed here that children use as patterns only the more prominent features of an array. For each such pattern or subdivision defined by SCAN, a SUBSCAN enumerates the corresponding subclass of items. Depending on the nature of the problem, an enumeration by SUBSCAN may be a precise numerical count, or it may be a roughly estimated count based on a perceptual judgment of length, area, or density. The flow-charts in Fig. 1 give more details.

Empirically, there is evidence that young children do, in fact, subdivide a single class according to mutually exclusive patterns (Potter & Levy, 1968) and that their counting is more likely to be accurate when the class can be subdivided than when it cannot (Schaeffer, Eggleston, & Scott, 1974). Subdivision presumably helps the child to remember which items have already been counted and which have not.

Using SCAN to accomplish a subdivision may be helpful to the child in an additional way. Suppose that the child wants to determine whether there
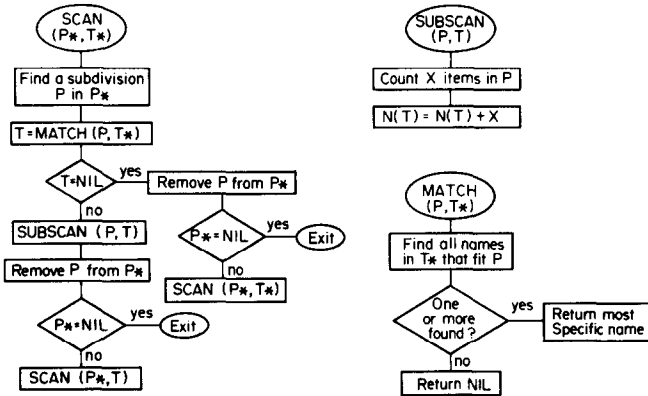
FIG. 1. Flow-charts for SCAN, SUBSCAN, and MATCH. Key to terms: P^* is a set of patterns P (e.g., the patterns dog-like and cat-like). T^* is a set of verbal targets T (e.g., the targets *animals* and *dogs*). NIL is the empty set. $N(T)$ is the cumulative count of T. Initially all $N(T) = 0$.

are more dogs or more cats, given the same picture as before. This problem requires a comparison of disjoint classes, so a subdivision to the mutually exclusive patterns P_1 (dog-like) and P_2 (cat-like) would clearly be appropriate. Strategy I could again be used to advantage.

Recognizing the utility of this strategy for solving counting problems of various types, perhaps the child would be inclined to use it indiscriminately. Suppose, as a final example, that the child is shown the same picture again and is asked whether there are more dogs or more animals. This is an inclusion problem. A subdivision by Strategy I would in this case yield a pattern P_1 (dog-like), which could properly be counted as the dogs, and a separate pattern P_2 (cat-like), which could coincidentally be counted as the animals, because cats are indeed animals. The unfortunate result would be a comparison between dogs and cats, not the intended one between dogs and animals. Strategy I, then, will produce a solution for an inclusion problem, but the solution is erroneous.

One way the child could achieve a correct solution would be to use a counting strategy that employs SCAN more than once. One SCAN could count the total class of animals, as in the first example discussed above; another SCAN could count the dogs; and the respective totals could then be compared. This procedure is the essence of Strategy II, which is assumed to be developmentally more advanced. Strictly speaking, this counting method does not require that any logical inference be made. Its repeatable use of SCAN, however, is an approximate analog of the Piagetian concept of full reversibility.

The greater complexity of Strategy II may be one reason for its later appearance during development. Necessitating an initial use of SCAN in one manner and its immediate reuse in another manner, Strategy II would

very likely require greater proficiency in problem-solving than Strategy I. More important, there is a very good reason why the young child might be reluctant to have anything like Strategy II available in a repertoire of frequently used problem-solving techniques, especially if the child's skills in counting are still rudimentary. By permitting more than one use of SCAN on a particular problem, any such strategy would make the child too often vulnerable to errors of double-counting.

This introduction to the problem processing model has intentionally omitted some details and ignored some difficulties that will be considered below. Hopefully, the essential features of the model are clear. The adequacy of the model was compared in three experiments with that of the Piagetian logical model. Consideration was restricted in these experiments to the causes of inclusion errors in very young children; the nature of older children's proficiency in class inclusion was not examined. The subjects were all preschool children, 4 and 5 years of age. None had reached their fifth birthday in time to qualify for kindergarten.

EXPERIMENT 1

Parts (a) and (b) of Fig. 2 illustrate the two types of inclusion problems used in Experiment 1. Type (a), called *concept inclusion*, is a standard problem in which a pattern P_1 noticeably marks the subclass A (the boys), and a different pattern P_2 similarly marks A' (the girls), but no equally prominent pattern identifies uniquely the supraclass B (the children). This problem corresponds to the inclusion question *Are there more boys or more children?* Type (b) in Fig. 2, called *percept inclusion*, corresponds to the
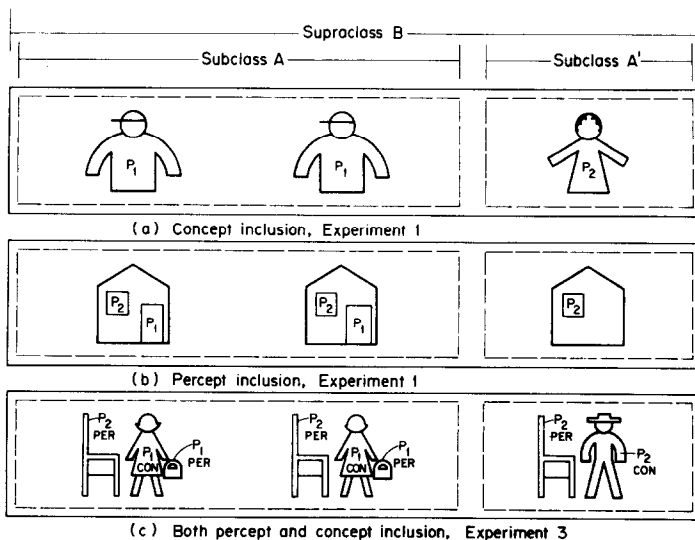

FIG. 2. Examples of problems for Experiments 1 and 3.

inclusion question *Are there more houses that have a door or more houses that have a window*? In this case, the subclass A (houses having a door) corresponds directly to P_1 (door), and the supraclass B (houses having a window) corresponds directly to P_2 (window).

A literal reading of the Piagetian model predicts that performance on percept and concept inclusion should be the same, because they have the same logical structure $(A + A' = B)$. The problem processing model, however, makes a different prediction. Since, in the case of percept inclusion, the subclass A is marked by two separable patterns P_1 and P_2, this class may be counted first as one pattern and then again as the other pattern. In this way, the subclass may be counted twice, even though each pattern is counted only once. There could be strict observance of the Strategy I constraint against the double-counting of patterns, yet the subclass A could still be correctly included in the enumeration of the supraclass B. Accordingly, it was anticipated that children who used Strategy I would perform well on percept inclusion, despite their poor performance on concept inclusion.[1]

A second hypothesis was also tested. Presumably young children prefer Strategy I, but it may be possible to induce them to adopt a more advanced strategy under facilitating circumstances. As an attempt to provide such circumstances, a special procedure was used with some of the children, in the expectation that it would improve their performance. The procedure required the children to execute, just before the inclusion question was asked, a pointing sequence that did not involve overt counting, but was, nevertheless, isomorphic with the sequence of attention deployment characteristic of Strategy II.

Method

Subjects

The subjects were 24 girls and 24 boys who attended nursery schools in Ann Arbor, Michigan. Six additional children from the same schools were excluded from the final sample because they failed to satisfy the performance criterion on a control task, as described below.

Materials

The stimulus materials were prepared on white cards, 13 × 20 cm, which comprised three distinct sets of problems.

1. Five concept inclusion cards each depicted three children, approximately 4 × 1 cm. The

---

[1] Although performance on percept inclusion was expected to be better than on concept inclusion, it was still not expected to be perfect. Erroneous answers could result even on percept problems if the child using Strategy I assigned patterns to classes in a manner different from that shown in Fig. 2(b). For example, P_1 could be defined as (having both a door and a window), and P_2 could be defined as (not having both), in which case the subclass A would mistakenly be found to be more numerous. In general, the assignment of patterns to classes is a problematic issue which the present model handles less well then would be desirable.

supraclass of three children included subclasses of either two identical boys and one girl as in Figure 2(a), or one boy and two identical girls.

2. Five percept inclusion cards each depicted three houses, approximately 4×3 cm. Like the example in Figure 2(b), there were always two classes (houses having a door and houses having a window), of which one represented a supraclass of all three houses, and the other an included subclass of two houses.

3. Four control cards were similar in construction. Their purpose was to assess the child's ability to compare the numerosity of a class of three items with that of a class of two items, given mutually exclusive classes that did not require double-counting of particular items. There were thus five items on each card rather than three. Two of the control cards had either three houses with a door and two with a window, or the reverse, but no house had both a door and a window. For these cards the child was asked, "Are there more houses that have a door or more houses that have a window?" just as in the case of the cards for percept inclusion. The other two control cards had either two boys and three girls, or the reverse, and for these cards the subject was asked, "Are there more boys or more girls?"

The restriction, in all problems, to comparisons of three vs. two was motivated by Gelman's (1972) report that young children can accurately compare exact numerosities only for very small numbers. A child's data were used in the analyses reported below only if that child demonstrated competence in making comparisons of this magnitude by correctly answering at least three of the four control questions. Of the 48 children who satisfied this requirement, eight gave three correct answers and 40 gave four.

In all three sets of problems, the depicted items were arrayed along a horizontal line, with equal spacing between items. Within each set, half the comparison problems had items of each subclass contiguously juxtaposed (e.g., boy–boy–girl), while the other half had them placed discontiguously (e.g., boy–girl–boy).

Procedure

All subjects were shown all three sets of problems. The first card in each of the percept and concept inclusion sets was a familiarization card whose purpose was to ensure that the child could correctly match the descriptive name for a class with its pictorial representation. The child was shown the familiarization card first, and was asked questions of the form: "How many of these houses have a door? How many have a window?" and "How many boys are there here? How many girls? How many children?" Most subjects answered these questions correctly on their first attempt, but if not, the questions were restated more explicitly as, e.g., "Is that *all* the children? Tell me how many are all the children?" One or at most two restatements of this type were always sufficient to elicit a spontaneously correct answer from each subject retained in the final sample.

For half the children of each sex, the order of administration was percept inclusion, concept inclusion, control; for the other half, the order of the first two sets was reversed. The control set was always given last in order to prevent any potential effect it might have had of biasing the child toward making exclusive comparisons on subsequent inclusion problems.

Half the children in each sex × order group were assigned to the no-pointing condition. In this condition, a child was simply asked the relevant comparison question for each card. The other half were given pointing instructions just preceding each inclusion question. The instructions were of this type: "Point and show me all the boys. (Pause) Now show me all the children." If the child pointed incorrectly, prompting instructions were given, such as: "Is that *all* the children? Show me *all* the children." The sequence of hand movements required by the pointing instructions was identical to the order in which patterns would be enumerated by means of Strategy II.

The order of mention for the subordinate and superordinate terms was fixed for a given problem, but was counterbalanced across problems in each set. This order of mention was the same for the pointing instructions (if given) as for the comparison question itself. Within the

percept set it was possible to counterbalance whether *houses that have a door* or *houses that have a window* represented the supraclass and was thus the correct answer. Of necessity, the term *children* always named the correct answer in the concept set.

Results and Discussion

Preliminary analyses investigated sex differences and effects due to order of presentation. No reliable differences were found, and .95 confidence limits indicated that the maximum population difference for any of the observed sex or order contrasts was not more than a single correct response.

Confidence limits were obtained in other analyses as well, and these limits are reported below in the notation of the example 5 (4, 6), where 4 and 6 are the .95 population limits for a statistic whose sample value was 5. All confidence levels are .95 unless indicated otherwise.

Comparison of Percept and Concept Performance

Or primary interest was the difference in performance within a subject on percept vs. concept problems. Achieved levels of performance were defined as the proportions of children reaching the criterial level of at least three correct answers for the four problems in an inclusion set. Table 1 shows these proportions by pointing instructions and inclusion type. Reading across each row separately in Table 1, it can be seen that for both rows the confidence interval in the percept column does not overlap with that in the concept column. This nonintersection indicates that under both pointing and no-pointing instructions, reliably more of the children achieved the performance criterion on percept than on concept problems.

As implied by these grouped data, individual children almost always performed better on percept inclusion than on concept inclusion. Eliminating the children who answered either all eight questions correctly or none correctly, the proportion of the remaining children who gave at least one more correct answer for percept than for concept inclusion was .88 (.73, .97), $N = 43$.

These findings contradict a strict reading of the Piagetian model, since percept and concept problems having identical logical structures elicited

TABLE 1

PROPORTIONS OF CHILDREN ACHIEVING THE PERFORMANCE CRITERION[a]

Instructions	Percept inclusion	Concept inclusion
Pointing	.83 (.60, .96)	.29 (.12, .56)
No-pointing	.63 (.40, .82)	.00 (.00, .16)

[a] $N = 24$ per cell. The values in parentheses are .95 confidence limits.

clearly different patterns of performance. It might be argued, however, that the percept problems were less difficult because the children did not actually make class comparisons. Perhaps they counted doors and windows without ever comparing the class of houses that have a door with the class of houses that have a window. Some recognition of the classes must have occurred, however, at least at the level of a grammatical analysis which identified the full description of each class as a distinct noun phrase. If the children subsequently ignored the class aspects of the problem, they may have done so because they retained only the information that was vital for their use of Strategy I. This interpretation of the children's performance is consistent with the view that inclusion questions actually test children's counting strategies, not their skill in manipulating the logic of class relations.

There were, however, methodological difficulties which made this interpretation uncertain. These difficulties were examined in Experiment 3, and are discussed below in the section which reports that experiment.

Effects of Pointing Instructions

The effects of the pointing procedure were first examined by combining the percept and concept questions and computing the total number of correct answers given by a subject for the combined set of all questions. The mean number correct was greater with pointing instructions (4.9, $s = 1.8$) than without them (3.2, $s = 1.4$). A Mann–Whitney test revealed that this difference was reliable, $z = 3.04, p < .01$. However, .99 confidence limits indicated that the gain due to pointing could be negligibly greater than zero, and could be no more than three additional correct responses out of the total of eight questions. So although reliable, the effect was not impressively large.

Returning to Table 1, the pointing procedure was also examined for its effect on the proportions of children achieving the performance criterion on each type of inclusion problem. For concept inclusion, reading down the right-hand column of the table, it can be seen that the confidence interval for pointing subjects overlaps somewhat with that for the no-pointing subjects. Despite this slight intersection, the improvement associated with pointing was reliable in the case of concept inclusion, $p < .01$ by Fisher's exact test (which is more powerful as a single test than the two separate confidence intervals). For percept inclusion, reading down the left-hand column of the table, the intersection of the confidence intervals was much larger, and the gain due to pointing was not trustworthy, Fisher's $p = .10$. There was no unequivocal indication in the data that the pointing instructions interacted in any additional way with sex, problem-type, or order of presentation.

The gain due to pointing may have reflected the children's momentary adoption of Strategy II, as hypothesized. If so, it is not surprising that the

effect was more reliable for concept inclusion, which virtually required the more advanced strategy, than for percept inclusion, which did not.

Unfortunately, a weakness in the experimental design permitted an alternative interpretation of these results. The weakness was that the pointing sequence immediately preceded each comparison question; consequently, the children's greater success could have resulted from memory for numbers counted covertly at the time of pointing. The hypothesis of immediate memory was unanticipated. Nevertheless, it is generally congruent with the problem processing model, for it implies that pointing prompted the children to count as they pointed, helping them in this way to separate the two classes as they counted and to keep the two totals separate in memory.

Thus the original hypothesis and the memory hypothesis both conform to the problem processing model, although for different reasons. Neither hypothesis, however, seems consistent with the Piagetian model. There is no obvious way in which the pointing sequence could have served to correct a logical deficit.

Of course, the two hypotheses are also inconsistent with each other, but a choice between them is not justified by the available data. It therefore remains to be demonstrated whether pointing provided a partial substitute for the advanced counting method, offering the children a manual analog which they could then use to guide their own discovery of the full strategy, or whether, instead, it provided a nearly complete substitute, requiring the children only to preserve the separate identities of previously computed totals.

EXPERIMENT 2

An essential feature of the SCAN operator, as illustrated earlier in Fig. 1, is its use of a MATCH routine.[2] MATCH enforces the requirement that whenever a single SCAN is used, for example, to count the two classes of dogs and animals, each dog-like pattern must always be counted as a dog, never as an animal. A pattern is assigned by MATCH to the class which names the pattern most specifically, considering only the target classes which are to be counted during the current application of SCAN.

An interesting prediction follows from this feature. It concerns a *coextensive* comparison in which the semantic supraclass and subclass are equivalent (e.g., all the available animals are dogs). In this case, Strategy I would first enumerate the subclass term, since every item would have to be MATCHed as more specifically identifiable by its subclass than by its supraclass name. But then the prohibition of Strategy I against

---

[2] The MATCH feature was motivated by a deficiency noted by Hayes (1972) in a production system designed as a model for class inclusion by Klahr and Wallace (1972a, 1972b; see also the modification of the model by Klahr, 1973).

double-counting would ensure that the count for the remaining supraclass term would erroneously equal zero. Children who use Strategy I should therefore resolve a coextensive comparison of this type by answering that the subclass is more numerous. Shown three dogs, for example, and asked whether there are more dogs, more animals, or the same number, they should always answer that there are more dogs, never that there are more animals or that the numbers are equal. Experiment 2 tested this prediction.

Method

Subjects

The 11 children in Experiment 2 included all but three of the children who, in Experiment 1, had given a correct answer to every percept question and an erroneous answer to every concept question. This criterion ensured that the subjects in Experiment 2 were children who consistently counted as if they were using Strategy I. They had shown no tendency to guess or to use a nonsystematic strategy.

Materials and Procedure

The stimulus materials were 15 cards similar to the ones used in Experiment 1. Three sets contained five cards each: (a) concept cards with children (boys and girls), (b) concept cards with animals (dogs and cats), and (c) percept cards with houses (having a door and/or a window). The sets were presented to all subjects in the order just given. The problems within each set were of three types.

2 + 1 Inclusion problems. Presented on the first and third card in each set, these problems had a supraclass of three items which included complementary subclasses of two and one items (e.g., two dogs and one cat = three animals). The 2 + 1 problems were comparable to the ones used in Experiment 1, and their purpose was to assess whether the child's current counting strategy would produce the same response patterns as before.

2 + 2 Problems. Presented on the second card in each set, these problems were designed to determine whether the answer *the same number* was available in the child's response repertoire. As the necessarily correct answer for any coextensive comparison, this answer must be available to the child as a minimum requirement for correct coextension performance. In the problems depicting animals and children, the 2 + 2 problems presented an inclusive comparison (e.g., dogs vs. animals, given two dogs and two cats = four animals). A child who did not double-count would give *the same number* as an answer to an inclusion question of this kind. This answer would also be expected for an exclusive comparison of the percept type (e.g., houses having a door vs. houses having a window, given two houses having only a door and two houses having only a window = four houses). The 2 + 2 problems for the percept set were designed in this way. Consequently, *the same number* was the predicted answer for all the 2 + 2 problems.

3 = 3 Coextension problems. Presented on the last two cards in each set, these problems had a supraclass of three items which was coextensive with a subclass of the same three items (e.g., three dogs = three animals). It was expected that the children would erroneously name the semantic subclass as their answer for each coextension problem in the two concept sets. However, because they had previously demonstrated competence in percept inclusion, the children were expected to correctly answer *the same number* for the coextension problems in the percept set.

The comparison question was always phrased, "Are there more (class 1), more (class 2), or the same number?" The respective order of mention for the subclass and the supraclass was counterbalanced across questions in each set. At the very beginning of the experiment, and

before presentation of any comparison problems, the experimenter demonstrated with his fingers the meaning of *the same number.*

Results and Discussion

As just explained, specific responses were predicted for each of the 15 questions asked of each child. The median number of the 15 questions answered by a child exactly as predicted was 14 (11, 15), $p = .99$.

For the 3 = 3 problems in the concept sets, the predicted (incorrect) answer was the name of the semantic subclass. The median number of times a child gave this predicted but incorrect answer to the four questions of this type was 4 (3, 4), $p = .99$.

Two lines of evidence implied that these consistently incorrect answers could not be attributed merely to the unavailability of the correct answer, which was *the same number.* First, for the comparable 3 = 3 coextension problems in the percept set, most children correctly answered *the same number,* as expected. There were only two of these problems, but eight of the 11 children answered both correctly. Second, *the same number* was also the predicted answer for the 2 + 2 problems, of which there were three. Here the median number of times a child gave the expected answer of equivalence was 3 (2, 3), $p = .99$.

Finally, there was no marked indication that any child's counting strategy had changed during the time between Experiments 1 and 2. In the present experiment, the 2 + 1 inclusion problems numbered six in all, and the median number of these six questions answered by a child in the expected manner was 6 (4, 6), $p = .99$.

These results offered consistent empirical support for two components of the problem processing model. In support of the SCAN component, children whose inclusion performance suggested their use of Strategy I were indeed operating under a constraint against double-counting, and this constraint applied to coextensive as well as inclusive comparisons. In support of the MATCH component, these children identified the patterns they were counting as corresponding to the more specific class name, if two class names were available to be used.

EXPERIMENT 3

The primary purpose of Experiment 3 was to clarify the reasons for the difference observed in Experiment 1 between the percept and concept problems. One explanation of this difference was that some idiosyncratic property of the particular category used for the percept problems in that experiment (i.e., houses) may have made them easier than the concept problems which were drawn from a different category (i.e., children). In Experiment 3 this confounding of problems and categories was eliminated. Four categories were used, and for each category both a percept and a

concept problem were prepared. Part (c) of Figure 2, presented earlier, provides an example that corresponds to the category of grown-ups. Associated with this category were both (a) the percept problem *Are there more grown-ups who have a picnic basket or more grown-ups who have a chair?* and (b) the concept problem *Are there more mothers or more grown-ups?* A percept and a concept problem which are related in this manner have virtually identical start-states or sources of contextual information, because their pictorial representations are the same. Such identity of start-states minimized the likelihood that idiosyncratic properties of pictorial representations could cause a difference in performance between the two types of problems.

Another possible explanation of the difference in performance concerned the child's comprehension of semantic hierarchies. Whether percept and concept problems have different start-states as in Experiment 1 or identical ones as in Experiment 3, their end-states are never the same because the target names for the supraclass and subclass always form a semantic hierarchy for concept inclusion, but never do for percept inclusion. So the difficulty may simply be that children are unable to reach an end-state which requires a semantic comparison. To test this possibility, additional inclusion problems were constructed in the form of simple stories. One such story was prepared for each of the four categories, and each story ended with a request for a semantic comparison identical to the one used in the corresponding concept problem. For example, at the end of one story was a question that asked whether there had been more mothers in the story or more grown-ups.

The story problems were actually of two slightly different types. In *story–picture* problems the semantic inclusion question was accompanied by the same picture that had been prepared for the corresponding percept and concept problems. The story–picture problems were included for clarificational purposes that are described below with the appropriate data analyses. Of more immediate interest were the *story-only* problems, in which the inclusion task was purely verbal and was not accompanied by a picture. The solution to a story-only problem must be found by examining hierarchical relations in semantic memory, since the absence of pictures would prevent the application of any counting method. Nelson (1974) has reported that young children do in fact have hierarchical relations stored in semantic memory. It was therefore expected that children would succeed on the story-only problems.

Two previous studies (Winer, 1974; Wohlwill, 1968) have reported that children's performance was, in fact, better on purely verbal inclusion problems than on numerically identical problems that were presented with pictures. In these studies, however, the class relations in the inclusion problems did not always form a simple semantic hierarchy. An example of a simple hierarchy is (oranges + bananas = fruit). Instead, complex noun

phrases were occasionally used, as in the example (oranges + carrots = things to eat). In addition, the children in these studies were told the number of members comprising each subclass in a problem. As a result of these procedures, the children's solutions may have been based in part on numerical comparisons or on contextual analyses of the meanings of noun phrases, rather than on inferences drawn exclusively from relations implicit in a semantic hierarchy. In the present experiment, the story problems contained only simple semantic hierarchies, and no numerical information was given verbally.

To summarize, the design crossed four categories with four problem-types. The Piagetian model did not differentiate the problem-types, assigning the same logical structure to all of them. The problem processing model did differentiate them, making two important predictions. (a) *Start*-states were identical for a concept problem and its corresponding percept problem, and the concept problem was predicted to be more difficult. (b) *End*-states were identical for a concept problem and its related story-only problem, and again the concept problem was predicted to be more difficult. Confirmation of both these predictions would clearly imply that the difficulty of concept inclusion must reside neither in the properties of its start-state nor in those of its end-state. Rather, the difficulty would have to be in the mediating strategy.

Method

Subjects

The subjects were 24 boys and 24 girls who attended nursery schools in Ann Arbor, Michigan. None had participated in Experiments 1 or 2. One additional child was dropped from the study because he failed on a warm-up comparison between a group of three items and a separate group of two items.

Materials

Four pictures were prepared, 13 × 13 cm, each representing a different semantic category. In each picture there were three members of the pertinent category. The inclusion question for a picture always required a comparison of this supraclass of three items with an included subclass of two items. The three items in each picture formed a nonlinear triangular array. Like the example in Figure 2(c), each picture could accompany either a percept or a concept question.

The four categories, and the names associated with their supraclasses and subclasses, were: (a) *grown-ups* (two mothers with both a picnic basket and chair, one father with only a chair), (b) *animals* (two rabbits with both a carrot and pink spot, one turtle with only a spot), (c) *fruit* (two bananas which were both situated in a bowl and being cut by a knife, one orange which was only being cut by a knife), and (d) *children* (two boys with both a hat and an ice-cream cone, one girl with only an ice-cream cone).

Procedure

The instructions for the four types of inclusion problems are illustrated by those used for the cateogry of grown-ups.

Percept inclusion. Before showing the picture, the experimenter said, "This is a picture of some grown-ups. Some of the grown-ups have a picnic basket, and some of the grown-ups have a chair." Then the experimenter presented the picture and said, "Do you see all the picnic baskets? Do you see all the chairs? Are there more grown-ups who have a picnic basket or more grown-ups who have a chair?"

Concept inclusion. Comparable preparatory remarks came first: "This is a picture of some grown-ups. Some of the grown-ups are mothers, and some of the grown-ups are fathers." Then came the pictures and these remarks: "Do you see all the mothers? Do you see all the grown-ups? Are there more mothers or more grown-ups?"

Story-only inclusion. No picture was shown. Instead, the following story was told:

> This is a story about two girls. These two girls went to the park one day, to have a picnic. When they arrived at the park, they saw that there were lots of grown-ups in the park. Some of the grown-ups were mothers, and some of the grown-ups were fathers. Now do you remember the two girls? Well, one of the girls said, "I have an idea. Let's go around the park and say 'Hello' to all the mothers." So this girl wanted to way "Hello" to who, to all the _____? Then the other girl said, "I have a different idea. Let's say 'Hello' to all the grown-ups." So this girl wanted to say "Hello" to who, to all the _____? Now which girl do you think wanted to say "Hello" more times, the one who wanted to say "Hello" to all the mothers or the one who wanted to say "Hello" to all the grown-ups, to all the mothers or to all the grown-ups?

At the two points in the story indicated by a blank line, the child was required to supply the correct answer. If necessary, preceding portions of the story were repeated or clarified, until the child could answer correctly. This procedure was intended to ensure that the semantic information relevant to the inclusion question had been stored in memory by the child, and was accessible.

Story–picture inclusion. Exactly the same story was told, but just preceding the inclusion question, the associated picture was presented with these instructions: "Now I'm going to show you who was in the park. Do you see all the mothers? Do you see all the grown-ups?" Then the inclusion question was asked with the same phrasing used in the story-only problem.

Each child was given only four inclusion problems, with each of the four categories and each of the four problem-types appearing just once for that child. There are 24 distinct ways in which the four categories could be matched to the four problem-types. Two children, one boy and one girl, were assigned randomly to each distinct matching.

It was necessary to counterbalance, across the children, both the presentation order of the categories and the necessarily related presentation order of the problem-types. An algorithm was employed which yielded the following scheme for the groups of 24 children of each sex. The categories appeared six times in each of four distinct orders which formed a Latin square, and the problem-types, which had previously been assigned to these categories, appeared just one time in each of the 24 distinct orders that were possible. Finally, for two categories the subclass was mentioned first for all children, and for the other two the supraclass was always mentioned first.

Results and Discussion

The four categories and four problem-types generated a contingency table with 16 cells. In Table 2 are shown the proportions of children assigned to each of these cells who answered the corresponding inclusion question correctly. When similar contingency tables were prepared separately for boys and for girls, the mean sex difference over the 16 cells was zero, and no difference between marginal proportions was greater than .13. Consequently, the data of boys and girls were combined in Table 2.

TABLE 2

PROPORTIONS OF CHILDREN ANSWERING CORRECTLY BY PROBLEM-TYPE AND CATEGORY[a]

Category	Problem-type			
	Percept	Concept	Story-only	Story–picture
Grown-ups	.75	.42	.75	.75
Animals	.42	.08	.58	.58
Fruit	.58	.33	.58	.33
Children	.67	.08	.75	.50
Marginal proportion	.60	.23	.67	.54
.95 limits	(.46, .72)	(.14, .36)	(.53, .79)	(.41, .67)

[a] Data from a given child appeared just once in each row and once in each column. Consequently, $N = 12$ (six boys and six girls) for each cell proportion, and $N = 48$ for marginal proportions.

Reading across each row in the table, it may be seen that the rank order of difficulty for problem-types was reasonably consistent from one category to another.

A more rigorous test of this consistency was made by subjecting the 16 entries in Table 2 to an analysis of variance. Problem-types were a fixed effect in this analysis, and categories (not subjects) were the random component.[3] The effect of problem-types was significant, $F(3,9) = 11.03$, $p < .01$. Because the denominator for this F-ratio was the interaction between problem-types and categories, this finding supported the claim that here and probably also in Experiment 1, the differences among problem-types were generalizable across categories.

The comparative difficulty of particular problem-types was investigated more precisely than in Table 2 by means of the pairwise comparisons shown in Table 3. Each two-way classification in Table 3 is balanced in the sense illustrated by the following example. Exactly as many children had, e.g., the category of fruit for percept inclusion and the category of animals for concept inclusion, as had the reverse.

Part (a) of Table 3 strongly confirms the prediction that the concept problems would be more difficult than the percept problems, $\chi^2 (1) = 13.5$, $p < .01$ (McNemar's test of symmetry). It could be argued, however, that this difference in difficulty only indicated that the children guessed randomly on percept problems, which they may have found confusing, while giving systematically false answers to the concept problems.

[3] There was a slight but unavoidable dependency among the errors in this analysis, since any given subject contributed to four of the 16 proportions. This dependency was minimized, however, by the counterbalancing scheme which assured that within each sex, no two subjects appeared together in the same cell more than once.

TABLE 3

Two-Way Classifications for Pairs of Problem-Types[a]

Pair of problem-types			
Percept (P) Concept (C)	Story-only (SO) Concept (C)	Story-only (SO) Story–picture (SP)	Story–picture (SP) Concept (C)

	P+	P−		SO+	SO−		SO+	SO−		SP+	SP−
C+	8	3	C+	5	6	SP+	21	5	C+	5	6
C−	21	16	C−	27	10	SP−	11	11	C−	21	16
	(a)			(b)			(c)			(d)	

[a] Each of (a) through (d) is a two-way classification of children giving correct (+) and erroneous (−) answers for two problem-types. Total $N = 48$ children for each two-way classification.

Returning to Table 2, presented earlier, it may be seen that the confidence interval of the marginal proportion for percept inclusion did not, in fact, reject the chance level of performance (.50).

Additional data were available, however, which caused the guessing hypothesis to be rejected. These additional data came from the responses of the no-pointing subjects in Experiment 1 to the first percept problem only. The first percept problem under the no-pointing treatment in the earlier experiment was methodologically comparable to the percept problems in Experiment 3. Five categories were thus available. The corresponding proportions were .75 ($N = 24$) for the category of houses in Experiment 1 and the four values in Table 2 (each $N = 12$) for the categories in Experiment 3. Only one of these values is less than .50. Because the Ns varied, the categories were not combined. Instead, the five proportions were treated as data for a sample of categories, and were found to reject the hypothesis that the grand mean for all possible categories is not greater than .50, $t(4) = 2.14$, $p < .05$ one-tailed. The comparatively high level of success on percept problems was thus not an artifact of guessing, but, more likely, a systematic consequence of the counting strategy used by the children on these problems.

Returning to Table 3, part (b) may be seen to provide clear confirmation of the prediction that the story-only problems would be easier than the concept problems, $\chi^2(1) = 13.4$, $p < .01$. Moreover, the earlier Table 2 shows that more than half of the children succeeded on the story-only problem for each category. The marginal proportion in that table eliminates the guessing hypothesis. These findings imply that the children's success on the story-only problems followed from their use of a semantic strategy based on the information implicit in simple hierarchies. Conversely, the children's failure on the concept problems could not have been caused by

misunderstanding of semantic hierarchies, since the same semantic categories were used for both concept and story inclusion.

If this interpretation is correct, then why did the children not use the same semantic strategy to find correct solutions for the concept problems, which employed the same semantic categories? One explanation is that when given the opportunity to use either the semantic strategy leading to a correct solution or a counting strategy leading to an erroneous one, the children simply made the mistake of preferring to count. If so, then the children should also have preferred counting over semantic inference on the story–picture problems, which permitted both strategies. These problems should then have caused poorer performance than the story-only problems, which permitted only the semantic solution. But part (c) of Table 3 shows that the two types of story problems did not differ reliably, χ^2 (1) = 2.25.

An alternative explanation is that on concept problems the children never even recognized the possibility of a purely semantic solution. This hypothesis is supported by part (d) of Table 3, which shows that performance was reliably better on story–picture than on concept problems, χ^2 (1) = 8.33, $p < .01$. Both of these problem-types supplied semantic as well as pictorial information in the presentation of the inclusion task. Their semantic end-states were also the same. So the difference between them must have been strategic. Apparently, the story procedure made the pertinent hierarchies in semantic memory more accessible, or their relevance more noticeable (cf. Winer, 1974).

IMPLICATIONS

The main implication of these findings is that inclusion errors in young children may be more precisely represented as the outcome of problem-solving strategies than as a reflection of logical deficits. Methodologically, the success of the problem processing model suggests that a similar model might prove useful in other studies.

Methodology

The experimentally coordinated manipulation of start-states and end-states appears to be particularly effective as a method for investigating problem-solving strategies. It could provide researchers of cognitive development with a useful complement to the commonly used paradigm of training studies. Use of the method depends upon a precise analysis of two factors: (a) the sources of relevant information external to the subject, which constitute the initial conditions of the task, and (b) the minimal requirements for a solution, which constitute the terminal conditions of the task. The initial and terminal conditions are both under experimental control. Systematic manipulation of both factors within a single ex-

periment may help to elucidate the nature of the internal resources which constitute the subject's problem-solving techniques. In essence, the suggested method is an extension to experimental design of an analysis originally applied by Simon (1969) to the methodology of artificial intelligence.

The methods used in the present study were comparable in another way to those of artificial intelligence. Although the problem processing model was not actually implemented on a computer, it could have been. Using methods proposed by Nilsson (1971), a more technical version of the model was developed than has been presented here. The precision demanded by this technical enterprise was instrumental in suggesting the hypotheses that were subsequently supported by the experiments. Indeed, even the less detailed model presented here may suggest additional hypothesis which could be tested. Arguments have often been advanced for the potential value of technical models in the study of cognition and cognitive development (e.g., Klahr, 1973; Reitman, 1965). The positive results of the present study give support to these arguments.

Developmental Theory

The study also has implications for development theory. In particular, it suggests three possible reasons for the difficulty of class inclusion.

First, although preschool children are able to use an appropriate semantic strategy, they do so only when its relevance is made more noticeable by elaboration of the verbal context in which the inclusion question is embedded. In this respect, children's inclusion errors are analogous to their failure, in memory tasks, to spontaneously use mnemonic strategies which are demonstrably within their competence (Hagen & Kingsley, 1969; Moely, Olson, Halwes, & Flavell, 1969). As in memory development, one developmental aspect of class inclusion may thus be the acquisition of skill in thoroughly searching a problem-space for possible solution strategies.

Second, the exigencies of learning to count appear to make concept inclusion difficult by predisposing children to use a counting strategy that forbids the double enumeration of patterns. What might change with development is the likelihood that the child will double-count by employing the SCAN operator twice, rather than only once. A process that could explain this change is simply the automation of counting skills. As the child becomes more experienced and proficient in counting, an application of SCAN may require less careful monitoring. The portion of the child's finite problem-solving capacity which is freed in this way might then be available to organize an additional SCAN. Examining the effects of preparatory activities, such as the pointing sequence in Experiment 1, is one way to study children's capacities for organizing a strategy and monitoring its execution.

Experimental study of such activities, with more careful controls, may be a fruitful area for additional research.

Finally, it may be that with age children become increasingly skillful in the interactive exchange of information between the two cognitive systems whose respective functions are semantic analysis and problem-solving. As they were described earlier, these two systems operated in a fixed order that began with grammatical analysis, proceeded to semantic analysis, and ended with problem-solving. There is no real necessity for this order to be so rigid; in fact, greater flexibility could be quite advantageous. For example, when children solve the inclusion problem by the semantic strategy (without counting), they must use semantic information in the service of the problem-solving system. Older children may use this approach with comparative ease, but the younger children in this study and elsewhere (Winer, 1974; Wohlwill, 1968) used it only with the inducement of a semantically elaborated statement of the problem.

Another example, ironically, might be the solution of an inclusion problem by a clever use of Strategy I. In the familiar case of the dogs and the animals, one might successfully restrict SCAN to a single application by using the pattern P_1 (four-legged) to count the animals and the pattern P_2 (has-a-long-snout) to count the dogs. To succeed in this way, however, a child would presumably have to decide upon Strategy I, but somehow hold it in reserve while searching a cognitive dictionary for the names of both a common feature shared by dogs and animals, and a distinct feature unique to dogs. This approach is clearly not used by preschool children, and it may or may not be used by older children. The suggestion here is that if older children use Strategy I in this more sophisticated way, it may be because they have acquired greater skill in coordinating reciprocal exchanges of information between the semantic and problem-solving systems.

To summarize, there are reasonably good indications that young children's errors on class inclusion follow from their inflexible use of Strategy I. Whether older children succeed on similar problems because they use Strategy II, a clever form of Strategy I, or some other approach, remains unknown. Nevertheless, it is possible to suggest three developmental processes which may contribute to children's eventual acquisition of competence in this domain. These processes are recognition of relevant strategies, automation of counting, and growth in the capability for interactive communication among cognitive systems. It remains for future research to determine the validity of each process as a dimension of cognitive growth and its specific importance to class inclusion.

REFERENCES

Flavell, J. H. *The developmental psychology of Jean Piaget.* New York: Van Nostrand Reinhold, 1963.

Gelman, R. The nature and development of early number concepts. In H. W. Reese (Ed.),

Advances in child development and behavior. New York: Academic Press, 1972. Vol. 7.

Hayes, J. R. The child's conception of the experimenter. In S. Farnham-Diggory(Ed.), *Information processing in children.* New York: Academic Press, 1972.

Inhelder, B., & Piaget, J. *The early growth of logic in the child.* New York: Harper, 1964.

Inhelder, B., & Sinclair, H. Learning cognitive structures. In P. H. Mussen, J. Langer, & M. Covington (Eds.), *Trends and issues in developmental psychology.* New York: Holt, Rinehart & Winston, 1969.

Isen, A. M., Riley, C. A., Tucker, T., & Trabasso, T. *The facilitation of class-inclusion by use of multiple comparisons and two-class perceptual displays.* Paper read at the convention of the Society for Research in Child Development, 1975.

Kingsley, P. R., & Hagen, J. W. Induced versus spontaneous rehearsal in short-term memory in nursery school children. *Developmental Psychology,* 1969, **1,** 40–46.

Klahr, D. An information processing approach to the study of cognitive development. In A. D. Pick (Ed.), *Minnesota symposia on child psychology.* Minneapolis: University of Minnesota, 1973. Vol. 7.

Klahr, D., & Wallace, J. G. Class inclusion processes. In S. Farnham-Diggory (Ed.), *Information processing in children.* New York: Academic Press, 1972. (a)

Klahr, D., & Wallace, J. G. Reply to Hayes: On the value of theoretical precision. In S. Farnham-Diggory (Ed.), *Information processing in children.* New York: Academic Press, 1972. (b)

Markman, E. The facilitation of part–whole comparisons by use of the collective noun "family." *Child Development,* 1973, **44,** 837–840.

Moely, B. E., Olson, F. A., Halwes, T. G., & Flavell, J. H. Production deficiency in young children's clustered recall, *Developmental Psychology,* 1969, **1,** 26–34.

Nelson, K. Variations in children's concepts by age and category. *Child Development,* 1974, **45,** 577–584.

Nilsson, N. J. *Problem-solving methods in artificial intelligence.* New York: McGraw-Hill, 1971.

Newell, A., & Simon, H. A. *Human problem solving.* Englewood-Cliffs, NJ: Prentice-Hall, 1972.

Piaget, J. *The child's conception of number.* New York: Humanities, 1952.

Piaget, J. Piaget's theory. In P. H. Mussen (Ed.), *Carmichael's manual of child psychology* New York: Wiley, 1970. 3rd ed., Vol. 1.

Potter, M. C., & Levy, E. I. Spatial enumeration without counting. *Child Development,* 1968, **39,** 265–272.

Reitman, W. *Cognition and thought.* New York: Wiley, 1965.

Schaeffer, B., Eggleston, V. H., & Scott, J. L. Number development in young children. *Cognitive Psychology,* 1974, **6,** 357–379.

Simon, H. A. *The sciences of the artificial.* Cambridge, MA: MIT Press, 1969.

Winer, G. A. An analysis of verbal facilitation of class-inclusion reasoning. *Child Development,* 1974, **45,** 224–227.

Winograd, T. Understanding natural language. *Cognitive Psychology,* 1972, **3,** 1–191.

Wohlwill, J. F. Responses to class-inclusion questions for verbally and pictorially presented items. *Child Development,* 1968, **39,** 449–465.