Preprint: To appear in *Cognitive Science, 18*, 87-122, 1994.

Representations in Distributed Cognitive Tasks

Jiajie Zhang
The Ohio State University
jiazhang@magnus.acs.ohio-state.edu

Donald A. Norman
Apple Computer, Inc.
dnorman@apple.com

ABSTRACT

In this paper we propose a theoretical framework of distributed representations and a methodology of representational analysis for the study of distributed cognitive tasks—tasks that require the processing of information distributed across the internal mind and the external environment. The basic principle of distributed representations is that the representational system of a distributed cognitive task is a set of internal and external representations, which together represent the abstract structure of the task. The basic strategy of representational analysis is to decompose the representation of a hierarchical task into its component levels so that the representational properties at each level can be independently examined. The theoretical framework and the methodology are used to analyze the hierarchical structure of the Tower of Hanoi problem. Based on this analysis, four experiments are designed to examine the representational properties of the Tower of Hanoi. Finally, the nature of external representations is discussed.

People behave in an information rich environment filled with natural and artificial objects extended across space and time. A wide variety of cognitive tasks, whether in everyday cognition, scientific practice, or professional life, require the processing of information distributed across the internal mind and the external environment. It is the interwoven processing of internal and external information that generates much of a person's intelligent behavior (e.g., Hutchins, 1990, in preparation; Norman, 1988, 1991, in press).

The traditional approach to cognition, however, often assumes that representations are exclusively in the mind (e.g., as propositions, schemas, productions, mental images, connectionist networks, etc.). External objects, if they have anything to do with cognition at all, are at most peripheral aids. For instance, written digits are usually considered as mere memory aids for calculation. Thus, because the traditional approach lacks a means of accommodating external representations in its own right, it sometimes has to postulate complex internal representations to account for the complexity of behavior, much of which, however, is merely a reflection of the complexity of the environment (e.g., Kirlik, 1989; Simon, 1981; Suchman, 1987).

This paper addresses the representational issues in *distributed cognitive tasks*—tasks that require the processing of information distributed across the internal mind and the external environ-

---

This research was supported by a grant to Donald Norman and Edwin Hutchins from the Ames Research Center of the National Aeronautics & Space Agency, Grant NCC 2-591 in the Aviation Safety/Automation Program, technical monitor, Everett Palmer. Additional support was provided by funds from the Apple Computer Company and the Digital Equipment Corporation to the Affiliates of Cognitive Science at UCSD.

We are very grateful to Ed Hutchins for many inspiring discussions throughout this study. We would also like to thank David Kirsh, Jean Mandler, Mark St. John, Hank Strub, and Kazuhiro Takahashi for many good comments during this study, James Greeno and Daniel Reisberg for thoughtful comments and suggestions on an earlier version of the manuscript, and Sandra Hong, Lin Riley, and Richard Warren for their assistance in the experiments.

Correspondence and requests for reprints should be sent to Jiajie Zhang, Department of Psychology, The Ohio State University, 1885 Neil Avenue, Columbus, OH 43210.

ment, focusing on three problems: (a) the distributed representation of information; (b) the interaction between internal and external representations; and (c) the nature of external representations. In the first part of this paper, we begin with an introduction to the *representational effect*, then we propose a theoretical framework of *distributed representations* and a methodology of *representational analysis* for the study of distributed cognitive tasks. In the second part, we illustrate the principles of distributed representations and the strategies of representational analysis through the analysis of the representational structure of the Tower of Hanoi problem. In the third part, we design four sets of Tower of Hanoi isomorphs for empirical investigations. In the last part, we summarize our major claims and discuss the general properties of distributed cognitive tasks.

PHENOMENON, THEORY, AND METHODOLOGY

The Representational Effect

The *representational effect* refers to the phenomenon that different isomorphic representations of a common formal structure can cause dramatically different cognitive behaviors. One obvious example is the representation of numbers (for cognitive analyses, see Nickerson, 1988; Norman, in press; Zhang, 1992; Zhang & Norman, 1993). We are all aware that Arabic numerals are more efficient than Roman numerals for multiplication (e.g., 73 × 27 is easier than LXXIII × XXVII), even though both types of numerals represent the same entities—numbers. The most dramatic case is probably the Copernican revolution, where the change from the geocentric representation of the solar system (Ptolemaic system) to the heliocentric representation (Copernican system) laid the foundation of modern science and fundamentally changed people's conception of the universe.

Psychological studies of the representational effect in problem solving and reasoning have focused on a few well-structured problems, including the Tower of Hanoi problem (e.g., Hayes & Simon, 1977; Kotovsky & Fallside, 1989; Kotovsky, Hayes & Simon, 1985; Simon & Hayes, 1976), the Chinese Ring puzzle (Kotovsky & Simon, 1990), the Hobbits-Orcs problem (e.g., Greeno, 1974; Jeffries, Polson, & Razran, 1977; Thomas, 1974), Wason's selection task (e.g., Cheng & Holyoak, 1985; Evans, 1983; Margolis, 1987; Wason, 1966; Wason & Johnson-Laird, 1972), and Wason's THOG problem (e.g., Griggs & Hewstead, 1982; O'Brien, Noveck, Davidson, Fisch, Lea, & Freitag, 1990; Wason & Johnson-Laird, 1969). The basic finding is that different representations of a problem can have dramatic impact on problem difficulty even if the formal structures are the same. One characteristic of these problems is that they all require the processing of both internal and external information. However, most of these studies either exclusively focused on internal representations or, when taking external representations into account, failed to separate them from internal representations.

In this paper, we argue that internal and external representations are two indispensable parts of the representational system of any distributed cognitive task. To study a distributed cognitive task, it is essential to decompose the representation of the task into its internal and external components so that the different functions of internal and external representations can be identified.

Distributed Representations

The basic principle of distributed representations is that the representational system of a distributed cognitive task can be considered as a set, with some members internal and some external. Internal representations are in the mind, as propositions, productions, schemas, mental images, connectionist networks, or other forms. External representations are in the world, as physical symbols (e.g., written symbols, beads of abacuses, etc.) or as external rules, constraints, or relations embedded in physical configurations (e.g., spatial relations of written digits, visual and spatial layouts of diagrams, physical constraints in abacuses, etc.). Generally, there are one or more internal and external representations involved in any distributed cognitive task.

Figure 1 shows the representational system of a task with two internal and two external representations. Each internal representation resides in a person's mind, and each external representation resides in an external medium. The internal representations form an internal representational space, and the external representations form an external representational space. These two spaces together form a distributed representational space, which is the representation of the abstract task space that describes the abstract structures and properties of the task.

We need to clarify our use of the term "representation". In our present study, the referent (the represented world) of a representation (the representing world) is an abstract structure. For simple tasks, a representation and its referent can be perceived by both theorists and task performers. For example, both theorists and task performers know that written numerals, whether they are Arabic or Roman, are the representations of abstract numbers. For complex tasks, however, a representation and its referent are usually only meaningful from the point of view of theorists. To a task performer, a representation does not represent anything: it is simply the medium (internal and/or external) on which the task performer performs the task. For example, for the Tower of Hanoi problem that we will consider later, the three problems in Figure 12 all represent the same abstract structure (i.e., problem space, see Figure 3) to a theorist. To a task performer, however, they are simply three different problems and they do not represent anything, though the task performer might notice some regularities across these problems. Our notion of representation is essential for our present studies of the representational properties of distributed cognitive tasks. By considering alternative representations of a common abstract structure, we can identify the factors that affect the cognitive behavior in distributed cognitive tasks.

Our current approach to distributed cognitive tasks demands: (a) the consideration of the internal and external representations of a distributed cognitive task as a representational system; (b) the explicit decomposition of the representational system into its internal and external components; and (c) the identification of the different functions of internal and external representations in cognition. The traditional approach to cognition is not appropriate for the study of distributed cognitive

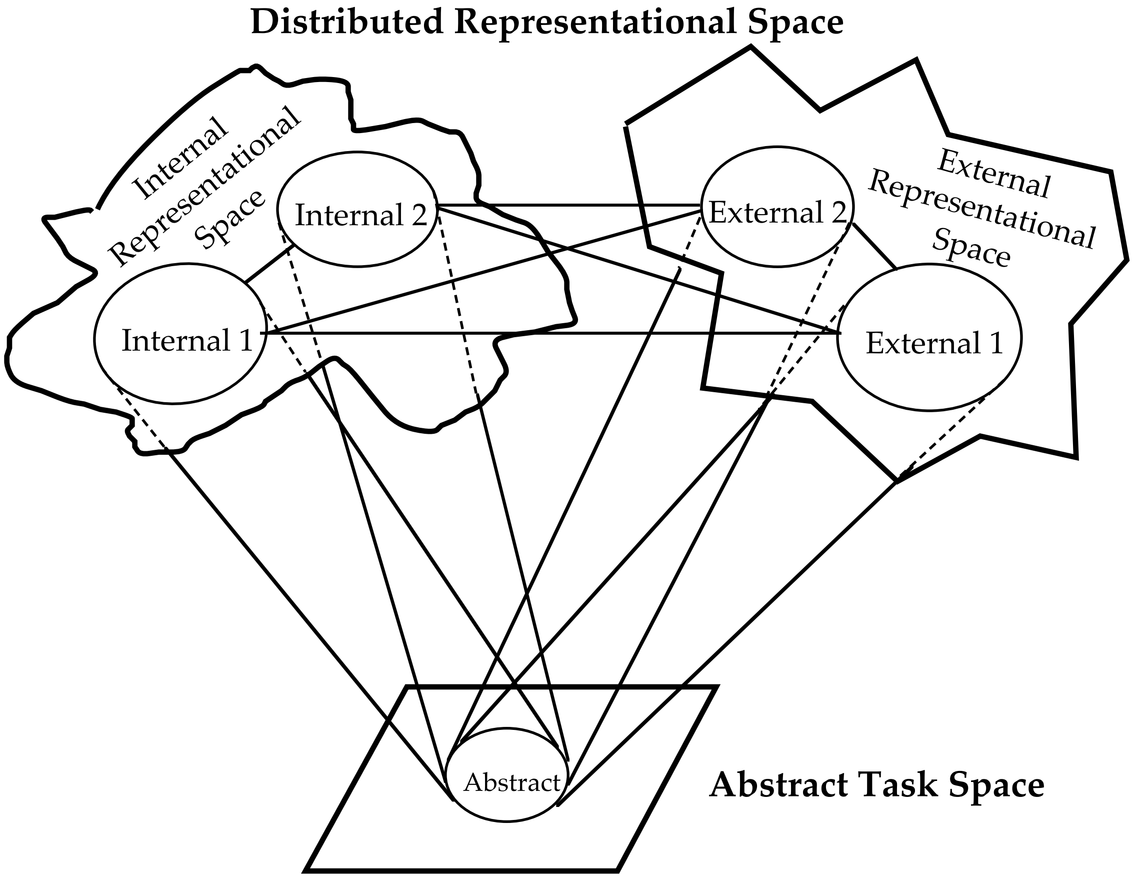

Figure 1. The theoretical framework of distributed representations. The internal representations form an internal representational space, and the external representations form an external representational space. The internal and external representational spaces together form a distributed representational space, which is the representation of the abstract task space.

tasks, because (a) it considers external representations as mere peripheral aids to cognition, and (b) it often mixes external representations with internal representations. Furthermore, the traditional approach often mistakenly equates a task's distributed representation that has both internal and external components to the task's internal representation. This confusion often leads one to postulate complex internal mechanisms to explain the complex structure of the wrongly identified internal representation, much of which is merely a reflection of the structure of the external representation.

Representational Analysis

Representational analysis is a methodology for the study of the representational effect in distributed cognitive tasks. It is based on hierarchical representations, isomorphic representations, and distributed representations.

Many distributed cognitive tasks have multi-level hierarchical representations (see Zhang, 1992, for a few examples). At each level of a task's hierarchical representation, there is an abstract structure that can be implemented by different isomorphic representations. For some levels, the isomorphic representations can be distributed representations. By decomposing the representation of a task into its component levels, we can identify the representational properties at each level that are responsible for a different aspect of the representational effect. It should be noted that different tasks usually have different hierarchical structures and have different representational properties at their component levels. The key of the methodology of representational analysis is the strategy of decomposing a task into its component levels and studying the representational properties at each level.

The methodology of representational analysis is fully illustrated in the analysis of the representational structure of the Tower of Hanoi problem in next section. It has also been used to study several real world problems, including numeration systems, relational information displays, and cockpit instrument displays (Zhang, 1992).

THE REPRESENTATIONAL STRUCTURE OF THE TOH

The Tower of Hanoi (henceforth TOH; see Figures 2 and 3) is a well-studied problem. Much of the research has focused on its isomorphs and problem representations (e.g., Hayes & Simon, 1977; Kotovsky & Fallside, 1989; Kotovsky, Hayes & Simon, 1985; Simon & Hayes, 1976). The basic finding is that different problem representations can have dramatic impact on problem difficulty even if the formal structures are the same. Many of these studies either explicitly or implicitly mentioned external representations. However, they did not consider internal and external representations as two indispensable parts of the representational system of the TOH, and they did not separate external representations from internal ones. Furthermore, the hierarchical structure of the TOH was not analyzed in these studies. In this section, we use the principles of distributed representations and the methodology of representational analysis described in last section to analyze the representational structure of the TOH in a systematic way.

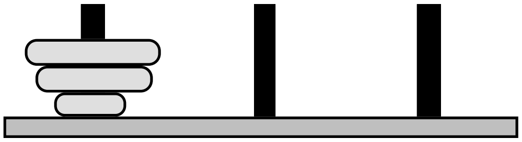

Figure 2. The standard Tower of Hanoi problem. The task is to move the three disks from one configuration to another, following two rules: (1) Only one disk can be transferred at a time; (2) A disk can only be transferred to a pole on which it will be the largest.

Rule Representations

The standard TOH[1] shown in Figure 2 has two internal rules: (1) only one disk can be transferred at a time; (2) a disk can only be transferred to a pole on which it will be the largest. These two rules have to be memorized. The TOH in Figure 2 also has one external rule: (3) only the largest disk on a pole can be transferred to another pole. In the representation shown in Figure 2, Rule 3 need not be stated explicitly because the physical structure of the disks and poles coupled with Rule 1 guarantee that it will be followed. But if the disks were not stacked on poles, explicit statement of Rule 3 would be necessary. In our studies we used four rules[2] :

---

[1] The disk sizes of the traditional TOH are the reverse of those shown in Figure 2: the largest disk is at the bottom and the smallest is at the top. The disk sizes have been reversed to make all experimental designs consistent. We call the size-reversed TOH with disks and poles in Figure 2 *our* standard TOH, because all other TOH isomorphs in our present study were derived from this version.

[2] Experiments 1A and 1B used all four rules. Experiments 2 and 3 only used Rules 1, 2, and 3, which are the three rules for the standard TOH.

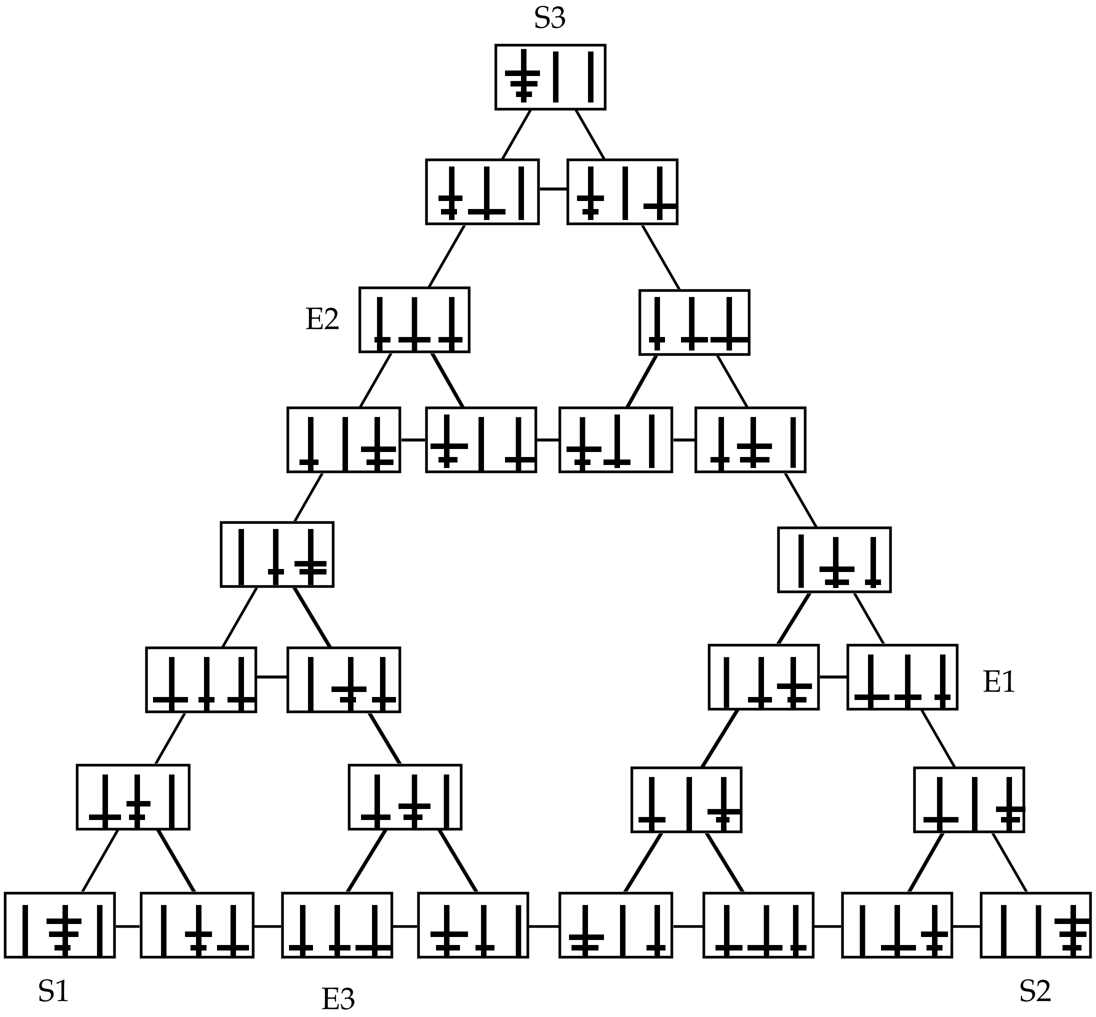

Figure 3. The problem space of the Tower of Hanoi problem. Each rectangle shows one of the 27 possible configurations (states) of the three disks on the three poles. The lines between the rectangles show the transformations from one state to another when the rules are followed. S1, S2, and S3 are three starting states, and E1, E2, and E3 are three ending states. They will be used later in the experiments.

Rule 1: Only one disk can be transferred at a time.
Rule 2: A disk can only be transferred to a pole on which it will be the largest.
Rule 3: Only the largest disk on a pole can be transferred to another pole.
Rule 4: The smallest disk and the largest disk can not be placed on a single pole unless the medium sized disk is also on that pole.

Any of these four rules can be either internal or external. Internal rules are memorized rules that are explicitly stated as written propositions in the instructions for experiments. External rules are not stated in any form in the instructions. They are the constraints that are embedded in or implied by physical configurations and can be perceived and followed without being explicitly formulated. Some external rules may sometimes depend on other internal rules and/or some background knowledge. For example, the external Rule 3 in the standard TOH depends on the internal Rule 1, because Rule 3 is a rule about the movement of a single object[3]. In addition, Rules 2 and 3, whether they are implemented as external or internal rules, require that only one hand can be used to move the objects and the objects can only be placed on the three poles (or plates, see Figures 11 & 12), because otherwise we can always use the spare hand or another place (e.g., table surface) as a temporary holder (equivalent to the addition of a fourth pole for the standard TOH) such that Rules 2 and 3 can be bypassed. These two extra requirements are embedded in the cover stories for the experiments (see Figure 9), and they are not counted as internal rules. Furthermore, Rule 2 in the *Waitress and Coffee* TOH (see Figure 11) is an external rule that needs some cultural knowledge: spilling coffee in front of a customer is not a good behavior for a waitress or a waiter (see Experiment 1B for details). Even though some external rules are not fully independent and not truly external, we still call them external rules in the sense that they are not stated and memorized but nevertheless functionally equivalent to those that are explicitly stated and memorized.

In the experiments that follow, we varied the number of external rules. In Condition *I1* (Figure 11A, in Experiment 1B) and Condition *I123* (Figure 12A, in Experiment 2), no rule is external. In Conditions *I1-E3* (Figure 11B, in Experiment 1B) and *I12-E3* (Figure 12B, in Experiment 2), Rule 3 is external. In Condition *I1-E23* (Figure 11C, in Experiment 1B; and Figure 12C, in Experiment 2), Rules 2 and 3 are external. In Condition *I1-E234* (Figure 11D, in Experiment 1B), Rules 2, 3, and 4 are external. Detailed explanations of these external rules are described in Experiment 1B and Experiment 2.

Problem Space Structures

A problem space of the TOH is composed of all possible states and moves constrained by the rules. Figures 4A-E show the problem spaces constrained by Rules 1, 1+2, 1+3, 1+2+3, and 1+2+3+4, respectively (see footnote 3) They are derived from the problem space shown in Figure 3. The rectangles (problem states) in Figure 3 are not shown in Figure 4 for the reason of clarity. Figures 4A-D have the same 27 problem states as in Figure 3. Figure 4E only has 21 problem states (shown by the dots), which are the outer 21 rectangles in Figure 3. Lines with arrows are unidirectional; lines without arrows are bidirectional. One important point is that these five spaces can represent internal, external, or mixed problem spaces, depending upon how the rules constructing them are distributed across internal and external representations. A problem space constructed by external rules is an external problem space, one constructed by internal rules is an internal problem space, and one constructed by a mixture of internal and external rules is a mixed problem space. Figure 4B is the internal problem space of the standard TOH because Rules 1 and 2 are internal. If the physical constraints imposed by the disks themselves are such that only one can be moved at a time (e.g., the disks are large or heavy), then Figure 4C is the external problem space of the standard TOH because under this cir-

---

[3] The four rules for the Tower of Hanoi are not fully orthogonal. Rules 2, 3, and 4 are orthogonal to one another and Rule 4 is orthogonal to Rule 1. However, Rules 2 and 3 are not orthogonal to Rule 1, because Rule 1 is the prerequisite of Rules 2 and 3.

cumstance Rules 1 and 3 are both external. These two spaces form the distributed problem space of the standard TOH (Figure 5), and their conjunction forms the abstract problem space, which is equivalent to the problem space shown in Figure 4D.

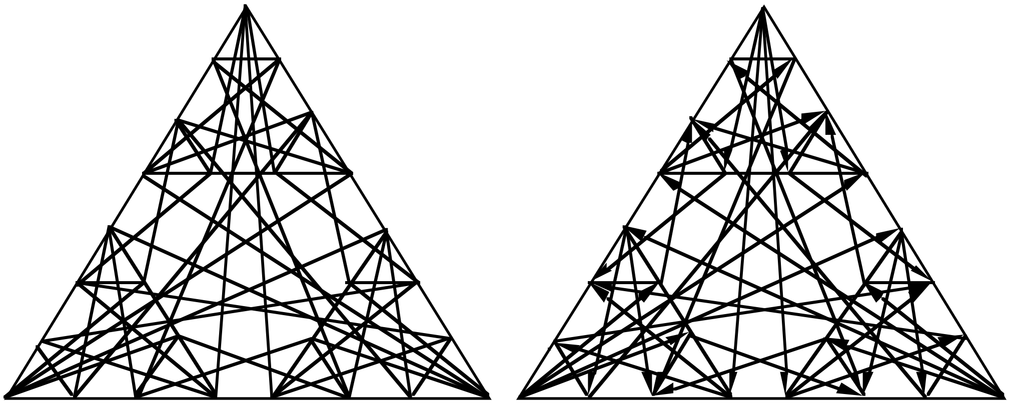

(A) Rule 1 **(B) Rules 1+2**

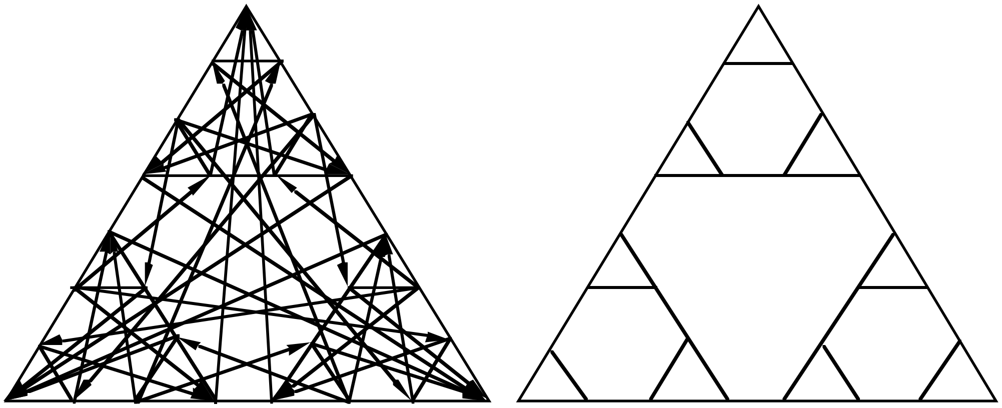

(C) Rules 1+3 **(D) Rules 1+2+3**

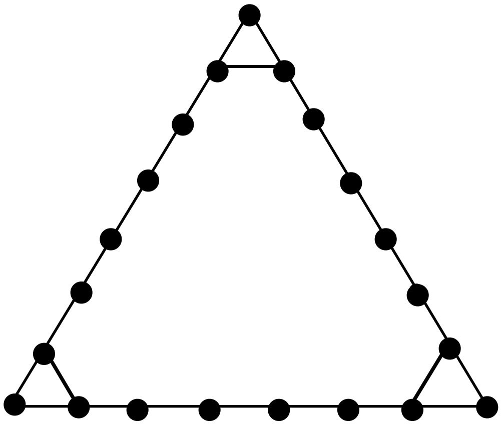

(E) Rules 1+2+3+4

Figure 4. Problem spaces constructed by five different sets of rules. They are derived from the problem space shown in Figure 3. Lines with arrows are uni-directional. Lines without arrows are bi-directional. The rectangles (problem states) are not shown here for the reason of clarity. (A)-(D) have the same 27 problem states as in Figure 3. (E) only has 21 problem states (shown by the dots), which are the outer 21 rectangles in Figure 3.

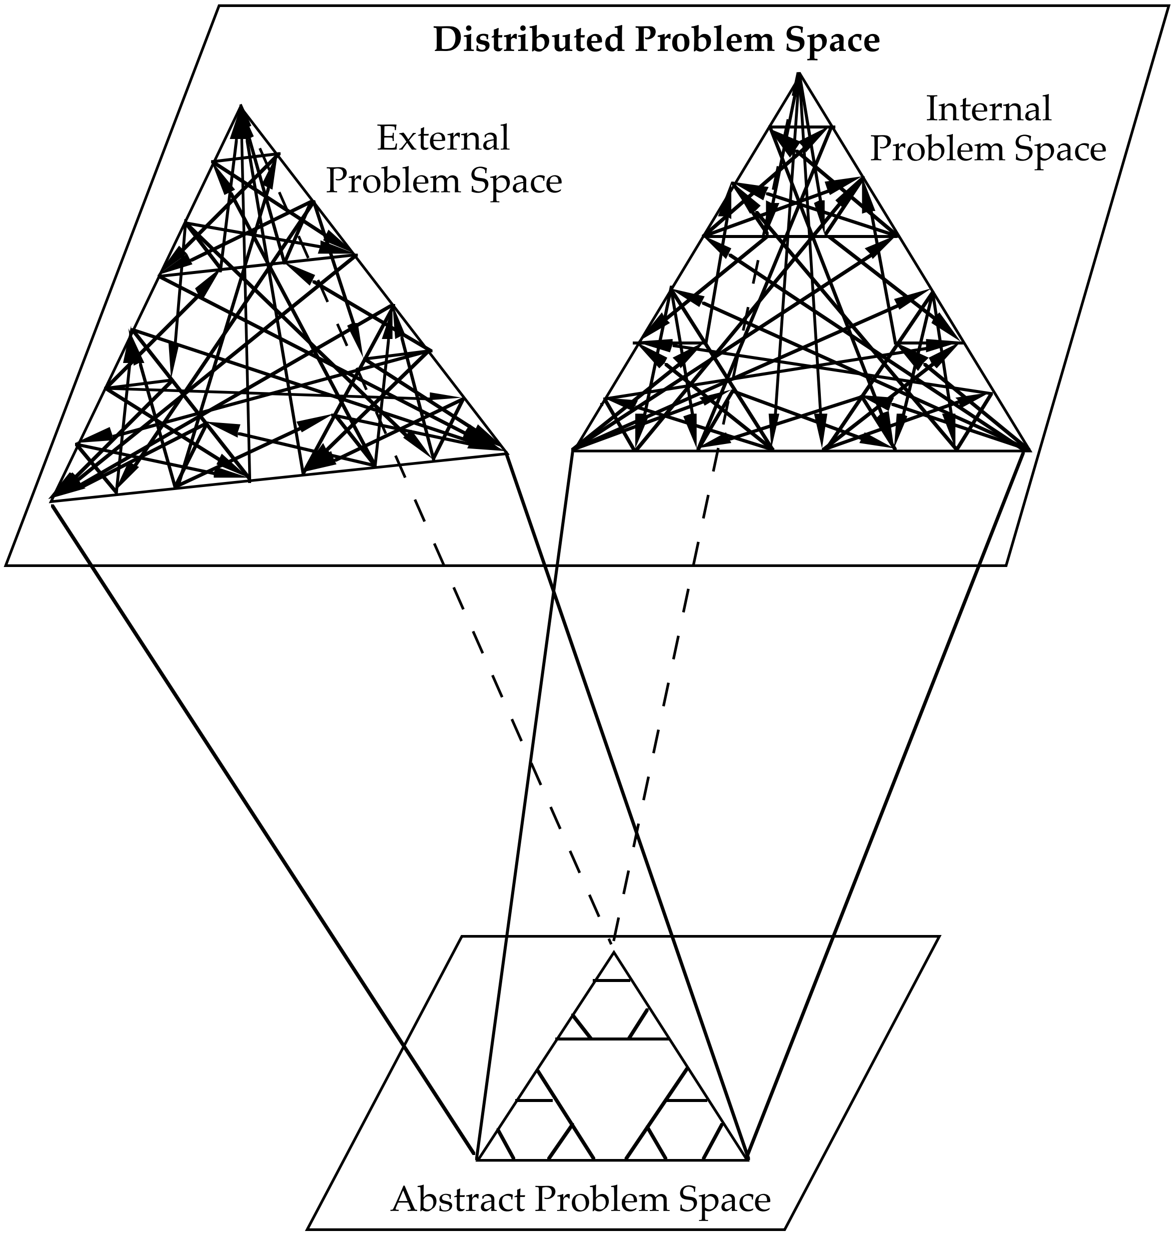

Figure 5. The distributed representation of the TOH. The distributed problem space is composed of the internal and the external problem spaces. The abstract problem space is the conjunction of the internal and the external problem spaces.

Dimensional Representations

The standard TOH (Figure 2) has three disks, which possess two dimensions of properties. The first is the ordinal dimension represented by the sizes of the three disks, which has three levels: large > medium > small. The ordinal information on this dimension is required by the rules. The values on this dimension are constants, that is, the sizes of the disks are fixed. The second is the nominal dimension represented by the locations of the disks, which also has three levels: left, middle, right. This dimension is called nominal dimension because only the categorical information of the three locations is needed for the rules. (For a description of psychological scales, see Stevens, 1946, 1951.) The values on this dimension are variables, that is, a disk can be at any of the three locations. The three poles in the standard TOH shown in Figure 2 are not essential: they are only used to construct the external Rule 3.

The ordinal and nominal dimensions do not have to be represented by sizes and locations of the disks, as in the standard TOH. They can be represented by any properties of any objects. For example, Figure 6 shows two TOH isomorphs whose ordinal and nominal dimensions are represented by different properties. In Figure 6A, the ordinal dimension is represented by the sizes of the balls (large > medium > small), and the nominal dimension is represented by the locations of the balls (left, middle, and right). In Figure 6B, the ordinal dimension is represented by the locations of the triangular cylinders (top > middle > bottom). Each cylinder has three different colors (R = red, G = green, Y = yellow) on its three sides, which represent the nominal dimension. The sizes of the balls in Figure 6A are mapped to the locations of the cylinders in Figure 6B, and the locations in Figure 6A are mapped to the colors of the cylinder sides in Figure 6B. For example, moving the largest ball from the right location to the left location in Figure 6A corresponds to rotating the top cylinder from the yellow side to the red side in Figure 6B. More examples of the dimensional representations of the TOH are shown in Figure 14.

The ordinal and nominal dimensions of the TOH can be represented either internally or externally. If the ordinal dimension is represented by physical properties that are on an ordinal (or higher) scale (e.g., sizes or ordered locations), it can be represented externally. For example, the sizes of the balls in Figure 6A are an external representation of the ordinal dimension because the ordinal relation is embedded in the physical properties (sizes) of the balls. If the ordinal dimension is represented by physical properties that are on a nominal scale, it must be represented internally. For example, if we use colors to represent the ordinal dimension, we must first arbitrarily assign an order to them (e.g., red > green > yellow) and then internalize this ordinal relation (see Figures 14G, 14H, 14I for examples). The nominal dimension can be represented externally by physical properties that are either on a nominal or on an ordinal scale, because an ordinal scale is also a nominal scale (but not vice versa).

Furthermore, the ordinal and nominal dimensions of the TOH can be represented by either visual properties (such as size, color, texture, and shape) or spatial properties (such as location). The separation of spatial properties from visual properties is significant because many studies have found that there are important anatomical and

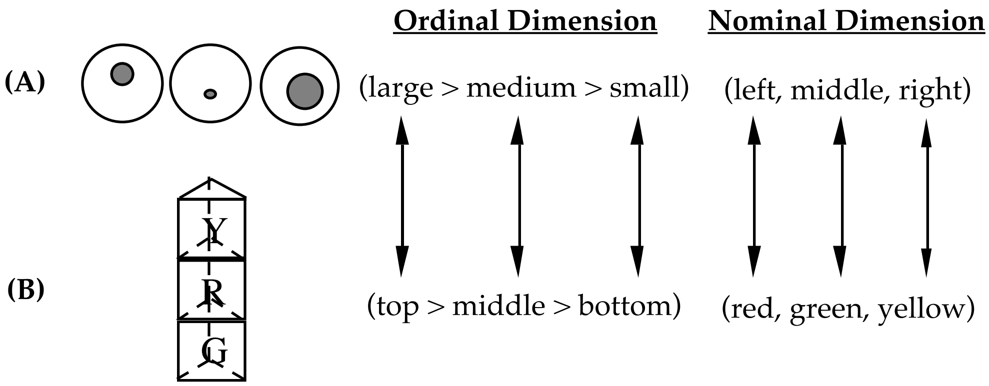

Figure 6. The mapping between dimensions of two TOH isomorphs. The three triangular cylinders in (B) can be rotated independently around their axis. The three sides of each cylinder have three different colors (R = red, G = green, Y = yellow). The sizes of the balls (ordinal dimension) in (A) are mapped to the locations of the cylinders (ordinal dimension) in (B), and the locations of the balls (nominal dimension) in (A) are mapped to the colors of cylinder sides (nominal dimension) in (B). For example, moving the largest ball from the right location to the left location in (A) corresponds to rotating the top cylinder from the yellow side to the red side.

functional differences between visual and spatial representations. Mishkin, Ungerleider, & Macko (1983) found that there are two separate systems for visual and spatial information processing: visual information processing follows a projection from the occipital to the temporal cortex, and spatial information processing follows a projection from the occipital to the parietal cortex. Goldman-Rakic (1992) showed that among the prefrontal neurons responsible for working memory, some only respond to spatial locations while others only respond to visual properties. Treisman & Gelade (1980) found that the visual and spatial properties of a stimulus are processed differently in perceptual tasks and that location is required for focused attention to conjunctive stimuli. The differences between visual and spatial representations may have important implications for complex cognitive tasks such as problem solving and reasoning. One of the purposes of Experiment 3 in the present study is to examine the different functions of visual and spatial dimensions in the TOH task.

The Abstract Structure of the TOH

The TOH has a goal, a set of rules, two dimensions, and three objects. Figure 7 shows the abstract structure of the TOH (see also Figure 8, which will be explained in the following section). The goal is defined by an initial and a final problem state. The rules can vary in numbers, that is, any subset of the four rules or all of them can be chosen for a given version of the TOH. The rules in a given set can be distributed across internal and external representations in different ways. The two dimensions can be represented by different physical properties. O_1, O_2, and O_3 are the three levels of the ordinal dimension, and N_1, N_2, and N_3 are the three levels of the nominal dimension. An object OBJ_i is described as $OBJ_i = (O_i, N_l)$, which can be at three different levels on the nominal dimension: (O_i, N_1), (O_i, N_2), (O_i, N_3). The three objects can be represented by different physical entities.

- Two property dimensions.
  - Ordinal dimension. 3 levels: $O_1 > O_2 > O_3$.
  - Nominal dimension. 3 levels: N_1, N_2, N_3.
- Object: $OBJ_i = (O_i, N_l)$. $i = 1, 2, 3; l = 1, 2, 3$.
- Problem state: $S(l, m, n) = ((O_1, N_l), (O_2, N_m), (O_3, N_n))$. $l, m, n = 1, 2, 3$.
- Operation: $OP(O_i, N_l) = (O_i, N_m)$. $l \neq m$.
- Rules:
  - 1: OP is a unary operator.
  - 2: When $OBJ_j = (O_j, N_m)$, $OP(O_i, N_l) = (O_i, N_m)$ is true if $O_i > O_j$.
  - 3: When $OBJ_i = (O_i, N_l)$ & $OBJ_j = (O_j, N_l)$, $OP(O_i, N_l)$ is true if $O_i > O_j$.
  - 4. $OBJ_1 = (O_1, N_m)$ & $OBJ_3 = (O_3, N_m)$ is true if $OBJ_2 = (O_2, N_m)$
- Goal: $S(l', m', n') \rightarrow S(l'', m'', n'')$.

Figure 7. The abstract structure of the TOH. O_1, O_2, and O_3 are the three levels of the ordinal dimension, and N_1, N_2, and N_3 are the three levels of the nominal dimension. An object OBJ_i is described as $OBJ_i = (O_i, N_l)$, which can be at three different levels on the nominal dimension: (O_i, N_1), (O_i, N_2),(O_i, N_3). All four rules or any subset of them can be chosen for a given problem.

The Hierarchical Representation of the Tower Of Hanoi

From the abstract structure of the TOH (Figure 7), we can identify four levels of representations: problem space structures, rule representations, dimensional representations, and object representations. At each level, there is an abstract structure that can be implemented by different representations. The different representations at each level are isomorphic in the sense that they all share the same abstract structure at that particular level. The hierarchical representation of the TOH is summarized in Figure 8, and the details are described as follows.

The Level of Problem Space Structures

The problem space of a problem is composed of all possible states and all moves constrained by the rules of the problem. The TOH is not bound to the three rules of the standard TOH. Different sets of rules construct different problem spaces. Figure 4 shows problem spaces constructed by different sets of rules: Rule 1, Rules 1+2, Rules 1+3, Rules 1+2+3, and Rules 1+2+3+4. Problems that have different problem spaces are isomorphic at this level if their goals (initial and final states) are the same. At this level, the formal properties of problem space structures (such as the connections between problem states and the total number of problem states) are the major factors that affect problem solving behavior.

In Figure 4, when the number of rules increases, problem space structures become more constrained and problem solving behavior might change accordingly. The effect of the structural change of a problem space on problem solving behavior might depend on the nature of the rules (whether internal or external). Experiments 1A and 1B were designed to examine these two factors. In

Experiment 1A, the change of the problem space structure was caused by internal rules, while in Experiment 1B, it was caused by external rules.

The Level of Rule Representations

Given a set of rules, say, Rules 1, 2, and 3, the abstract problem space constructed by them is fixed. However, the rules can be distributed across internal and external representations in different ways. For example, Rule 1 may be internal and Rules 2 and 3 external, or Rules 1 and 2 internal and Rule 3 external, or all three rules internal (see Figure 12). Problems that have the same set of rules but different distributions of rules are isomorphic to each other at this level in the sense that they all have the same abstract problem space. Experiment 2 was designed to study the effect of the distributed representation of rules on problem solving behavior.

The Level of Dimensional Representations

The ordinal and the nominal dimensions of the TOH do not have to be represented by sizes and locations as in the standard TOH. They can be represented by any properties (see Figure 6 and Figure 14 for examples). Problems whose ordinal and nominal dimensions are represented by different properties are isomorphic to each other at this level in the sense that the different properties represent the same ordinal information on the ordinal dimension and the same nominal information on the nominal dimension.

There are two factors at this level that might affect the problem solving behavior of the TOH. The first factor is whether the ordinal dimension of the TOH is represented internally (e.g., by colors) or externally (e.g., by sizes and locations). The second factor is whether the ordinal and nominal dimensions of the TOH are represented by visual properties (e.g., size and color) or spatial properties (e.g., location). The different representational properties of internal and external dimensions and those of visual and spatial dimensions may produce different processing strategies for the problem solving of the TOH. Experiment 3 was designed to examine these two factors.

The Level of Object Representations

The problems that are isomorphic at all previous three levels can be isomorphic at still another level—the level of object representations. At this level, different objects can be used. For example, when the ordinal dimension is represented by sizes, we can use different sized disks or different sized balls. For the same reason, when the nominal dimension is represented by colors, we can use different sets of colors to represent the nominal dimension, e.g., red, green, and yellow, or purple, blue, and pink. The representational properties at this level usually do not have as large effects as those at the first three levels.

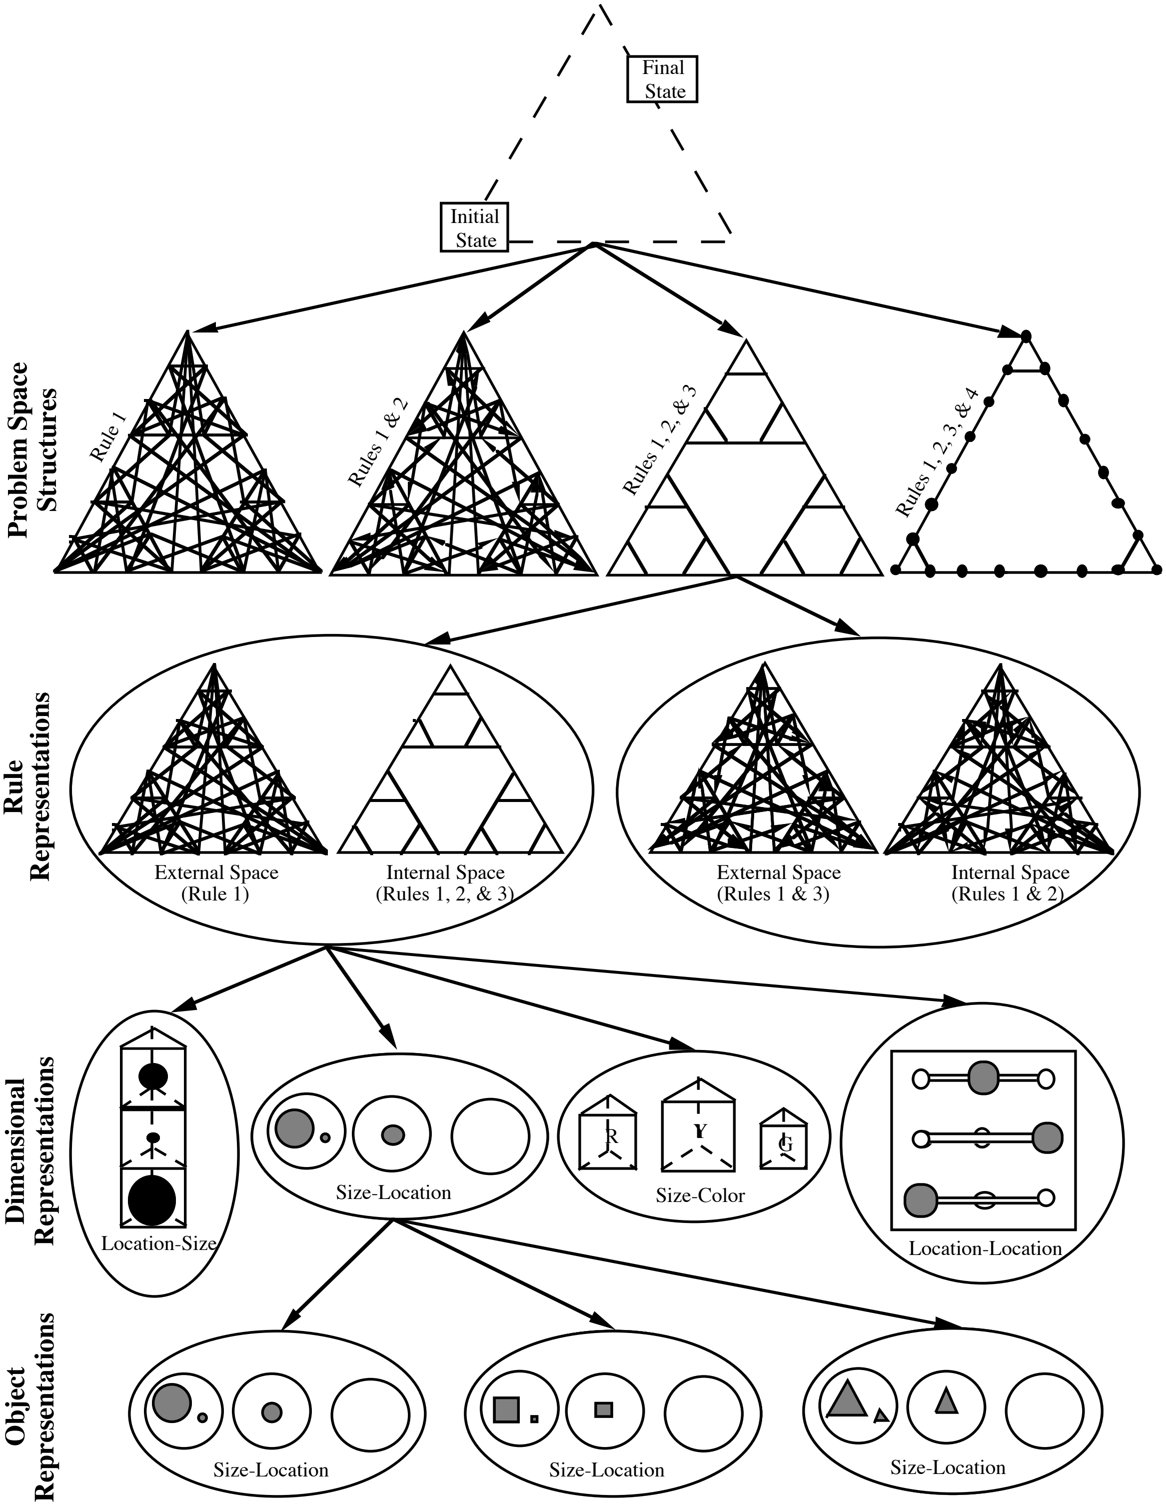

Figure 8. The hierarchical representation of the TOH. At the level of problem space structures, different sets of rules construct different problem spaces. At the level of rule representations, the same set of rules can be distributed across internal and external representations in different ways. At the level of dimensional representations, the ordinal and nominal dimensions can be represented by different properties. At the level of object representations, different objects can be used.

EXPERIMENT 1A: PROBLEM SPACE STRUCTURES (INTERNAL)

In Experiments 1A and 1B, we study isomorphic representations at the level of problem space structures. Isomorphic representations at this level have the same goal but different problem spaces constructed by different sets of rules. Figure 4 shows that the problem spaces become more constrained with the increase of the number of rules. How does the structural change of the problem space affect problem solving behavior? There are at least two rival factors. On the one hand, the fewer rules, the more paths there are from an initial state to a final state. Hence, fewer rules might make the problem easier. On the other hand, the more rules, the fewer the choices. Subjects can simply follow where the highly constrained structure forces them to go. So, more rules might make the problem easier. This implies that the problem difficulty might not change monotonically with the number of rules. The effect on problem solving behavior caused by the structural change of the problem space might also depend on the nature of the rules (whether internal or external). Experiments 1A and 1B examine these effects. In Experiment 1A, all rules are internal. We change the number of internal rules. In Experiment 1B, all but Rule 1 are external. We change the number of external rules.

Experiment 1A has four conditions, in which all rules were internal. Condition *I1* has Rule 1, Condition *I13* has Rules 1 and 3, Condition *I123* has Rules 1, 2, and 3, and Condition *I1234* has Rules 1, 2, 3, and 4. We made a restaurant story (*Waitress and Oranges*) for the instructions. The instructions for Condition *I1234* are shown in Figure 9. The instructions for Conditions *I123*, *I13*, and *I1* were the same as for *I1234*, except that different sets of rules were stated in the instructions.

Method

Subjects

The subjects were 24 undergraduate students enrolled in introductory psychology courses at the University of California, San Diego, who volunteered for the experiment to earn course credit.

Materials

Three plastic orange balls of different sizes (small, medium, and large) and three porcelain plates were used for all four conditions.

Design

Each subject played all four games, once each. There were twenty-four possible permutations for the four games. The twenty-four subjects were assigned to these permutations randomly. Due to a limitation in the number of subjects available, the first, second, third, and fourth games always started at positions S1, S2, S3, and S1 and ended at positions E1, E2, E3, and E1, respectively (see Figure 3). This treatment was not expected to cause significant systematic deviation because the task structures of the four problems each subject solved were different from each other, and the games were randomized.

Procedure

Each subject seated in front of a table and read the instructions aloud slowly. Then the subject was asked to turn the instruction sheet over and to attempt to repeat all the rules. If the subject could recite all the rules twice

Waitress and Oranges

A strange, exotic restaurant requires everything to be done in a special manner. Here is an example. Three customers sitting at the counter each ordered an orange. The customer on the left ordered a large orange. The customer in the middle ordered a medium sized orange. And the customer on the right ordered a small orange. The waitress brought all three oranges in one plate and placed them all in front of the middle customer (as shown in Diagram 1). Because of the exotic style of this restaurant, the waitress had to move the oranges to the proper customers following a strange ritual. No orange was allowed to touch the surface of the table. The waitress had to use only one hand to rearrange these three oranges so that each orange would be placed in the correct plate (as shown in Diagram 2), following these rules:

- Only one orange can be transferred at a time. (Rule 1)
- An orange can only be transferred to a plate on which it will be the largest. (Rule 2)
- Only the largest orange in a plate can be transferred to another plate. (Rule 3)
- The smallest orange and the largest orange can not be placed on a single plate unless the medium sized orange is also on that plate (Rule 4).

How would the waitress do this? That is, you solve the problem and show the movement of oranges the waitress has to do to go from the arrangement shown in Diagram 1 to the arrangement shown in Diagram 2.

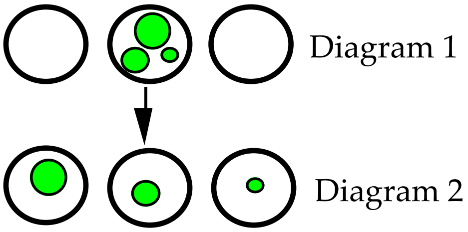

Figure 9. The instructions for the *I1234* condition in Experiment 1A. The instructions for other three conditions were the same as the one shown here, except that different sets of rules were stated.

without error, the subject was instructed to start the games. Otherwise the subject reread the instructions and was again tested. The cycle continued until the subject reached the criterion. The final states were presented to the subject in diagrams. A subject's hand movements and speech were recorded by a video camera. The solution time, which was from the time the experimenter said "start" to when the subject finished the last move, was recorded by a timer synchronized with the video camera.

Results

The results are shown in Figure 10. The minimum numbers of steps from the starting state to the final state are 2, 4, 7, and 8 for Conditions *I1*, *I13*, *I123*, and *I1234*, respectively. In order to make meaningful comparisons, solution times, solution steps, and errors for each condition were normalized by being divided by the number of minimum steps for each condition. The *p* values of the main effects and multiple comparisons are shown in the upper half of Table 1. When solution times and errors are used as the difficulty measurements, the difficulty order was, from easiest to hardest: $I1 < I13 < I1234 \leq I123$. The difference between *I1234* and *I123* was not statistically significant. When solution steps are used as the difficulty measurement, the difficulty order remained the same ($I1 < I13 \leq I1234 < I123$), but in this case the difference between *I13* and *I1234* was not statistically significant.

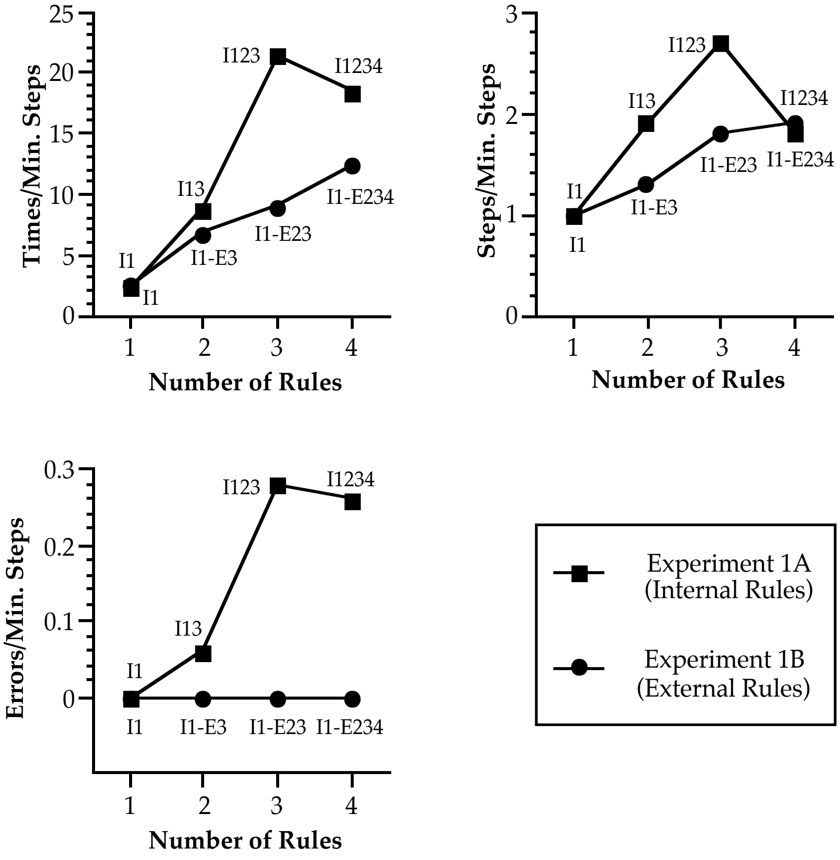

Figure 10. The results of Experiments 1A and 1B. Solution times were measured in seconds. The minimum numbers of steps to solve the problems are 2, 4, 7, and 8 for Conditions *I1*, *I13*, *I123*, and *I1234*, respectively. In order to make meaningful comparisons, solution times, steps, and errors for each condition were normalized by being divided by the number of minimum steps for each condition.

Table 1. *p* Values of Experiments 1A and 1B

	Measurements		
Comparisons	Times	Steps	Errors
	Experiment 1A		
Main Effect	< .0001	< .0001	< .0001
I1 vs. I13	< .05	< .005	< .06
I1 vs. I123	< .00001	< .00001	< .001
I1 vs. I1234	< .00001	< .01	< .001
I13 vs. I123	< .02	< .01	< .005
I13 vs. I1234	< .005	> .6	< .01
I123 vs. I1234	> .36	< .003	> .75
	Experiment 1B		
Main Effect	< .0001	< .0001	—
I1 vs. I1-E3	< .01	< .1	—
I1 vs. I1-E23	< .0001	< .0001	—
I1 vs. I1-E234	< .00001	< .00001	—
I1-E3 vs. I1-E23	> .18	< .01	—
I1-E3 vs. I1-E234	< .0001	< .0001	—
I1-E23 vs. I1-E234	< .03	> .4	—

NOTE. Fisher PLSD test was used for the multiple comparisons.

EXPERIMENT 1B: PROBLEM SPACE STRUCTURES (EXTERNAL)

Experiment 1B was exactly the same as Experiment 1A, except that Rules 2, 3, and 4 were external rather than internal. There were also four conditions. In Condition *I1* (*Waitress and Oranges*, Figure 11A), Rule 1 was the only internal rule. In Condition *I1-E3* (*Waitress and Straws*, Figure 11B), Rule 1 was internal and Rule 3 was external. A smaller straw could be dropped into a larger straw and a larger straw could be placed outside (around) a smaller straw (that is, Rule 2 was not implemented here). However, the diameters of the three straws were so small (1 cm, 1.5 cm, and 2 cm, respectively) that a smaller straw inside a larger one could not be moved out without the larger straw being moved away first (external Rule 3). In Condition *I1-E23* (*Waitress and Coffee*, Figure 11C), Rule 1 was internal and Rules 2 and 3 were external. All cups were filled with coffee. A smaller cup could not be placed on the top of a larger cup (external Rule 2), as this would cause the coffee to spill. All subjects understood that spilling coffee was not a good behavior because they had to imagine that they were waitresses or waiters working in a restaurant (see the cover story in Figure 9). In this sense, Rule 2 was not a truly external rule with rigid physical constraints, but an external rule grounded in cultural knowledge. Rule 3 in this condition was external because a cup could not be moved if there was another cup on its top. In Condition *I1-E234* (*Waitress and Tea*, Figure 11D), Rule 1 was internal and Rules 2, 3, and 4 were external. All cups were filled with tea. Rules 2 and 3 in this condition were external for the same reason described in Condition *I1-E23*: a smaller cup could not be placed on the top of a larger one (Rule 2), and only the largest (topmost) cup could be moved (Rule 3). Rule 4 were external because the bottom of the largest cup was smaller than the top of the smallest cup and the bottom of the smallest cup was smaller than the top of the largest cup, that is, the largest cup and the smallest cup could not be placed on the top of each other.

The instructions for these four conditions were the same as for the *I1234* condition in Experiment 1A, except that different words for different materials were used and only internal Rule 1 was stated in the instructions (all other rules were not stated because they were external).

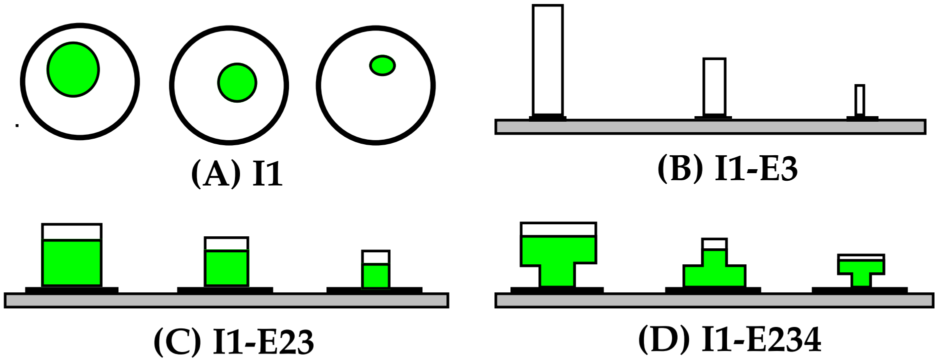

Figure 11. The four conditions of Experiment 1B. See text for explanations. (A) *I1*: Rule 1 is internal. (B) *I1-E3*: Rule 1 is internal and Rule 3 is external. (C) *I1-E23*: Rule 1 is internal and Rules 2 and 3 are external. (D) *I1-E234*: Rule 1 is internal and Rules 2, 3, and 4 are external.

Method

Subjects

The subjects were 24 undergraduate students enrolled in introductory psychology courses at the University of California, San Diego, who volunteered for the experiment to earn course credit.

Materials

Materials for Condition *I1* were the same as in Experiment 1A. In Condition *I1-E3*, the straws and tiny plates were made from paperboard. The diameters and heights of the small, medium, and large straws were approximately 1 cm, 1.5 cm, and 2 cm, and 3 cm, 6 cm, and 9 cm, respectively. In Condition *I1-E23*, three plastic cups of different sizes (small, medium, and large) and three paper plates were used. All three cups were filled with coffee. In Condition *I1-E234*, three cups made from metal cans and three paper plates were used. All three cups were filled with tea.

Design and Procedure

The design and procedure were the same as in Experiment 1A.

Results

The results are shown in Figure 10. The solution times, solution steps, and errors for all four conditions were normalized by being divided by the minimum number of steps for each condition, as in Experiment 1A. The p values of the main effects and multiple comparisons are shown in the lower half of Table 1. If solution times are used as the difficulty measurement, the difficulty order was, from easiest to hardest: *I1* < *I1-E3* ≤ *I1-E23* < *I1-E234*. The difference between *I1-E3* and *I1-E23* is not statistically significant. If solution steps are used as the difficulty measurement, the difficulty order remained the same (*I1* < *I1-E3* < *I1-E23* ≤ *I1-E234*), but the difference between *I1-E23* and *I1-E234* is not statistically significant. Subjects didn't make any errors in this experiment.

Discussion

From the results of Experiments 1A and 1B, we can see that problem space structure is an important factor of problem difficulty. For example, the solution time difference between the easiest problem (*I1*) and the hardest problem (*I123*) can be as large as a ratio of one to ten. We can also see that the effect of the structural change of a problem space on problem solving behavior depended on the nature of the rules. When all but one rule were external (Experiment 1B), a problem became more difficult when its problem space became more constrained (more rules). However, when all rules were internal (Experiment 1A), the hardest problem was neither the most constrained one (*I1234*) nor the least constrained one (*I1*), but the intermediately constrained problem (*I123*).

Two factors, working memory and problem structure, might contribute to the different difficulty orders caused by the nature of rules. In Experiment 1A, all rules were internal. The working memory load was high for both *I123* and *I1234* (three and four internal rules, respectively). One possible explanation is that subjects could not do much planning because most of the working memory was loaded by the processing of the internal rules. In this case, the structure of the problem space might be the dominant factor of problem difficulty. Moves were to a large extent guided by the structure of the problem space. *I1234* was easier than

I123 problem because it was more constrained than *I123* . In Experiment 1B, all but Rule 1 were external. In this case, planning was probably the dominant factor of problem difficulty because there was little load on working memory. This explains why *I1-E23* was easier than *I1-E234,* even though *I1-E234* was more constrained than *I1-E23*.

All four rules in Experiment 1A were internal and three of the four rules in Experiment 1B were external. Comparing the results in these two experiments, we found that the conditions with external rules in Experiment 1B were easier than their corresponding conditions with internal rules in Experiment 1A. In other words, external rules could make problems easier. This was explicitly tested in Experiment 2, which follows.

EXPERIMENT 2: THE DISTRIBUTED REPRESENTATION OF RULES

In this experiment, we study isomorphic representations at the level of rule representations. We have shown that a problem space is constructed by a set of rules. Given the same set of rules, the abstract problem space is fixed. However, the rules can be distributed across internal and external representations in different ways. Different distributions may have different effects on problem solving behavior, even if the formal structures are the same. This experiment examines these effects. Our hypothesis is that the more rules are distributed in the external representation, the easier the problem.

There were three conditions in this experiment, which all had three rules. Though the three rules were the same in their abstract forms, they were distributed across internal and external representations in different ways. In the *I123* (*Waitress and Oranges*) condition (Figure 12A), Rules 1, 2, and 3 were all internal. In the *I12-E3* (*Waitress and Donuts*) condition (Figure 12B), Rules 1 and 2 were internal, and Rule 3 was external. This is the standard TOH. Rule 3 was external because the physical constraints (coupled with Rule 1) guaranteed that it was followed (see the previous discussion in the section *Rule Representations* for more explanations). In the *I1-E23* (*Waitress and Coffee*) condition (Figure 12C), Rule 1 was internal, and Rules 2 and 3 were external. This was identical to the *I1-E23* condtion in Experiment 2B. The instructions for all three conditions were the same as for the *I1234* condition in Experiment 1A, except that different words suitable for the materials used in the current experiment were used and that only internal rules were stated in the current instructions (external rules need not be stated explicitly).

Method

Subjects

The subjects were 18 undergraduate students enrolled in introductory psychology courses at the University of California, San Diego who volunteered for the experiment to earn course credit.

Materials

The materials in the *I123* condition were the same as in the *I123* condition in Experiment 1A. In the *I12-E3* condition, three plastic rings of different sizes (small, medium, and large) and three plastic poles were used. The materials in the *I1-E23* condition were the same as in the *I1-E23* condition in Experiment 1B.

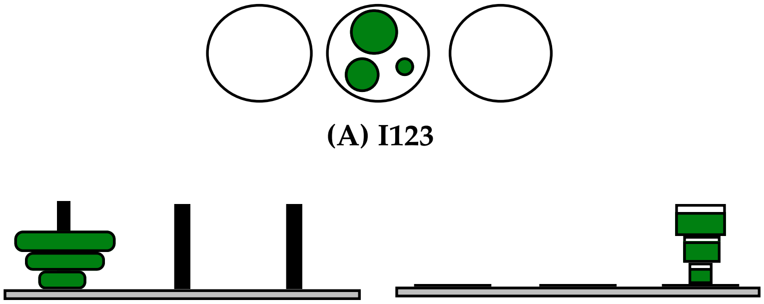

(A) I123

(B) I12-E3 **(C) I1-E23**

Figure 12. The three conditions of Experiment 2. See text for explanations. (A) *I123*: All three rules are internal. (B) *I12-E3*: Rules 1 and 2 are internal and Rules 3 is external. (C) *I1-E23*: Rule 1 is internal and Rules 2 and 3 are external.

Design

Each subject played all three games, one for each of the three conditions, once in a randomized order (e.g., *I1-E23, I123, I12-E3*). There were six possible permutations for the three games. Each permutation was assigned to a subject randomly. There were a total of eighteen subjects. Due to a limitation in the number of subjects available, the starting and ending positions were not randomized. That is, for each subject, the first, the second, and the third games always started at positions S1, S2, and S3 and ended at positions E1, E2, and E3, respectively (see Figure 3). The starting and ending positions should not cause significant systematic deviation because the three pairs of starting and ending positions were exactly symmetric, and the order of the three games played by each subject was randomized.

Procedure

The procedure was the same as in Experiment 1A.

Results

The average solution times, solution steps, and errors are shown in Figure 13. The *p* values for the main effects and multiple comparisons are shown in Table 2. Problem difficulty measured in solution times, steps, and errors for the three problems was consistent. The more rules were external, the easier the problem. The order of difficulty was, from hardest to easiest: *I123* > *I12-E3* ≥ *I1-E23*. The difference between *I12-E3* and *I1-E23* was not statistically significant. All errors made were for internal rules: none were for external rules.

Discussion

Two of the three conditions in the present experiment, *I123* and *I12-E3*, were modifications of the *Dish-move* and *Peg-move* problems of Kotovsky, Hayes, and Simon (1985). The results from the present study are consistent with their results: subjects took more time to solve *I123* than *I12-E3*. Kotovsky, Hayes, and Simon only reported solution times in their study. The numbers of steps and errors in the present study are all consistent with solution times. In the present experiment, the more rules externalized, the easier the problem. In addition, external rules seem to be error proof: subject did not make any errors for external rules. This effect might be due to the fact that the external rules were either perceptually available or physically constrained. The errors for internal errors might be caused by the load of working memory.

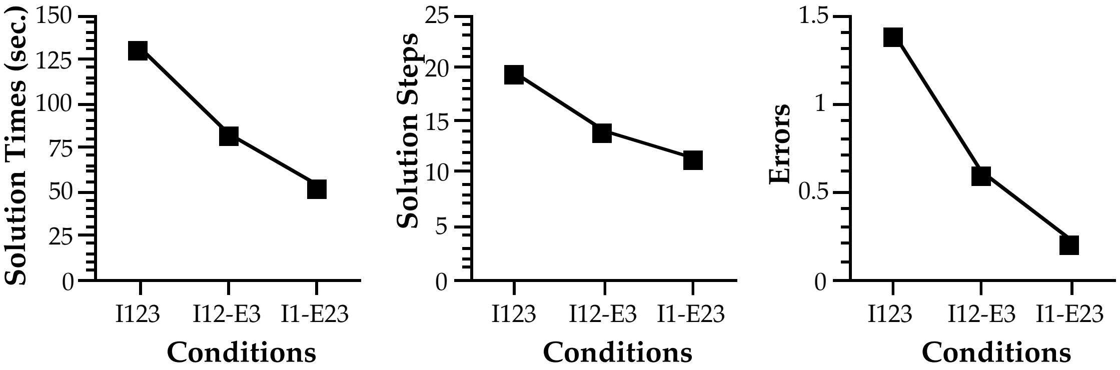

Figure 13. The results of Experiment 2. Problem difficulty decreased with the increase of the number of external rules.

Table 2. The *p* Values of Experiment 2

	Measurements		
Comparisons	Times	Steps	Errors
Main Effect	< .05	< .05	< .005
I123 vs. I12-E3	< .1	< .1	< .03
I123 vs. I1-E23	< .01	< .02	< .001
I12-E3 vs. I1-E23	> .3	> .4	> .2

NOTE: Fisher PLSD test was used for the multiple comparisons.

Why did external rules make problems easier? There might be two factors. The first is the checking of rules before each move. External rules can be checked by perceptual inspection, while internal rules must be checked mentally. The processing of internal rules in the mind, which demands more resources of working memory, might interfere with other processes critical for problem solving, such as planning. The second factor is the recursive strategy—the strategy to reduce a three-object problem to a two-object problem, or in other words, move the smallest object to its destination first. From the problem space of the TOH (Figure 3) we can see that for any pair of starting and ending states, as long as either the small or the medium object is in its final state, only three more steps are needed to solve the problem. Among the fifty-four games played by the eighteen subjects, the last three steps were solved in three steps in fifty-two games and in five steps in only two games. Therefore, if either the small or the medium object is in the final state, the game is virtually solved. Rule 2 might be critical for the discovery of the recursive strategy. For the *I1-E23* condition, Rule 2 was external: a smaller cup of coffee could not be placed on the top of a larger one. This external representation might have prompted the subjects that the smallest cup had to be moved to its final state first. Out of the eighteen games for the *I1-E23* condition, sixteen were solved by moving the smallest cup to its final state first. The other two were solved by moving the medium sized cup to its final state first.

For the *I123* and the *I12-E3* conditions, Rule 2 was internal. Thus, the recursive strategy was harder to discover. Out of the eighteen games for the *I123* condition, only eleven were solved by moving the smallest object to its final state first. The other seven were solved by moving the medium sized object to its final state first. Similarly, out of the eighteen games for the *I12-E3* condition, the numbers for the two cases were ten and eight, respectively. Thus, the ease of discovering the recursive strategy in the *I1-E23* condition might have also contributed to the difficulty order.

EXPERIMENT 3: DIMENSIONAL REPRESENTATIONS

In this experiment, we study isomorphic representations at the level of dimensional representations. The ordinal and nominal dimensions of the TOH can be represented by any properties. We chose size, location, and color to represent these two dimensions, which generated nine isomorphic representations of the TOH (Figure 14). (See Figure 6 as well as Figure 14 for explanations on the dimensional representations and the mappings between dimensions).

The nine isomorphs in Figure 14 were the nine conditions of this experiment, which had the same three internal rules. Among these nine isomorphs, some might be easier than others. We consider two factors of problem difficulty. The first factor is the nature (internal or external) of the ordinal dimension. The ordinal dimension is represented externally by sizes and (ordered) locations, because the ordinal relation is embedded in the physical properties of sizes and locations. But it is represented internally by colors, because the order of colors is arbitrary and has to be internalized. (The nominal dimension is represented externally by sizes, locations, and colors.) The second factor is whether the ordinal and nominal dimensions are represented by visual properties (sizes and colors) or spatial properties (locations). The instructions for the O(size)-N(size) condition are shown in Figure 15. The instructions for other conditions were the same as for the O(size)-N(size) condition, except that modifications were made for different materials used in different conditions.

We hypothesize that external dimensions need less mental processing than internal dimensions, and that spatial dimensions support more efficient perceptual processing than visual dimensions. Thus, we have the following predictions. For the ordinal dimension, sizes should be better than colors because the former represents the ordinal dimension externally while the latter represents it internally, and locations should better than sizes because the former is a spatial dimension while the latter is a visual dimension. In addition, locations should be better than colors because the former is not only an external dimension but also a spatial dimension. For the nominal dimension, locations should be better than sizes and colors because the former is a spatial dimension while latter two are visual dimensions. Sizes and colors should not differ from each other because they are both visual dimensions.

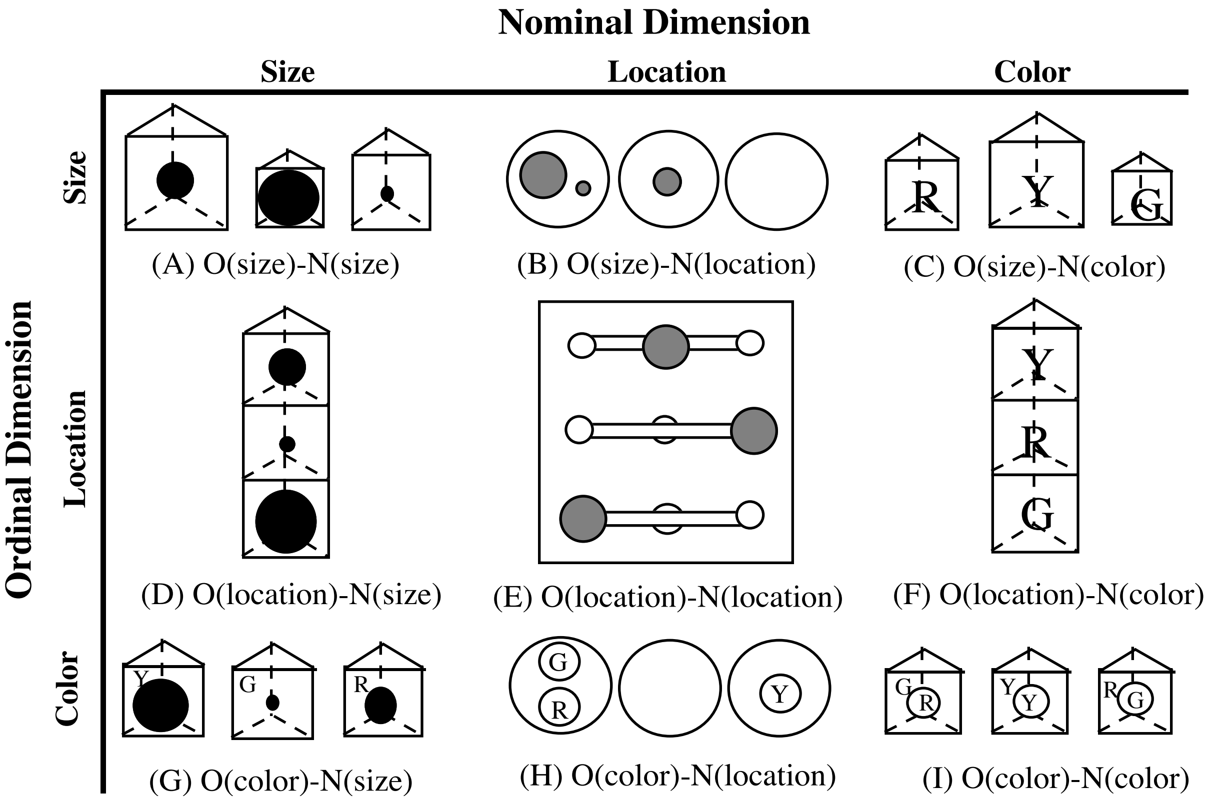

(A) O(size)-N(size) (B) O(size)-N(location) (C) O(size)-N(color)

(D) O(location)-N(size) (E) O(location)-N(location) (F) O(location)-N(color)

(G) O(color)-N(size) (H) O(color)-N(location) (I) O(color)-N(color)

Figure 14. Nine isomorphs at the level of dimensional representations. O = Ordinal Dimension, N = Nominal Dimension. For ordinal dimensions, the colors have a priority: red > green > yellow. All triangular cylinders can be rotated around their axes. (A) Each cylinder has three different sized circles on its three sides. O = sizes of the cylinders; N = sizes of the circles. (B) O = sizes of the balls; N = locations of the balls. (C) Each cylinder has three colors (red, green, yellow) on its three sides. O = sizes of the cylinders. N = colors of the sides. (D) Each cylinder has three different sized circles on its three sides. O = locations of the cylinders (top, middle, bottom). N = sizes of the circles. (E) The three disks can be moved horizontally among three locations. O = vertical locations (top, middle, bottom). N = horizontal locations (left, middle, right). (F) Each cylinder has three colors (red, green, yellow) on its three sides. O = locations of the cylinders (top, middle, bottom). N = colors of the sides. (G) Each cylinder has a color (red, green, or yellow) on all three sides. On the three sides of a cylinder there are three different sized circles. O = color priority. N = sizes of the circles. (H) O = color priority. N = locations of the balls. (I) Each cylinder has a color (red, green, or yellow) on all three sides. On the three sides of each cylinder there are three circles of different colors (red, green, yellow). O = the color priority of the cylinders. N = colors of the circles.

Legend has it that in ancient India there was a Buddhist scripture in a locked crypt hidden in a cave. Many people tried to acquire it because it possessed the secret of life. However, none of them succeeded because if the crypt was not opened in a very short time, it would disappear. The lock of the crypt was remotely controlled by three magic triangular cylinders of different sizes (small, medium, and large). On the three sides of each cylinder, there were three different sized circles (also small, medium, and large). In order to open the crypt, one had to rotate the three cylinders to a specific configuration, strictly following a set of rules. One day, a wise monk entered the cave. Inspired by the wisdom of Buddha, he skillfully opened the crypt within the specified time. How did he do it? You are given a replica of the magic cylinders. Try to solve this puzzle as fast as you can. The rules are given below. The initial configuration you will start with and the final configuration that opens the lock will be given to you when you have memorized these rules.

Rule 1: Only one cylinder can be rotated at a time.

Rule 2: After rotating a cylinder, if the new facing side has a same sized circle as that of another cylinder, the size of the cylinder you rotated must be larger.

Rule 3: Before rotating a cylinder, if any cylinders have matching sized circles facing you, only the largest cylinder can be rotated.

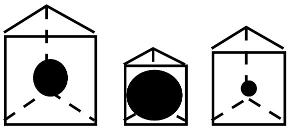

Figure 15. The instructions for the O(size)-N(size) condition. The instructions for other conditions were the same as the one shown here, except that modifications were made for different objects in different conditions.

Method

Subjects

The subjects were 36 undergraduate students at the University of California, San Diego who were paid for participating in this experiment.

Materials

The triangular cylinders in Figure 14 were made from paperboard and could be rotated around their vertical axes. In Figures 14D and 14F, the three cylinders were stacked on a vertical rod. In Figure 14B, three different sized plastic balls were used. In Figure 14E, the three plastic disks could only be moved horizontally. In Figure 14H, three sponge balls of different colors (red, green, yellow) were used.

Design

This was a mixed design. The within-subject factor was the ordinal dimension. It had three levels: size, location, and color. The between-subject factor was the nominal dimension. It also had three levels: size, location, and color. There were three games across the three within-subject levels at each between-subject level. Each subject played these three games once in a randomized order (e.g., O(color)-N(size), O(size)-N(size), and O(location)-N(size)). For the three games at each between-subject level, there were six possible permutations. There were a total of 18 possible permutations across the three between-subject levels. Each permutation was assigned to two subjects randomly. There were a total of 36 subjects. Due to a limitation in the number of subjects available, the starting and ending positions were not randomized. That is, for each subject, the first game always started at position S1 and ended at position E1 (see Figure 3), the second at position S2 and at position E2, and the third at position S3 and at position E3. The starting and ending positions were not expected to cause much systematic deviation because the three pairs of starting and ending positions were exactly symmetric, and the three games played by each subject were randomized.

Procedure

For each game, the subjects were asked to read the instructions aloud once and given two minutes to memorize the rules. The subjects were then asked to recite all the rules. If the subjects could recite all the rules without errors, they were shown two examples for each rule. Otherwise, the subjects read the instructions again and were again tested until they could recite all the rules without errors. Before the subjects started the games, they were given the initial and final states. The final state for each game was presented by a set of identical objects. Subjects' hand movements and speech were recorded by a video camera.

Results

The solution times, steps, and errors are shown in Figure 16. The *p* values are shown in Table 3. For the ordinal dimension, when solution times and errors were used as the difficulty measurement, locations were significantly better than sizes and colors. Though the solution times and errors for sizes were smaller than for colors, the differences were not statistically significant. When solution steps were used as the difficulty measurement, locations were significantly better than sizes, but not statistically better than colors, though the solution steps for locations were fewer than for colors. The difference of solution steps between colors and sizes was not statistically significant.

For the nominal dimension, when solution times and errors were used as the difficulty measurement, locations were significantly better than sizes and colors. The difference between sizes and colors was not statistically significant. When solution steps were used as the difficulty measurement, none of the differences between sizes, locations, and colors were statistically significant, though solution steps for locations were fewer than for sizes and colors.

Discussion

We identified two factors that might affect problem solving behavior at the level of dimensional representations. The first factor is whether the ordinal dimensional is represented internally or externally. The results showed that although solution times and errors for colors, which represent the ordinal dimension internally, were larger than for sizes, which represent the ordinal dimension externally, the differences were not statistically significant. This might be because, compared with other mental activities during the problem solving tasks, internalizing the color priority (red>green>yellow) was a trivial task. The second factor, whether a dimension is represented by visual or spatial properties, played a major role in determining problem difficulties. Locations were much better than sizes for the ordinal dimension and than sizes and colors for the nominal dimension. The superiority of locations over colors for the ordinal dimension was an effect caused by both factors.

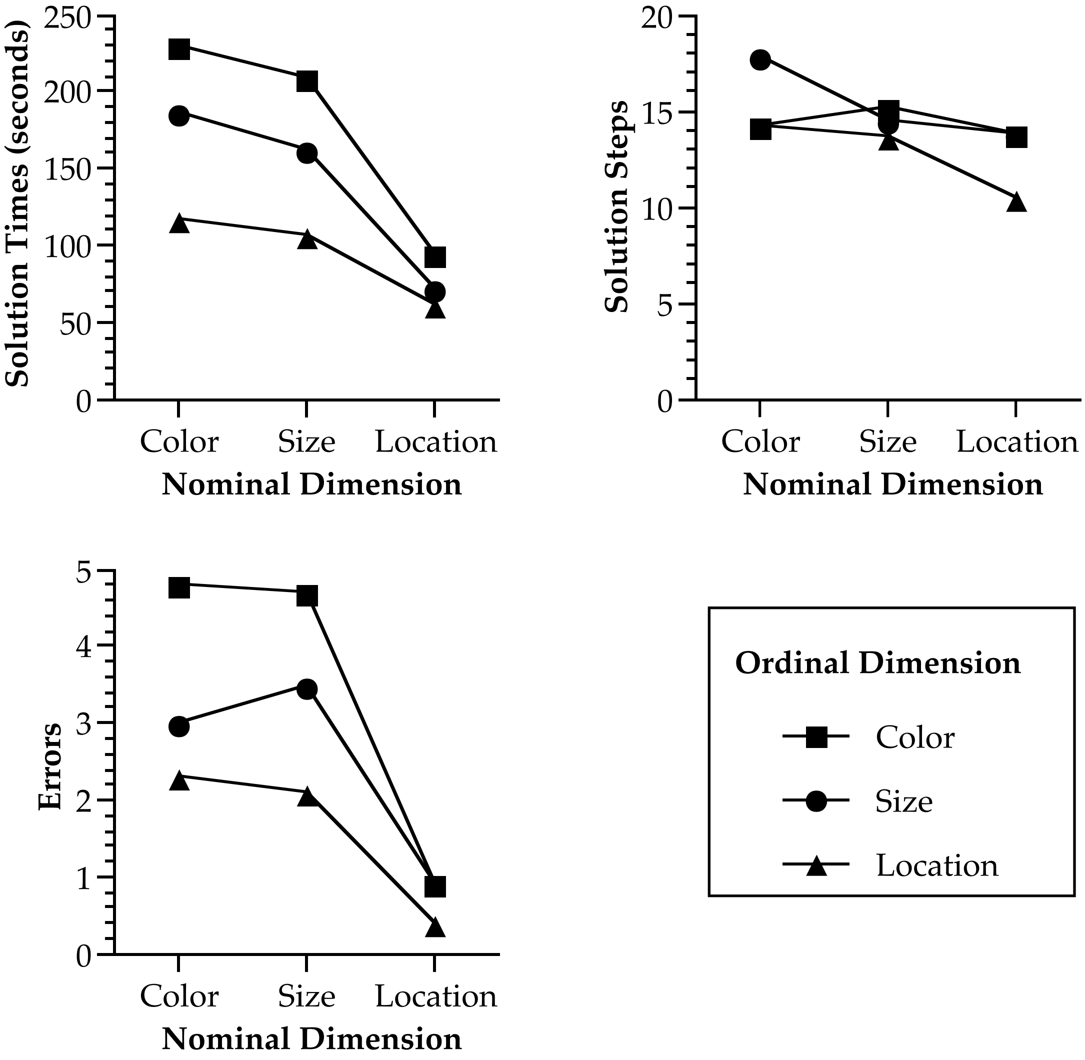

Figure 16. The results of Experiment 3.

Table 3. The *p* Values of Experiment 3

	Measurements		
Comparisons	Times	Steps	Errors
O(size) vs. O(location)	< .01	< .05	< .03
O(size) vs. O(color)	> .17	> .53	> .15
O(location) vs. O(color)	< .0007	> .2	< .004
N(size) vs. N(location)	< .003	> .24	< .0001
N(size) vs. N(color)	> .66	> .62	> .96
N(location) vs. N(color)	< .016	> .17	< .007

NOTE. The comparisons were between dimensions, not between individual conditions. For example, O(size) vs. O(location) was the comparison between the ordinal dimension represented by sizes and the ordinal dimension represented by locations. None of the interactions were statistically significant.

Why were locations so special? In the analysis of the dimensional representations of the TOH, we reviewed some empirical studies showing the important anatomical and functional differences between visual and spatial representations (Goldman-Rakic 1992; Mishkin, Ungerleider, & Macko 1983; Treisman & Gelade, 1980). Though the results from the present experiment do not allow us to offer a direct explanation of why locations (spatial dimension) were better than sizes and colors (visual dimensions), we can compare our results with some perceptual tasks. The ordinal and nominal dimensions in the present experiment are conjunctive dimensions. Treisman & Gelade (1980) proposed that location information is necessary for the perception of conjunctive stimuli. Nissen's (1985) study of conjunctive stimuli based on visual maps is more directly related to the present study. According to her, color and shape are registered in separate maps. When a conjunctive stimulus was composed of color and location, she showed that selecting an item by color and reporting its location was as accurate as selecting by location and reporting color. In this case, selection by location did not have special advantage, because location itself was part of the conjunctive stimulus. However, when a conjunctive stimulus is composed of color and shape, location is necessary for cross-referencing between the separate shape and color maps. She showed that when subjects reported the shape and location of an item cued by its color, the accuracy of shape judgments depended on the accuracy of locating the cued color. Comparing Nissen's study with the present study, we have parallel results. If neither of the two conjunctive dimensions (ordinal and nominal dimensions) is represented by locations, an extra step of invoking location information is needed to link the two separate visual maps of size and color. In this case, the three values on the ordinal dimension and the three values on the nominal dimensions, as well as the three locations, have to be checked before an operation (moving or rotating) can be made. If one of the two conjunctive dimensions is represented by locations, the extra step of invoking the location information is not needed, because the location information itself is in the conjunctive dimensions. In this case, only the three values on the dimension not represented by locations and the three values on the dimension of locations have to be checked before an operation can be made. Thus, if one of the ordinal and nominal dimensions or both of them are represented by locations, the problem is easier.

GENERAL DISCUSSION

In this paper, we developed a theoretical framework of distributed representations and a methodology of representational analysis for the study of distributed cognitive tasks—tasks that require the processing of information distributed across the internal mind and the external environment, and applied them to the empirical studies of the Tower of Hanoi problem. Our approach to distributed cognitive tasks demands (a) the consideration of the internal and external representations of a task as a representational system, (b) the explicit decomposition of the representational system into its internal and external components, and (c) the identification of the different functions of internal and external representations in cognition. The traditional approach to cognition is not appropriate for the study of distributed cognitive tasks, because it often ignores the important functions of external representations in cognition. In addition, it often confuses a task's distributed representation that has both internal and external components with the task's internal representation. This confusion often leads to unnecessarily complex accounts of cognition.

Formal Structures, Representations , and Processes

Any distributed cognitive task can be analyzed into three aspects: its formal structure, its representation, and its processes. Our present study focused on representations and their relation to formal structures, for the following reasons. First, in order to understand the processes involved in a distributed cognitive task, we first have to understand what information is processed and how the information to be processed is represented. Second, different processes are activated by different representations, but not vice versa. For example, perceptual processes are activated by external representations, while cognitive processes are usually activated by internal representations. Third, from a representational perspective, tasks that look dramatically different may in fact have a common structure. However, if we start our analysis with processes, we may fail to capture this common structure, because different representational formats of the common structure can activate completely different processes.

The three aspects of distributed cognitive tasks are closely interrelated: the same formal structure can be implemented by different isomorphic representations, and different isomorphic representations can activate different processes. Though an representational analysis should be the first step for the study of distributed cognitive tasks, a process model is also essential. In order to understand the nature of distributed cognitive tasks, we need to study all of the three aspects. Two interesting issues worth of further studies are the relationship between representations and processes and the interplay between perceptual and cognitive processes.

The Nature of External Representations

External objects are not just peripheral aids to cognition, they provide a different form of representation—an external representation. By decomposing the representational system of a distributed cognitive task into its internal and external representations, we can separate the functions of external representations from those of internal representations. The empirical studies reported here on the Tower of Hanoi problem suggest the following properties of external representations.

1. *External representations can provide memory aids*. For example, in all of the TOH experiments reported here, the goal problem states didn't need to be memorized, because they were placed in front of the subjects either by diagrams or by physical objects. This is the most acknowledged property of external representations. To many people, this is the only one.

2. *External representations can provide information that can be directly perceived and used without being interpreted and formulated explicitly*. For example, in the *I1-E23* version of the TOH, Rules 2 and 3 were not told to the subjects: they were built into the physical constraints. They could be perceived and followed directly by the subjects. When the subjects were asked to explic-

itly formulate the rules after the games, few could do it. External representations seem to provide affordances (Gibson, 1979).

3. *External representations can anchor and structure cognitive behavior*. The physical structures in external representations constrain the range of possible cognitive behaviors in the sense that some behaviors are allowed and others prohibited. For example, in the *I1-E23* version of the TOH, external Rules 2 and 3 could not be violated. They constructed an action space in which only the actions that did not violate these two rules were permitted. This action space is the external problem space of the *I1-E23* problem.

4. *External representations change the nature of a task*. Norman (1991) proposed that there are two different views of cognitive artifacts. From the system's view (internal + external representations), external representations can make a task easier; from the person's view (internal representations only), external representations change the nature of the task. For example, in the *I123* version of the TOH, a subject had to process three internal rules, while in the *I1-E23* version the subject only had to process one internal rule. Though the two versions had the same abstract structure, their cognitive processes were different. Nevertheless, when considered as systems, *I1-E23* was much easier than *I123*.

5. *External representations are an indispensable part of the representational system of any distributed cognitive task*. This property is a direct reflection of the nature of distributed cognitive tasks, which require the processing of information distributed across internal and external representations.

Related Approaches

Although the mainstream approach to cognition has focused on internal representations, the role of the environment in cognition has long been acknowledged by several alternative approaches. For example, Gibson (1966, 1979) argued that perception is the direct pickup of environmental information (invariants) in the extended spatial and temporal patterns of optic arrays, and that information in the environment is sufficient for perception and action. The sociohistorical approach to cognition (Leontiev, 1981; Luria, 1976; Vygotsgy, 1978, 1986) argues that it is the continuous internalization of the information and structure from the environment and the externalization of internal representations into the environment that produce high level psychological functions.

More recently, the role of the environment in cognition has become the central concern in several fields of cognitive science. In the studies of the relationship between images and pictures, it has been shown that external representations can give us access to knowledge and skills that are unavailable from internal representations (e.g., Chambers & Reisberg, 1985; Reisberg, 1987). The situated cognition approach argues that the activities of individuals are situated in the social and physical contexts around them and knowledge can be considered as a relation between the individuals and the situation (e.g., Barwise & Perry, 1983; Greeno, 1989; Lewis, 1991; Suchman, 1987). Studies on diagrammatic reasoning have also focused on the functions of external representations in cognition (e.g., Chandrasekaran & Narayanan, 1990; Larkin, 1989; Larkin & Simon, 1987). For example, Larkin & Simon (1987) argue that diagrammatic representations support operators that can recognize features easily and make inferences directly.

Although our current approach shares the same interest with others in the function of the environment in cognition, it differs in several aspects. First, our approach focuses on the representational properties of distributed cognitive tasks: how the information needed for a distributed cognitive task is represented across the internal mind and the external environment. Second, our approach demands the consideration of the internal and external representations of a distributed cognitive task as a representational system. We argue that external representations are an indispensable part of any distributed cognitive task. Third, our approach demands the explicit decomposition of the representational system of a distributed cognitive task into its internal and external components. With such a decomposition, we can identify the different functions of internal and external representations in cognition. Fourth, we suggested that in order to understand the nature of a distributed cognitive task, we need to study the task's formal structure, representation, and processes and the interrelations among them. Finally, the principles of distributed representations and the methodology of representational analysis developed in the present study have been applied to several real world problems, including numeration systems, relational information displays, and cockpit instrument displays (Zhang, 1992).

Conclusion

The distributed representations approach offers a novel perspective for the study of cognition. It has both theoretical and practical implications. Theoretically, it can shed light on some issues regarding the nature of cognition, such as whether cognition is solely in the mind or distributed across the mind and the environment and whether people reason on formal structures or on content-specific representations. Practically, it can provide design principles for effective representations.

REFERENCES

Barwise, J. & Perry, J. (1983). Situations and attitudes. Cambridge, MA: The MIT Press.

Chambers, D. & Reisberg, D. (1985). Can mental images be ambiguous? *Journal of Experimental Psychology: Human Perception & Performance, 11,* 317-328.

Chandrasekaran, B. & Narayanan, N. H. (1990). Towards a theory of commonsense visual reasoning. In K. V. Nori & C. E. V. Madhavan (Eds.): *Lecture notes in computer science 472*. Berlin, Germany: Springer-Verlag.

Cheng, P. W. & Holyoak, K. J. (1985). Pragmatic reasoning schemas. *Cognitive Psychology, 17,* 391-416.

Evans, J. S. B. T. (Ed.). (1983). *Thinking and reasoning.* London: Routledge & Kegan Paul.

Gibson, J. J. (1966). *The senses considered as perceptual systems*. New York: Houghton Mifflin Company.

Gibson, J. J. (1979). *The Ecological Approach to Visual Perception.*

Boston: Houghton Mifflin.

Greeno, J. G. (1974). Hobbits and Orcs: Acquisition of a sequential concept. *Cognitive Psychology, 6,* 270-292.

Greeno, J. G. (1989). Situations, mental models, and generative knowledge. In D. Klahr & K. Kotovsky (Eds.), *Complex information processing: The impact of Herbert A. Simon.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Griggs, R. A. & Newstead, S. E. (1982). The role of problem structure in a deductive reasoning task. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 8 (4),* 297-307.

Hayes, J. R., & Simon, H. A. (1977). Psychological differences among problem isomorphs. In N. J. Castellan, D. B. Pisoni, & G. R. Potts (Eds.), *Cognitive theory*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Hutchins, E. (1990). The technology of team navigation. In J. Galegher, R. E. Kraut, & C. Egido (Eds.), *Intellectual teamwork: Social and technical bases of collaborative work*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Hutchins, E. (in preparation). *Distributed Cognition: a cognitive ethnography of ship navigation*. San Diego: University of California, Department of Cognitive Science.

Jeffries, R., Polson, P. G., Razran, L., & Atwood, M. E. (1977). A process model for Missionaries-Cannibals and other River-Crossing problems. *Cognitive Psychology, 9,* 412-440.

Kirlik, A. (1989). *The organization of perception and action in complex control skills.* Unpublished PhD Dissertation. Columbus, OH: The Ohio State University, Department of Industrial and Systems Engineering.

Kotovsky, K. & Fallside, D. (1989). Representation and transfer in problem solving. In D. Klahr & K. Kotovsky (Eds.), *Complex information processing: the impact of Herbert A. Simon*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Kotovsky, K. & Simon, H. A. (1990). What makes some problems really hard: explorations in the problem space of difficulty. *Cognitive Psychology, 22,* 143-183.

Kotovsky, K., Hayes, J. R., & Simon, H. A. (1985). Why are some problems hard? Evidence from Tower of Hanoi. *Cognitive Psychology, 17,* 248-294.

Larkin, J. (1989). Display-based problem solving. In D. Klahr & K. Kotovsky (Eds.), *Complex Information Processing: The Impact of Herbert A. Simon*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Larkin, J. H. & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science , 11,* 65-99.

Leontiev, A. N. (1981). *Problems in the development of mind*. Moscow: Progress.

Lewis, C. M. (1991). *Representational aiding*. Unpublished manuscript. Pittsburgh, PA: University of Pittsburgh, Department of Information Science.

Luria, A. R. (1976). *Cognitive development: its cultural and social foundations*. Cambridge, MA: Harvard University Press.

Margolis, H. (1987). *Patterns, thinking and cognition: A theory of judgment*. Chicago, IL: The University of Chicago Press.

Mishkin, M., Ungerleider, L. G. and Macko, K. A. (1983). Object vision and spatial vision: Two cortical pathways. *Trends in Neuroscience, 6*, 414-417.

Nickerson, R. (1988). Counting, computing, and the representation of numbers. *Human Factors, 30*, 181-199.

Nissen, M. J. (1985). Accessing features and objects: is location special? In M. I. Posner & O. S. M. Martin (Eds.), *Attention and performance XI*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Norman, D. A. (1988). *The psychology of everyday things*. New York: Basic Books.

Norman, D. A. (1991). Cognitive artifacts. In J. M. Carroll (Ed.), *Designing interaction: Psychology at the human-computer interface*. New York: Cambridge University Press.

Norman, D. A. (in press). *Things that make us smart*. Reading, MA: Addison-Wesley.

O'Brien, D. P., Noveck, I. A., Davidson, G. M., Fisch, S. M., Lea, R. B., & Freitag, J. (1990). Sources of difficulty in deductive reasoning: The THOG task. *Quarterly Journal of Experimental Psychology, 42A, 2*, 329-351.

Reisberg, D. (1987). External representations and the advantages of externalizing one's thoughts. *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Simon, H. A. (1981). *The Sciences of the Artificial (2nd ed.)* . Cambridge, MA: The MIT Press.

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

Stevens, S. S. (1946). On the theory of scales of measurement. *Science, 103 (2684)*, 677-680.

Stevens, S. S. (1951). Mathematics, measurement, and psychophysics. In S. S. Stevens (Ed.), *Handbook of Experimental Psychology*. New York: Wiley.

Suchman, L. (1987). *Plans and situated actions: The problem of human-machine communication*. New York: Cambridge University Press.

Thomas, J. C. (1974). An analysis of behavior in the Hobbits-Orcs problem. *Cognitive Psychology, 6*, 257-269.

Treisman, A. M. & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology, 12*, 97-136.

Vygotsky, L. S. (1978). *Mind in society*. Cambridge, MA: Harvard University Press.

Vygotsky, L. S. (1986). *Thought and language*. Cambridge, MA: The MIT Press.

Wason, P. C. (1966). Reasoning. In B. M. Foss (Ed.), *New horizons in psychology*. Harmondsworth: Penguin.

Wason, P. C. & Johnson-Laird, P. N. (1969). Proving a disjunctive rule. *Quarterly Journal of Psychology, 21*, 14-20.

Wason, P. C. & Johnson-Laird, P. N. (1972). *Psychology of reasoning: Structure and content*. Cambridge, MA: Harvard University Press.

Zhang, J. (1992). *Distributed representation: The interaction between internal and external information*. Unpublished PhD Dissertation. San Diego: University of California, Department of Cognitive Science.

Zhang, J. & Norman, D. A. (1993). A cognitive taxonomy of numeration systems. *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates.