A Distributed Representation Approach to Group Problem Solving

Jiajie Zhang[1]
Department of Psychology
The Ohio State University

Abstract

This article develops a theoretical framework of distributed representations to explore the representational properties in group problem solving. The basic principle of distributed representations is that the representational system of a group problem solving task is distributed across the representations of individuals, which together represent the abstract structure of the task. The framework was used to analyze the distributed representation of the Waitress and Orange task. From this analysis, an experiment was designed to examine group problem solving behaviors under different distributed representations. The experiment shows that (1) different distributed representations across two individuals produced dramatically different group problem solving behaviors even if they had the same abstract structure, and (2) two minds could be better than, not different from, or even worse than one mind, depending on how representations were distributed across the two minds. These results further support the interactionist view of group problem solving, which is that the interactions among individuals can produce group cognitive properties that can neither be reduced to nor be inferred from the cognitive properties of individuals.

Group problem solving refers to problem solving activities that involve interactions among a group of individuals. One critical issue in group problem solving is the nature of group properties. One view is that the cognitive properties of a group can be entirely determined by the properties of individuals. In this reductionist view, to understand group behavior, all we need is to understand the properties of individuals. Another view is that the interactions among the individuals can produce emergent group properties that cannot be reduced to the properties of the individuals. In this interactionist view, to study group behavior, we not only need to examine the properties of individuals but, more importantly, we also need to consider the interactions among the individuals as the basic units of analysis. The interactionist view has been evidenced by much of the research, such as the recent study of distributed cognition—the study of cognitive tasks that are distributed across the internal mind and the external environment, among a group of individuals, and across space and time (e.g., Hutchins, 1990, 1994, 1995; Norman, 1990, 1993; Zhang, 1996; Zhang & Norman, 1994, 1995). For instance, Hutchins (1995) has shown that the cognitive properties of a distributed system such as the airplane cockpit can differ radically from the cognitive properties of the individuals, and they cannot be inferred from the properties of the individuals alone, no matter how detailed the knowledge of the properties of those individuals may be. Examples of emergent group properties in other domains include group affect (George, 1990), collective efficacy (Bandura, 1986; Guzzo & Shea, 1993), and transactive memory systems (Wegner, 1987).

Another critical issue in group problem solving is the group effectiveness problem (Foushee & Helmreich, 1988). Most people would argue that two minds are better than one because in a group there are much more resources, task load and memory load are shared and distrib-

---

[1] This article will appear in the *Journal of American Society of Information Science*. Correspondence and requests for reprints should be sent to Jiajie Zhang, Department of Psychology, The Ohio State University, 1885 Neil Avenue, Columbus, OH 43210. Email: zhang.52@osu.edu.

uted, errors are cross-checked, and so on. This is a condition that can result in what Steiner (1972) called *process gain*. However, there are also conditions that can result in *process loss* (Steiner, 1972)—a phenomenon that the performance of a group is worse than that of an individual because in a group communication takes time, knowledge may not be shared, different strategies may be used by different individuals, and so on (McNeese, Zaff, & Brown, 1992). Most studies of the group effectiveness problem have focused on social and personality factors (e.g., Hackman & Morris, 1975; Hare, 1972; Janis, 1972; Lanzetta & Roby, 1960; Latane, Williams, & Harkins, 1979). Representational factors, in comparison, have not been extensively studied.

This article addresses the above two issues about group properties and group effectiveness from a purely cognitive perspective, focusing on the representational properties in group problem solving. It is divided into four sections. The first section develops a theoretical framework of distributed representations for group problem solving. The second section uses this framework to analyze the distributed representation of the Waitress and Orange problem. From this analysis, the third section designs an experiment to examine the effects of representation distribution on group properties and group effectiveness. Finally, the last section discusses the general implications of the framework of distributed representations for group problem solving.

Distributed Representations in Group Problem Solving

The basic principle is that the representational system for a group problem solving task is distributed across the representations of individuals. Figure 1 shows a representational system for a group problem solving task with four individuals. Each individual has a representation for the task. The task has a single abstract task space that represents the abstract structure of the task. On one hand, the representations of the individuals jointly determine the abstract task space; on the other hand, the abstract task space can be decomposed and distributed across the representations of the four individuals. The representations of the individuals interact with each other and together form a distributed representation space, which is the actual space in which the group problem solving task is performed.

This framework of distributed representations for group problem solving is interactionist in nature. The representation of a group problem solving task is not in any individual's mind, but distributed across all individuals. In this view, a group problem solving task requires dynamic, interactive, and integrative processing of the information distributed across individual representations. The abstract task space is an emergent group property jointly determined by individual representations: it does not belong to any individual.

The abstract task space of a group problem solving task can be distributed across individual representations in different ways. Let us consider a few cases. The first case is that none of the individual representations is a complete representation of the abstract task space but they together represent the abstract task space. In this case, individual representations may or may not overlap with each other, that is, there may or may not be redundancy in the distributed representation. The second case is that some individual representations each represent the complete abstract task space whereas others each only represent part of the abstract task space. In this case, there is always redundancy in the distributed representation because the individual representations that represent the complete abstract task space are always supersets of the individual representations that only represent part of the abstract task space. The third case is that every individual representation is a complete representation of the abstract task space. In this case the redundancy in the distributed representation is maximum.

Different distributed representations of the same abstract task space may cause dramatically different group problem solving behaviors even if they all represent the same structure at an abstract level. In addition, group problem solving behavior may be considerably different from individual problem solving behavior even if their structures are the same. These two issues are examined in the following sections in the representational analysis of the Waitress and Orange problem and the corresponding experiment.

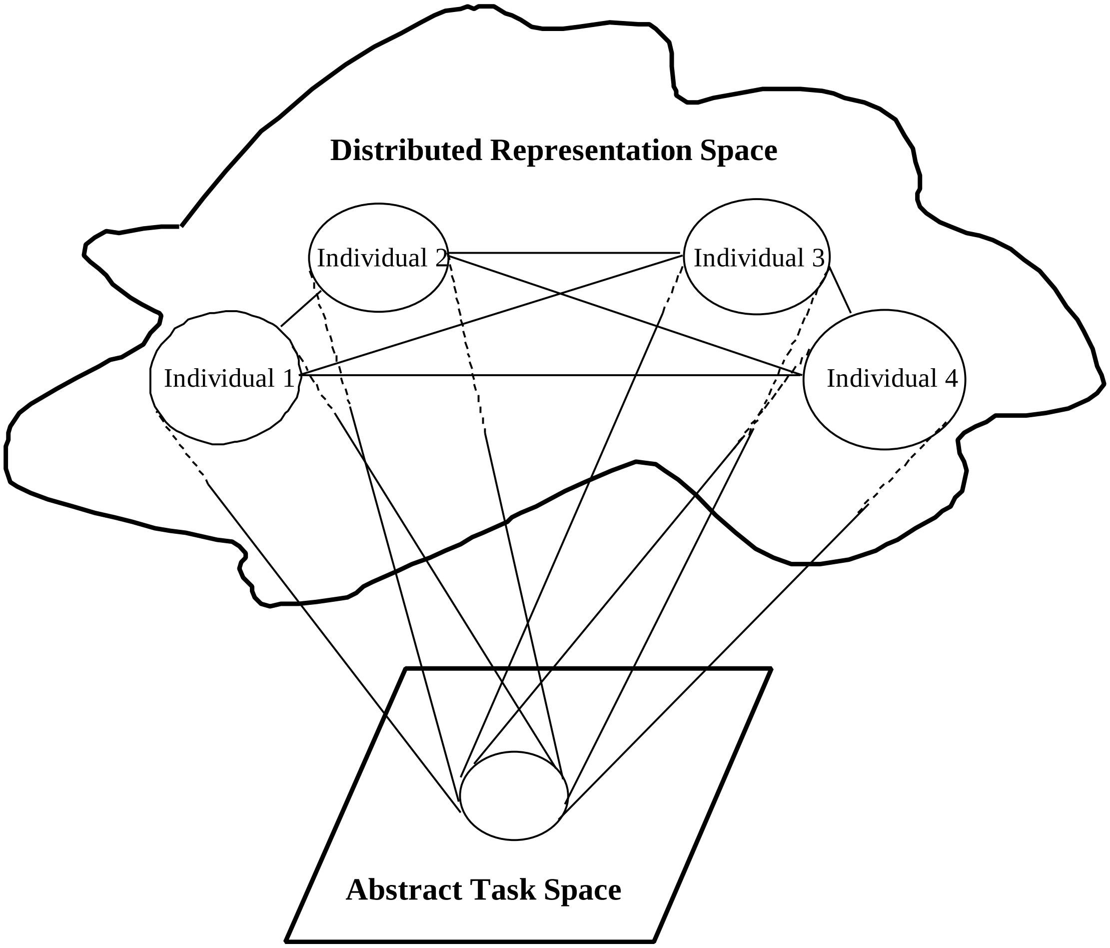

Figure 1. The framework of distributed representation for group cognitive tasks. The abstract task space, which represents the abstract structure of the task, is distributed across the representations of individuals.

The Waitress and Orange Problem

In this section, the framework of distributed representations for group problem solving is used to analyze the distributed representation of the Waitress and Orange problem (Zhang & Norman, 1994), which is an isomorph of the Tower of Hanoi problem (see Hayes & Simon, 1977; Kotovsky, Hayes, & Simon, 1985). Figure 2 shows the Waitress and Orange problem. The task is to move the oranges from one configuration to another, following the three rules stated in Figure 2. Figure 3 shows the problem space of this problem. Each rectangle shows one of the 27 possible configurations of the three oranges on the three plates. The lines between the rectangles show the transformations from one state to another when the three rules are followed.

Figure 3 is a problem space generated by all of the three rules of the Waitress and Orange problem. In general, any subset of these three rules can generate a problem space. Figures 4 shows the problem spaces generated by Rules 1, 1+2, 1+3, and 1+2+3, respectively. Lines with arrows are uni-directional. Lines without arrows are bi-directional. One important point is that these four problem spaces can be held by different individuals.

Figure 5 shows how the three rules of the Waitress and Orange problem is distributed across two individuals. Individual 1 only knows Rules 1 and 3, which generate Individual 1's problem space. Individual 2 only knows Rules 1 and 2, which generate Individual 2's problem space. Although neither of the two individuals alone knows all three rules, they together know all of them. The two problem spaces of the two individuals form the distributed problem space, which is the actual space in which problem solving takes place. The distributed problem space is mapped to the abstract problem space, which is jointly determined by the combined rules of the two individuals (Rules 1, 2, and 3).

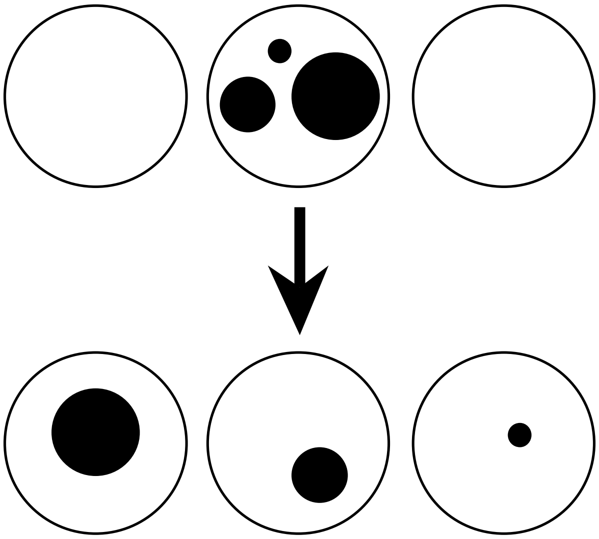

Rule 1: Only one orange can be moved at a time.

Rule 2: An orange can only be moved to another plate on which it will be the largest.

Rule 3: Only the largest orange on a plate can be moved to another plate

Figure 2. The Waitress and Orange problem. The task is to move the three oranges among the three plates from one configuration to another, following the three rules.

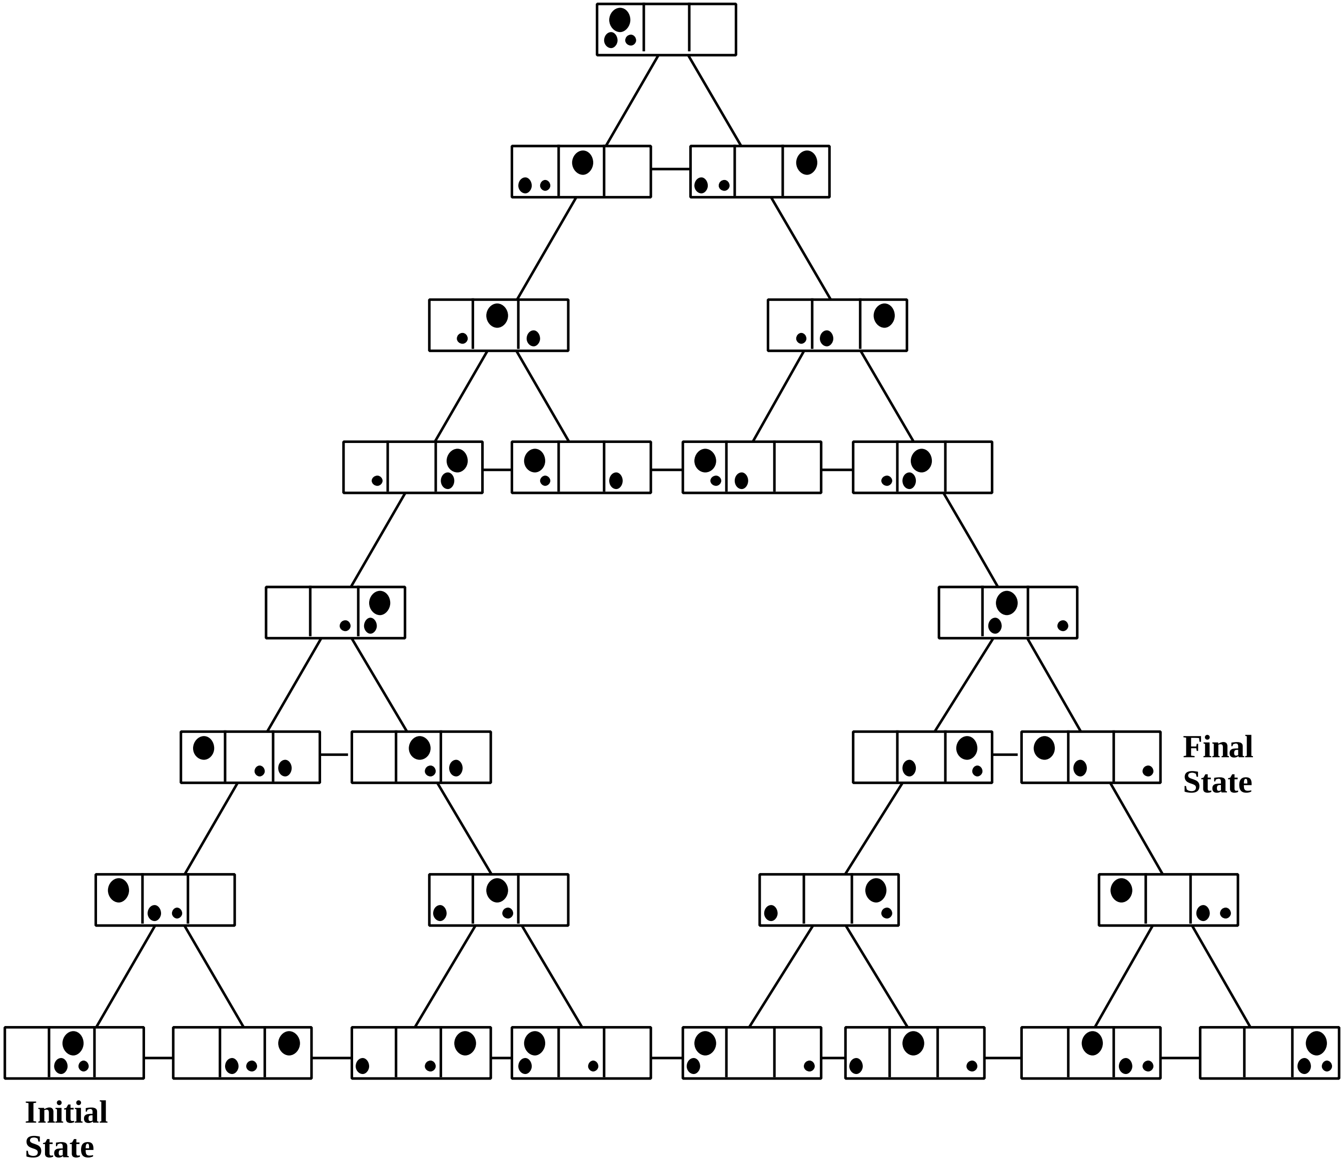

Figure 3. The problem space of the Waitress and Orange problem. Each rectangle shows one of the 27 possible configurations (problem states) of the three oranges on the three plates. The lines between the rectangles show the transformations from one state to another when the three rules are followed.

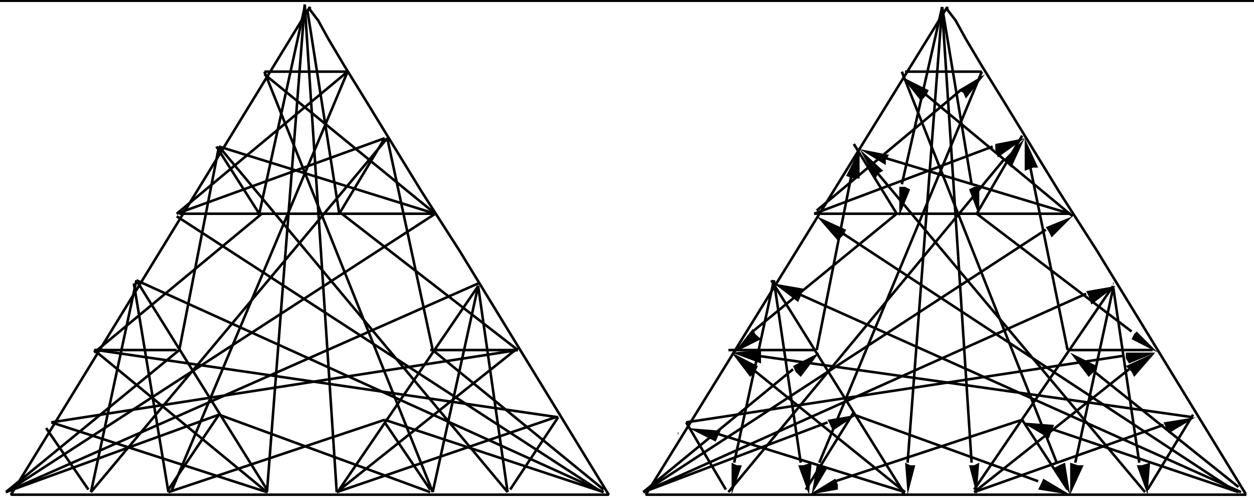

(A) Rule 1 **(B) Rules 1+2**

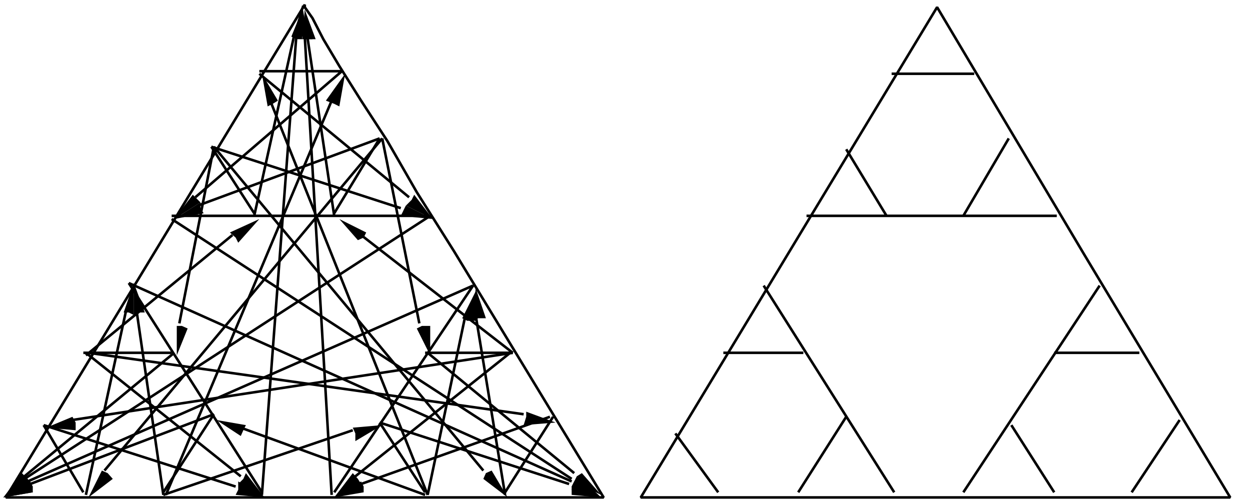

(C) Rules 1+3 **(D) Rules 1+2+3**

Figure 4. Problem spaces generated by four sets of rules. (A) Rule 1. (B) Rules 1+2. (C) Rules 1+3. (D) Rules 1+2+3. They are derived from the problem space in Figure 3. Lines with arrows are uni-directional, and lines without arrows are bi-directional. The rectangles (problem states) are not shown in this figure for clarity.

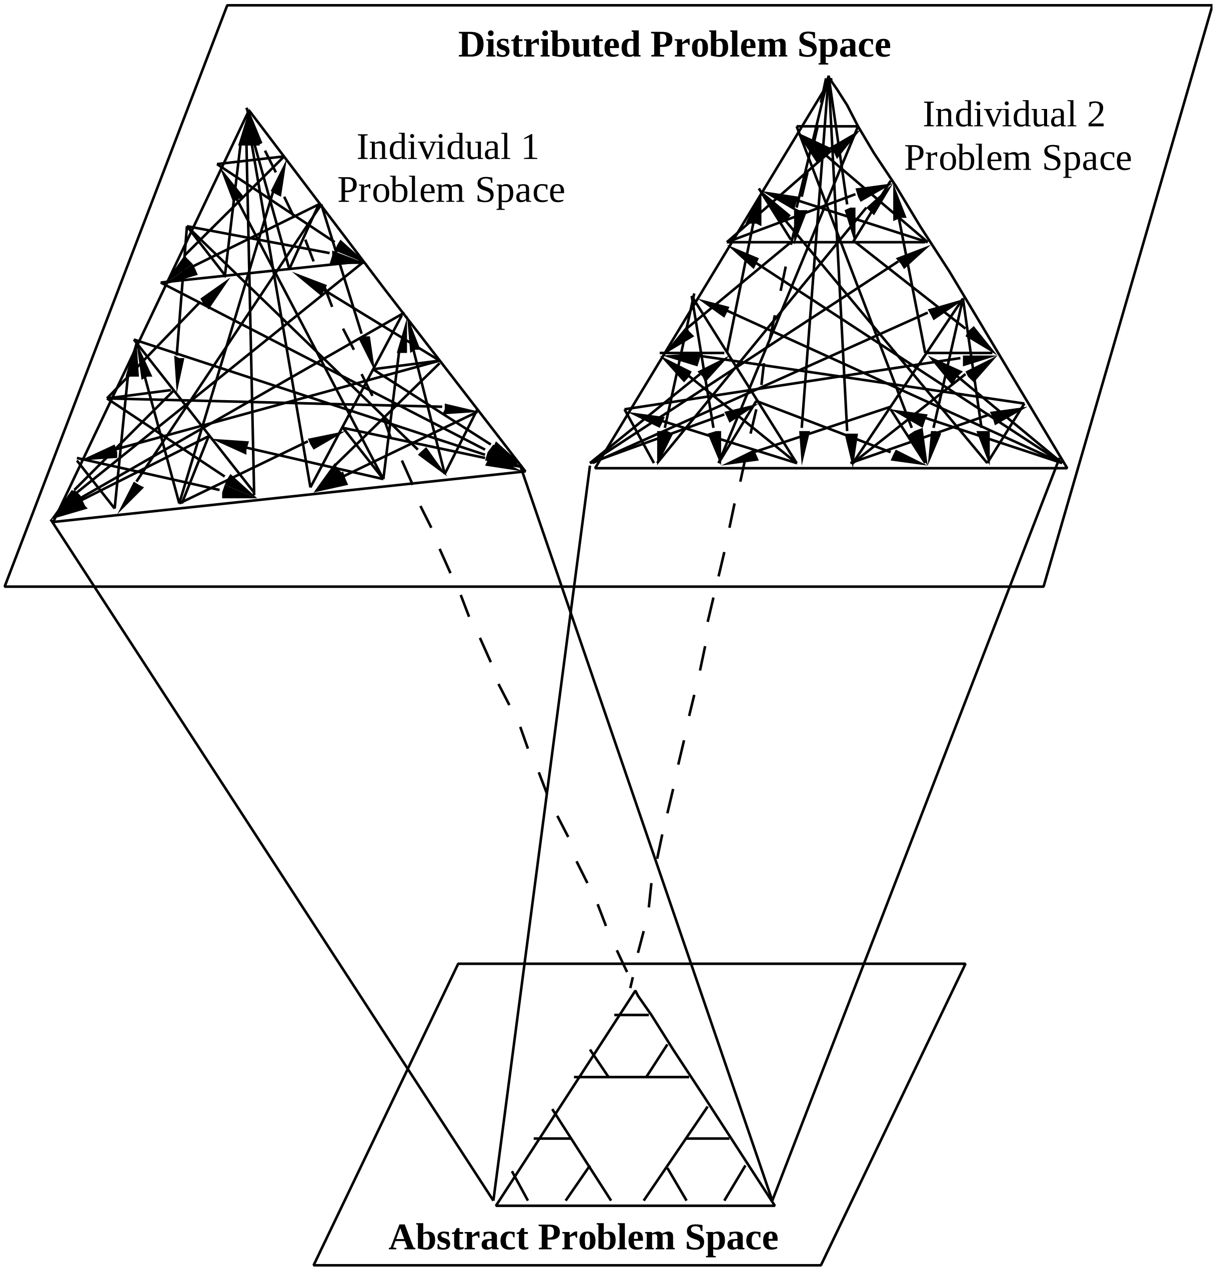

Figure 5. Distributed representation of the Waitress and Orange problem with two individuals. The three rules of the problem are distributed across two individuals: Individual 1 knows rules 1 and 3 whereas Individual 2 knows rules 1 and 2. The two problem spaces of the two individuals form the distributed problem space, which is mapped to the abstract problem space.

Experiment

This experiment examines (a) the effects of distributed representations on group problem solving behavior and (b) the relation between group and individual problem solving behaviors. The Waitress and Orange problem was used to generate four restaurant stories for four conditions (see Figure 6 for an example). The four conditions all had the same abstract structure, i.e., they all had the same set of three rules of the Waitress and Orange problem. Three of the four conditions were two-player tasks in which the three rules were distributed across two individuals in different ways. The fourth condition was a single-player task in which the three rules were represented in one individual.

In Condition *R123-R1*, Subject 1 was given Rules 1, 2, and 3, and Subject 2 was given Rule 1 only. This condition is also called *Expert and Novice* condition because Subject 1 (expert) knew all three rules whereas Subject 2 (novice) only knew one rule. Expert is defined here as a subject who has more knowledge of the task, and novice as a subject who has less knowledge.

In Condition *R12-R13*, Subject 1 was given Rules 1 and 2, and Subject 2 given Rules 1 and 3. This condition is also called *Two Middle-Level Players* condition because neither of the two players knew all three rules but both of them knew two rules.

In Condition *R123-R123*, both subjects were given all three rules. This condition is also called *Two Experts* condition because both subjects knew all three rules.

In Condition *R123*, there was only one subject who knew all three rules. This condition is also called *Single Expert* condition.

Method

Subjects. The subjects were 63 undergraduate students at the University of California, San Diego, who participated in the experiment to get course credits.

Materials. Three plastic orange balls of different sizes (small, medium, and large) were used for the oranges in the restaurant stories. The three plates were porcelain plates.

Design. For the three two-player conditions, two subjects formed a group in each experimental session. Each subject played the role of one of the two waitresses, Kathy and Mary. (For male subjects, John and Mark were used as the names of the two waiters). Which subject played Kathy or Mary was decided randomly. In condition *R123-R1*, Kathy was given Rules 1, 2, and 3, and Mary was given Rule 1 only. In *R12-R13*, Kathy was given Rules 1 and 3, and Mary given Rules 1 and 2. In *R123-R123*, Kathy and Mary were both given all three rules. The instructions for Kathy in *R123-R1* are shown in Figure 6. The instructions for Mary in *R123-R1* were the same as for Kathy, except that only Rule 1 was given to Mary and the names *Kathy* and *Mary* were switched. The instructions for Kathy and Mary in *R12-R13* were the same as for Kathy and Mary in *R123-R1*, except that the rules for Kathy were Rules 1 and 3 and the rules for Mary were Rules 1 and 2. The instructions for Kathy and Mary in *R123-R123* were the same as for Kathy and Mary in *R123-R1*, except that the rules for both Kathy and Mary were Rules 1, 2, and 3.

For the one-player condition *R123*, one subject was run in each experimental session. The instructions for *R123* were identical to the instructions given to Kathy in *R123-R1*, except that before starting the game the subjects were told that they would play the game by themselves without any partners.

There were 9 pairs of subjects for each of the three two-individual conditions and 9 subjects for the one-individual condition. The assignment of subjects to the four conditions was random.

Procedure. For the three two-player conditions, the two subjects in an experimental session were given three minutes to read their instructions in two separated sound-proof booths. After they finished reading their instructions, they were individually asked to recite the rules. Which subject was chosen first to recite rules was decided randomly. If a subject couldn't recite the rules the instructions were read again until all the rules could be recited without any errors. The experimenter then showed the subject two examples of each rule with the real objects that would be used in the experiment. While one subject was being tested by the experimenter, the other subject was instructed to study the rules. After the two subjects understood the game and memorized their rules, they were escorted to an experimental room, where they sat side-by-side in front of a table. Once the experimenter said "go," they started to play the game, making moves in turn. They were forbidden to communicate with each other by any means (verbal or nonverbal) throughout the experiment, except that they could say the word "no" when one person made a move that was not consistent with the rules of the other person. The minimization of communication between the two subjects was to make sure that subjects' behavior was mainly affected by the representational properties of the task.

For the one-player condition *R123*, the procedure was the same as that for the two-player conditions except that each subject played the game alone without a partner.

The initial and final states for all four conditions was identical, which are shown in Figure 3. The final state was shown in a diagram in front of the subjects. Any illegal move that violated one of a subject's own rules was promptly corrected by the experimenter. Subjects' hand actions and speech were recorded by a video camera. The solution time, which was from the time the experimenter said "go" to when the subjects finished the last move, was recorded by a timer synchronized with the video camera.

Results

The average solution times for the four conditions are shown in Figure 7A. The difficulty order was, from hardest to easiest: *R12-R13* > *R123-R1* ≅ *R123* > *R123-R123*. The main effect was significant (F(3, 32) = 8.50, p = 0.0003). Fisher LSD pairwise comparisons showed that except of between R123-R1 and R123 (p = 0.89), the differences between all other pairs were significant (0.00002 < p < 0.05). The average solution steps for the four conditions are shown in Figure 7B. The difficulty order was, from hardest to easiest: *R12-R13* ≅ *R123-R1* ≅ *R123* > *R123-R123*. The main effect was significant (F(3, 32) = 4.73, p = 0.008). Fisher LSD pairwise comparisons showed that the solution steps for *R123-R123* were significantly smaller than those for all other three conditions (0.002 < p < 0.03), which did not differ from each other significantly (smallest p = 0.27). The average errors for the four conditions are shown in Figure 7C. An error was a move that violated one of a subject's own rules. The main effect was not significant (F(3, 32) = 0.70, p = 0.56).

Waitresses 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. Two waitresses, Kathy and Mary, brought all three oranges in one plate and placed them all in front of the middle customer. Because of the exotic style of this restaurant, Kathy and Mary had to move the oranges to the proper customers following a strange ritual. No orange was allowed to touch the surface of the table.

Kathy and Mary moved the oranges in turn. Because they may have different rules in mind, they have to absolutely respect each other's rules. If Kathy makes a move which is not consistent with Mary's rules, Mary should say "NO" and Kathy should try another move until it is consistent with the rules of both of them. Mary has to respect Kathy's rules in the same way. No discussion was allowed between them. The only word they could say was "NO" when a move made by one of them violated the rules of the other person.

Suppose you are Kathy. You have to use only one hand to rearrange these three oranges in cooperation with Mary so that each orange will be placed in the correct plate, following these rules:

- **Only one orange can be transferred at a time.** (*Rule 1*)
- **An orange can only be transferred to a plate in which it will be the largest.** (*Rule 2*)
- **Only the largest orange in a plate can be transferred to another plate.** (*Rule 3)*

How would you do this? That is, you and Mary solve the problem and show the movements of oranges you have to make so that each customer will get his own orange, that is, from the configuration shown in Diagram 1 to that in Diagram 2.

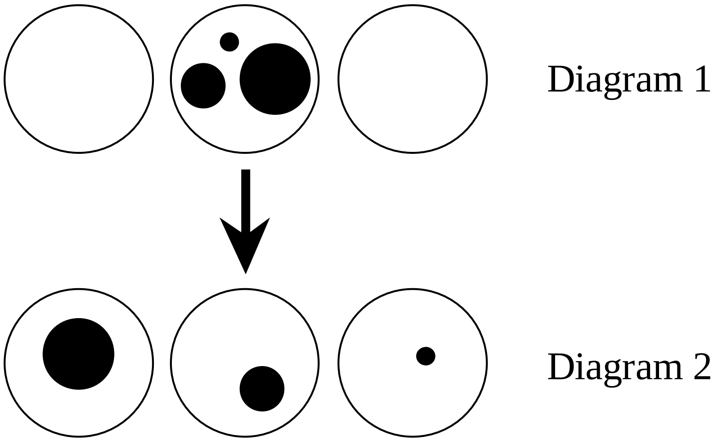

Figure 6. The instructions for Kathy in *R123-R1*.

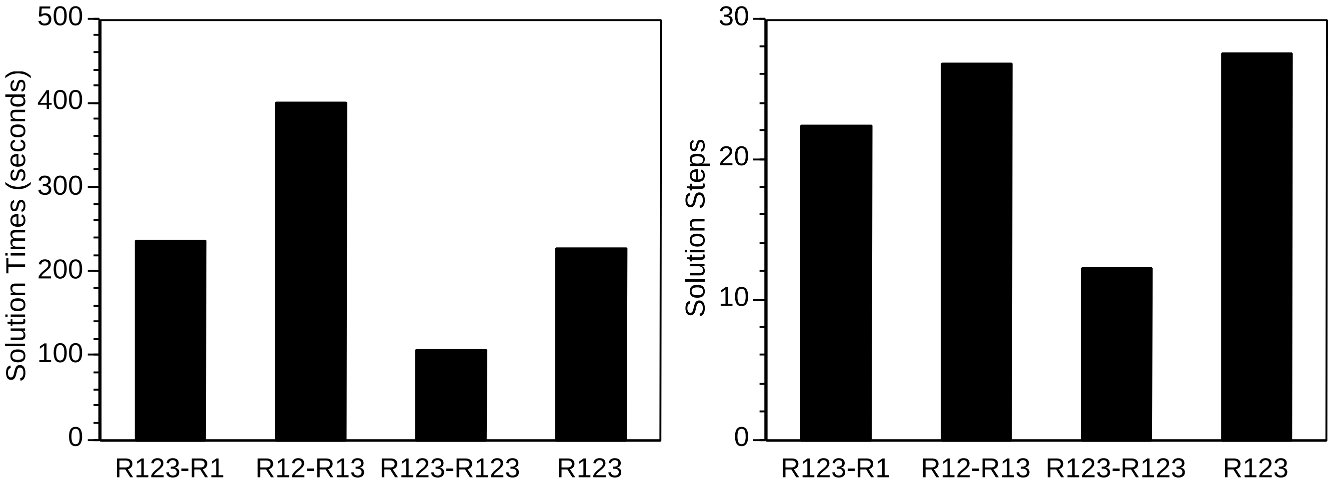

(A) Solution Times

(B) Solution Steps

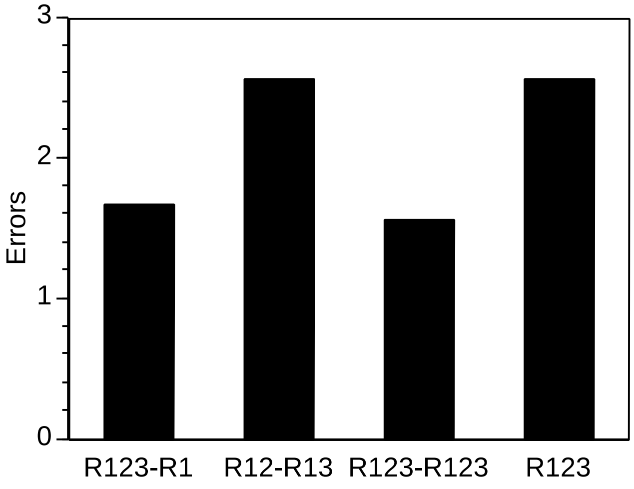

(C) Errors

Figure 7. Solution times, solution steps, and errors for the four conditions.

Discussion and Conclusion

Interpretation of Experimental Results

The abstract structures of all four conditions in the experiment are identical: they all have the same set of three rules. When solution times were used as the difficulty measure, two experts (*R123-R123*) were better than one expert and one novice together (*R123-R1*), which in turn were better than two middle-level players (*R12-R13*). Therefore, when the three rules were distributed across two individuals in different ways, problem solving behavior changed dramatically. When the single-player condition was compared with the three two-player conditions, we found that whether two minds were better than one depended on how representations were distributed across the two minds. Two experts (*R123-R123*) were indeed better than one expert (*R123*). However, one expert and one novice together (*R123-R1*) were not different from one expert (*R123*). In addition, two middle-level players (*R12-R13*) were even worse than one expert (*R123*).

When solution steps were used as the difficulty measure, two experts (*R123-R123*) were still better than one expert and one novice together (*R123-R1*) and better than two middle-level players (*R12-R13*), which, however, did not differ from each other. In addition, two experts (*R123-R123*) were still better than one expert (*R123*). However, two middle-level players (*R12-R13*) were no longer worse than one expert (*R123*).

Combining the above analyses on solution times and steps, we have the following results. First, different distributions of representations across two individuals produced different group problem solving behaviors, even if the abstract structures of the group problem solving tasks were identical. Second, two minds were better than one only when the two minds were both experts. When one or both of the two minds were not experts, two minds were not better and sometimes even worse than one mind.

Communication Hypothesis & Representation Sharing Hypothesis

The differences among the three two-player conditions can be explained by two hypotheses: communication hypothesis for the differences in solution times and representation sharing hypothesis for the differences in solution steps.

The communication hypothesis is that the less communication required among individuals, the better the performance of the distributed system in terms of solution times. This is because less communication takes less time. In *R123-R123*, no communication was necessary. In *R123-R1*, the R123 subject needed time to communicate to the R1 subject about R1 subject's violations of the R123 subject's rules. This is a one-way communication. In *R12-R13*, both subjects needed time to communicate to each other about each other's rule violations. This is a two-way communication. Thus, the order of communication amount is, from most to least: *R12-R13* > *R123-R1* > *R123-R123*. This order corresponds to the difficulty order in solution times.

The representation sharing hypothesis is that the more representation shared among individuals, the better the performance of the distributed system in terms of solution steps. The shared representation between the two individuals in *R123-R1* was Rule 1, that in *R12-R13* was also Rule 1, but that in *R123-R123* was Rules 1, 2, and 3. Thus, the order of the amount of representation sharing is, from least to most: *R12-R13* = *R123-R1* < *R123-R123*. This order corresponds to the difficulty order in solution steps. Although more shared representation can make a task easier, one important point we should keep in mind is that individuals who have little or no shared representation can still work together to perform a task even if none of them have complete representation of the task. In many real world tasks, complementary representation is more common than shared representation in group problem solving tasks, due to many factors such as knowledge specializations.

Communication and representation sharing are properties of groups, not properties of individuals. The above discussion shows that these two group properties can affect group behaviors in a dramatic way. The three two-player conditions all had the same structure. However, due to their differences in communication and representation sharing, they were very different in difficulty levels and group behaviors. This set of results further supports the interactionist view of group problem solving, that is, the behavior of group problem solving is not only affected by properties of individuals but also by properties of groups.

Are Two Minds Better Than One?

The group effectiveness problem was clearly illustrated by the experiment, which showed that two minds could be better than (*R123-R123* > *R123*), not different from (*R123-R1* = *R123*), or even worse than one mind (*R12-R13* < *R123*). We propose a two-factor explanation for this group effectiveness problem. The first factor is the cross-check of loops. From the problem space in Figure 3 we can see that moves can be made within loops. More loops certainly have more moves and thus cause longer solution times. Informal observations showed that single subjects had more loops than paired subjects. This might be because it was more likely for two players to consider two different paths at any given time so that it was less likely for them to be trapped in a loop. The second factor is communication. As shown in last section, more communication needs longer time.

Two minds (*R123-R123*) were better than one mind (*R123*) when the two minds were both experts. This might be because there is cross-check of loops for *R123-R123* but not for *R123* and the communication demands for both *R123-R123* and *R123* were the same (neither required communication). Two minds (*R123-R1*) were not different from one mind (*R123*) when only one of the two minds was an expert. This might be because the benefit of the cross-check of loops canceled the cost of communication for R123-R1. Two minds (*R12-R13*) were worse than one mind (*R123*) when neither of the two minds was an expert. This might be because the benefit of the cross-check of loops was smaller than the cost of communication for R12-R13 (two-way communication).

Implications

The experimental task used in the present study is simple, well-defined, and explicit, whereas real world tasks such as those performed by aircrew members are usually complex, ill-defined, and vague (e.g., Foushee & Helmreich, 1988; Hutchins, 1995; Norman, 1991). Nevertheless, the regularities found in real world tasks such as group properties emerging from interactions and the group effectiveness problem were also observed in the experimental task.

The present study only examined the representational properties in group problem solving. This is not denying the important roles of social, personality, and emotional factors, which are often crucial. As discussed by Foushee & Helmreich (1988), these non-cognitive factors were responsible for several tragic airplane crashes (NTSB, 1972, 1979, 1982). Field studies of group problem solving usually consider non-cognitive as well as cognitive factors. Laboratory studies in the future should focus more on multiple factors than single factors.

The framework of distributed representations for group problem solving is not a process model that has explicitly specified mechanisms. However, as a general principle, it states that the representation of a group problem solving task is distributed across individual representations, which jointly represent the abstract structure of the task. As a methodology, it demands (a) the consideration of the individual representations in a group problem solving task as a distributed representation system, (b) the explicit decomposition of a group problem solving task into its individual representations, (c) the identification of the abstract task structure and its relation to each of the individual representations, and (d) the emphasis on the interactions among individual representations. Although this framework was developed for group problem solving, the notion of distributed representations is a general idea that can be potentially applied to other domains as well, such as group development, shared cognition, and team mental models (Crespin, 1996; Klimoski & Mohammed, 1994).

References

Bandura, A. (1986). *Social foundations of thought and action: A social cognitive theory*. Englewood Cliffs, NJ: Prentice-Hall.

Crespin, T. (1996). *Cognitive convergence among members of newly formed groups*. PhD dissertation. Department of Psychology, The Ohio State University.

Foushee, H. C. & Helmreich, R. L. (1988). Group interaction and flight crew performance. In E. L. Wiener & D. C. Nagel (Ed.), *Human factors in aviation.* San Diego: Academic Press.

George, J. M. (1990). Personality, affect, and behavior in groups. *Journal of Applied Psychology, 75*, 107-116.

Guzzo, R., & Shea, G. P. (1993). Group performance and intergroup relations in organizations. In Dunnette & Campbell (Eds.), *The thinking organization: Dynamics of organizational social cognition*. San Francisco: Jossey-Bass.

Hackman, J. R. & Morris, C. G. (1975). Group tasks, group interaction process, and group performance effectiveness: A review and proposed integration. In L. Berkowitz (Ed.), *Advances in experimental social psychology*

(Vol. 8, pp. 45-99). New York: Academic Press.

Hare, A. P. (1972). Bibliography of small group research: 1959-1969. *Sociometry, 35*, 1-150.

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

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

Hutchins, E. (1994). *Cognition in the Wild.* Cambridge, MA: The MIT Press.

Hutchins, E. (1995). How a cockpit remembers its speed. *Cognitive Science, 19,* 265-288.

Janis, I. L. (1972). *Victims of groupthink: A psychological study of foreign policy decisions and fiascos*. Boston: Houghton Mifflin.

Klimoski, R., & Mohammed, S. (1994). Team mental model: Construct or metaphor? *Journal of Management, 20,* 403-437.

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

Lanzetta, J. T. & Roby, T. B. (1960). The relationship between certain group process variables and group problem solving efficiency. *Journal of Social Psychology, 52,* 135-148.

McNeese, M. D., Zaff, B. S., & Brown, C. E. (1992). Computer-supported collaborative work: A new agenda for human factors engineering. *Proceedings of the IEEE National Aerospace and Electronics Conference (NAECON), (681-686)* . Dayton, OH: IEEE, Aerospace and Electronic Systems Society.

Norman, D. A. (1990). Cognitive artifacts. In J. M. Carroll (Ed.), *Designing interaction: Psychology at the human-computer interface.* New York: Cambridge University Press.

Norman, D. A. (1991). Cognitive Science in the cockpit. *CSERIAC Gateway, 11* (2), 1-6.

Norman, D. A. (1993). *Things that make us smart*. Reading, MA: Addison-Wesley.

NTSB. (1972). *Aircraft accident report. Eastern Air Lines, Inc. L-1011, N310EA, Miami, Florida, December 29, 1972*. (Report No. NTSB-AAR-73-14). National Transportation Safety Board. Washington, DC.

NTSB. (1979). *Aircraft accident report. United Airlines, Inc. McDonnell-Douglas DC-8-61, N8082U, Portland, Oregon, December 28, 1978*. (Report No. NTSB-AAR-80-1). National Transportation Safety Board. Washington, DC.

NTSB. (1982). *Aircraft accident report. Air Florida, Inc. Boeing B-737-222, N62AF, Collision with 14th Street Bridge, Bear Washington National Airport, Washington, D. C., January 13, 1982.* (Report No. NTSB-AAR-82-8). National Transportation Safety Board. Washington, DC.

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

Steiner, I. D. (1972). *Group process and productivity*. New York: Academic Press.

Wegner, D. M. (1987). Transactive memory: A contemporary analysis of the group mind. In B. Mullen & G. R. Goethals (Eds.), *Theories of group behavior*. New York: Springer-Verlag.

Zhang, J. (1996). A representational analysis of relational information displays. *International Journal of Human-Computer Studies, 45,* 59-74.

Zhang, J., & Norman, D. A. (1994). Representations in distributed cognitive tasks. *Cognitive Science, 18* (1) , 87-122.

Zhang, J. & Norman, D. A. (1995). A representational analysis of numeration systems. *Cognition, 57,* 271-295.