

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Learning and Problem Solving Under a Memory Load

Permalink

<https://escholarship.org/uc/item/5fb552wf>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 14(0)

Authors

Reber, Paul J.

Kotovsky, Kenneth

Publication Date

1992

Peer reviewed

Learning and Problem Solving Under a Memory Load

Paul J. Reber
Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15217
pr18@andrew.cmu.edu

Kenneth Kotovsky
Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15217
kotovsky@psy.cmu.edu

Abstract

A problem solving experiment is described where the difficulty Ss experienced in solving a particular puzzle is manipulated using a dual task paradigm. Although Ss show impaired performance solving the puzzle the first time, performance improves considerably on a second trial and Ss are not impaired by a second trial memory load. In spite of the improvement in performance, Ss are unable to report virtually any information about the problem or their solution strategies. A model is presented that describes the pattern of performance across the levels of memory load and across the two trials. The theoretical implications of this model are discussed.

Introduction

Problem solving has traditionally been thought of as search through a problem space (Newell & Simon, 1972). Recent work with "isomorphic" puzzles which have identical underlying problem spaces but different surface features has shown that the working memory demands of the surface features plays an important role in predicting how difficult a particular problem will be (Kotovsky, Hayes & Simon, 1985; Kotovsky & Simon, 1990; Kotovsky & Kushmerick, 1991).

The following experiment demonstrates that the difficulty of a particular puzzle can be manipulated by reducing the working memory available for problem solving using a dual task paradigm. We find that the performance is impaired in proportion to the demands of the secondary task. However, when repeating the problem solving task, the memory load ceases to interfere with problem solving.

The pattern of results indicates that it is a learning process associated with problem solving that is impaired by the secondary task. However, an interesting problem arises in attempting to assess what and how the subjects are learning while problem solving as the subjects are

unable to describe what they know about the puzzle or how they solved it.


A reasonably simple mathematical model is presented to attempt to bring together some of the phenomena observed in this experiment. Focusing on weak methods and modeling the learning as an incremental function, the model gives a reasonable first-order account of the structure of the problem solving process for this puzzle.

Method

Seventy-six Carnegie Mellon University undergraduates were asked to solve an isomorph of the "Chinese Ring Puzzle" (Kotovsky & Simon, 1990) called the Balls and Boxes puzzle while performing a secondary task designed to impose a load on working memory. The Balls and Boxes puzzle was presented via computer and is described below. While solving the puzzle, subjects were also asked to listen to a tape containing a stream of letters at a rate of one letter every three seconds with beeps mixed in occasionally. Subjects were divided into four groups, a control group and three different levels of memory load. The three memory load groups were instructed to listen to the letters and remember either the last letter heard (group 1-back), the last two letters heard (group 2-back) or the last three letters heard (group 3-back). Since the tape contains a continuous stream of letters, each time a new letter is heard by the subject, the set of letter(s) to be remembered changes. These groups of subjects were instructed that when they heard a beep they were to write down the first letter of the letter set being remembered, i.e. subjects remembering the last letter wrote this at the beep, subjects remembering two letters wrote the letter before the last letter heard, etc. The control group was instructed to ignore the letters but listen for the beeps and when a beep was heard, write down a random letter. To be sure that the subjects performed the secondary task as well as possible, the importance of listening to the tape carefully was stressed with all subjects and the experimenter monitored subjects performance to be sure they wrote down a letter whenever a beep was heard.

An experimental session consisted of brief instruction on the puzzle and secondary task and two solutions of the puzzle. Between solutions, subjects

This research was supported by NIMH training grant 5T32 MH-19102. The authors would like to extend thanks to Gerry Starret who aided greatly in collecting the experimental data; also Chris Schunn and Frank Ritter for insightful comments that facilitated development of the model.



Start state



After moving rightmost ball

Figure 1: Appearance of the Puzzle to Subjects

were asked to write "how they solved the puzzle, how the puzzle works and especially anything they could say that would help somebody else solve the puzzle." Only two levels of memory load were used on the second trial, the control condition and the two-letter memory load condition.

The Balls and Boxes Puzzle

The Balls and Boxes puzzle is an isomorph of the "Chinese Ring Puzzle" which is a particularly difficult puzzle studied by Kotovsky & Simon (1990). The Balls and Boxes isomorph was designed to be a simpler version of the puzzle where the surface structure was modified so that all legal moves can easily be seen and the operators were "digitized" to reduce the working memory load associated with moving within the problem space. This version of the puzzle was found to be much easier to solve, generally taking 5-10 minutes. The puzzle was presented on a MicroVAX II computer

and is shown in Figure 1. This puzzle consists of five balls and five boxes. The object is to get all five balls out of their boxes. A ball may only be moved if its box top is open. As balls are moved in and out of their boxes, the box tops open and close according to a rule that defines the problem. The rule dictates that a ball may be moved (and hence the box top is open) if and only if the ball immediately to its right is in its box and all other balls to the right are out of their boxes. For the subjects, the trick to solving the puzzle is to figure out how to move the balls to get the right boxes to open up so that you can get all the balls out of their boxes.

The problem space, shown in Figure 2, is rather small, containing only 31 possible states, but the starting position was chosen to be 21 moves from the goal (21 moves is the minimum number required to solve the puzzle). Additionally, the problem space is linear, meaning that from any place in the problem space (except the top state) there are exactly two moves that can be made, one leading toward the solution (moving toward lower numbered states in Figure 2) the other leading directly away (moving toward the higher numbered states). Hence, after making a move, the only choices one has are to undo that move or make a new move. If one never retracts or undoes a move when it is not necessary, then one is guaranteed to easily solve the puzzle. However, it is extremely rare for a subject to do this when attempting to solve the puzzle for the first time.

Results

The effect of the memory load the first time the puzzle is solved is clearly seen in Figure 3. As the memory load increases, the number of moves required to solve the puzzle increases. A regression analysis shows this increase is significant ($\beta=25.82$, $F=10.82$, $p<0.01$). The results from the second solution of the puzzle are shown in Figure 4. This figure is shown on the same




Figure 2: Problem Space of the Balls and Boxes Puzzle

scale as Figure 3. It is clear that there is significant speed up in that the number of moves to solve the puzzle drops dramatically from trial 1 to trial 2 (paired t-test, $t=4.57$, $p<0.001$). These results are also broken down by memory load on trial one to demonstrate that there is no difference between the groups in the amount learned on the first trial. It is clear that the memory load did not impair performance on the second trial. This is verified by an ANOVA of trial 2 performance by trial 1 condition and trial 2 condition where neither condition nor the interaction are significant ($F=1.5$, $p>0.22$, $F=1.85$, $p>0.17$, $F=1.76$, $p>0.33$ for the effects of trial 1 condition, trial 2 condition and their interaction respectively).

Examining the progress of a single subject on the puzzle shows that the process of solving the puzzle can be broken down into two phases: *exploratory* and *final path*. The *exploratory* phase contains most of the moves made by the subject and little or no actual progress toward the goal is made during this phase. The *final path* is a sequence of rapid totally error-free moves directly toward the goal. Progress made by subjects over the course of solving the puzzle on the first trial can be seen in Figure 5 (the x-axis is move number and the y-axis measures distance from the goal). The average length of the final path was 18.5 moves and does not vary across groups. The second time the puzzle is solved, 66 of 76 subjects show both exploratory phase and final path behavior while 10 subjects solved the puzzle perfectly the second time (immediate final path). The average length of the final path on the second trial was 19 moves.

The written protocols collected from each subject in between trials were analyzed and rated on a five point scale with 1 indicating that the description had no relevant information about the puzzle; a rating of 2 indicates very little information in the description, i.e. one statement of value; and so on up to 5, indicating that the subject provided a complete description of how to solve the puzzle. The median value of these ratings was 1.5. The average rating did not vary across groups indicating that all groups were equally poor at describing the structure of the puzzle or the strategies that they used to solve it.

Most of the advice subjects gave was very general in nature. The most common statement in the protocols (29 occurrences in 76 protocols) was general advice to remember the combinations or patterns that lead to different boxes opening. In spite of this advice, virtually no subjects described any of these patterns or combinations correctly. Other general advice such as "Be systematic" and "Use trial and error" also occurred frequently (20 times). Subjects did seem to be aware of certain aspects of the problem space: 21 subjects reported that the left ball was hardest, 12 noted that it was important to work from left to right and 10 realized that it was occasionally necessary to put balls back into boxes. No other statement occurred as often as 10 times. Only one subject stated that not undoing the previous move would be a good strategy and two others hinted that not reversing is important. In summary, the protocols were surprisingly uninformative, especially given the fact that most subjects went on to solve the puzzle quite rapidly on the second trial.




Figure 3: Number of Moves to Solve the Puzzle on Trial 1




Figure 4: Number of Moves to Solve the Puzzle on Trial 2

The protocols present an interesting contrast to the final path behavior of the subjects. The sudden shift to error-free solution initially seems to suggest insight, yet the protocols suggest that subjects have not settled on a strategy even as simple as consistently avoiding reversals. The fact that most subjects also show exploratory behavior on the second trial also indicates that while subjects know something about the puzzle after the first solution, they don't understand it completely.




Figure 5: Sample move record

A Model

Although the pattern of results across the two trials is counter-intuitive, a relatively simple model provides a reasonable approximation of the subjects' behavior. The model is based on two basic problem solving strategies: (1) hill-climbing, that is preferring to remove balls from their boxes; and (2) no-reversal, a general preference not to undo the most recent move. Of course, since the problem space is linear, relying solely on the second strategy would lead to optimal solution of the puzzle. Although hill-climbing based on the increasing number of balls out of their boxes may initially appear useful, reliance on this strategy actually impedes solution of the puzzle. The existence of the barrier areas suggest that at least at the beginning, subjects appear to rely on hill-climbing.

These strategies come into conflict on moves where balls must be replaced to make progress. If there is a decision between two balls is to be replaced (e.g. state 24, Figure 2), then the model will prefer not to make the reversal. However, when the choice is between taking a ball out that was just put in and putting a different ball in, then there is a higher probability of a reversal. An example of this situation would be arriving at state 25 (see Figure 2) from state 24 (the last move was to place the rightmost ball in). The hill climbing strategy indicates taking the rightmost ball out (and returning to state 24) while the no-reversal strategy

indicates continuing on to state 26 and putting the secondmost ball from the right back into its box.

The relationship between these two strategies is controlled by a parameter, bp (for backup-penalty) describing the relative importance of the no-reversal strategy. The value of the hill-climbing strategy is held constant at +1 for removing a ball and -1 for replacing a ball. To select a move, the values of the two possible moves are calculated. The value of a move is the +1 if a ball is being moved out, -1 if a ball is being replaced. The move that reverses or undoes the previous move is penalized by an additional $-bp$. A small amount of normally distributed noise (mean=0, sd=2) is added to each value. The noise makes the model's behavior more variable and is meant to capture factors that affect subjects' move choice that are not captured by the model. The higher of these values is selected and this move is made. This process repeats until the goal state is reached.

When $bp=2$, these parameters exactly oppose each other in the difficult states described above (choosing between replacing a ball and removing a ball that was replaced on the previous move). When bp is less than 2, the model relies more on hill-climbing and when bp is greater than 2, the model makes fewer reversals. Since fewer reversals leads to more rapid problem solving, average number of moves to solve the puzzle decreases as bp increases.

When bp is 2 (or close to 2), the model essentially "wanders" around the problem state. At most states the model prefers moves that avoid reversals, but a number of situations provide problems (generally where two balls must be replaced in succession). These states correspond to the barrier areas observed in the subjects' behavior. The fit between barrier areas can be assessed by tracing subjects' move records with the model and at each point recording whether the model predicts a high, medium or low chance for reversal. This is then compared to the actual rate of reversal for each subject in each of these categories. The mean reversal rates across subjects was 9.7%, 13.6%, and 14.2%, for the states where the model predicted low, medium, and high reversal rates respectively. A repeated measures ANOVA across the reversal rates in each of the three categories for all 76 subjects shows a significant within subject effect ($F=13.6$, $p<0.001$). This indicates that subjects are indeed making reversals at different rates in the three categories and that the hill-climbing nature of the model reflects this strategy in the subject data. From this we can conclude that the *exploratory* behavior of the model is similar to the *exploratory* behavior shown by subjects.

In order to model the data pattern for the experiment across the groups and across trials, the model learns. In the subjects' data, the existence of the *final path* behavior initially appears to suggest insightful behavior, but the stunning lack of information in the protocols strongly suggests otherwise. Correspondingly, the model is constructed using a

simple, incremental function to model learning effects. The theoretical implications of this decision are discussed below.

The learning across the two trials is modeled by increasing bp incrementally throughout the first trial solution. This has the effect of reducing the model's reliance on hill-climbing over the course of solving the puzzle. Since these strategies compete with each other, the model cannot disambiguate between learning that reduces the perceived value of hill-climbing or increases the perceived value of not making a reversal.

Increasing the value of bp increases the likelihood of achieving a final path sequence of moves and solving the puzzle. A higher value of bp speeds problem solution but does not guarantee immediate final path. Hence, carrying the value of bp across from the end of trial one to trial two reduces the number of moves needed to solve the puzzle but does not guarantee immediate final path.

The rate at which bp increases also affects the average number of moves it takes the model to solve the puzzle on the first trial. The difference across the memory load conditions can therefore be modeled by having bp increase at different rates for the different memory load groups. Since all groups perform with similar efficiency on the second trial, it follows that the value of bp carried over to the second trial should be approximately equal for all groups. Accordingly, bp increases incrementally, per move, according to the following general formula:

$$(1) \quad bp_t = bp_0 + bp_{max} \cdot ce^{-\alpha t}$$

This formula describes exponential learning from the initial bp_0 toward some maximum value of bp , bp_{max} . The overall shape of the learning curve is determined by the quantities c and α . The parameter t , indexes the number of moves made. Several forms of the learning functions were considered. The function used mirrors the learning function previously used to describe classical conditioning (Rescorla & Wagner, 1972).

The parameters bp_0 and bp_{max} and c are the same for all runs of the model. Differences across the groups are captured by using different values of α . The higher memory load groups have smaller values of α , indicating that it takes more moves to reach the same level of bp . This implies that the groups are learning at different rates toward the same endpoint.

These parameters were fit to the data pattern with $bp_0=1.5$, $bp_{max}=2.6$, $c=2.0$, $\alpha=0.0428$, 0.0331 , 0.0228 , 0.0120 for the control, 1-back, 2-back and 3-back groups respectively. These parameter settings start the model out with a reliance on hill-climbing, gradually shifting to relying more on not reversing any moves. The performance of the model on the first and second trial is shown in Figures 7 and 8. The same regression analyses used for the subjects show a significant effect of memory load on the first trial ($\beta=37.4$, $F=56.9$, $p<0.001$) and no effect of either first or second trial memory load on the second trial ($F<1$ in all cases). The model also shows final path behavior with average final path being 22.9 moves on the first trial and 26.6 moves on the second trial.




Figure 7: Model data, Trial 1




Figure 8: Model data, Trial 2

Discussion

The experiment demonstrated that it is possible to manipulate problem difficulty by reducing available working memory using a dual task paradigm. The fact that the effect of the first trial memory load does not carry over into the second trial and that the second trial memory load does not impair second trial performance indicates that a learning process is impaired on trial one. It follows that subjects from different groups are at essentially the same level of knowledge going into the second trial. The knowledge they have allows them to solve the puzzle much more efficiently on the second trial; but subjects in the high memory load groups on trial one take much longer to reach this level of knowledge.

Interestingly, the knowledge that is acquired about the puzzle is not easy for the subjects to communicate. Other work with this puzzle has similarly found that subjects are unable to describe the puzzle retrospectively and requiring subjects to give a concurrent protocol has also been unenlightening (Reber & Kotovsky, in preparation).

The inability of subjects to report what they are learning or describe the structure of the puzzle after they have demonstrated that they can solve it implies that they are learning to solve the puzzle implicitly or automatically. It follows then that it is critically important to understand when these processes are an integral part of problem solving and how this type of learning is affected by environmental factors such as an additional load on working memory. The research presented here demonstrates that the learning process is impaired by an external memory load, but later application of this knowledge is not impaired by a memory load (witness the lack of an effect of memory load on the second trial).

We also present a model that captures the overall data pattern fairly effectively. The reliance on hill-climbing with a general tendency not to reverse moves generates "search" behavior like the *exploratory* phase found in the subject data. The asymptotic nature of the learning function of the model guarantees that all runs of the model will be at essentially the same "level of knowledge" after the first trial. Changing the learning rate causes the model to take different amounts of time to reach this asymptotic level and hence take different number of moves to solve the puzzle. Thus the model shows differential performance on the first trial based on the learning rate, but performance on the second trial is not affected by changing the learning rate. The match of the model to the subject data further strengthens our claim that the memory load impairs the process of learning to solve the puzzle.

The model also provides a convincing demonstration of how a gradual, incremental learning function can capture the apparently sudden shift in subject behavior from exploratory to final path behavior. The success of the model provides additional support to the idea that subjects may be learning to

solve the puzzle by some implicit, incremental function that gradually acquires enough information about the problem space to lead to improved performance.

While the model has been tailored to capture the mathematical structure of the data pattern, it also highlights some important theoretical points. First, the model provides support for the notion that subjects rely heavily on hill-climbing with some attention to not reversing the last move made. Second, the incremental learning function shows how a gradual learning process can result in *final path* behavior. The shape of the learning function further demonstrates how the effect of the memory load can be eliminated across the two trials. Although the model is designed to mathematically describe the data, rather than provide an explicit account of the underlying cognitive processes, it heavily constrains the development of a process model and expands our understanding of the process of solving this puzzle.

Further experimental work is underway that is aimed at uncovering the knowledge acquired on the first trial, together with constraints imposed by the model presented here should allow us to develop a more detailed model of the cognitive processes involved in this and other related tasks. This research represents progress toward a good understanding of the interaction between working memory and learning in problem solving that promises to greatly enhance our overall understanding of problem solving processes.

References

- Kotovsky, K., Hayes, J. R. & Simon, H. A. (1985). Why are some problems hard: evidence from the tower of hanoi. *Cognitive Psychology*, 17, 248-294.
- Kotovsky, K. & Simon, H. A. (1990). Why are some problems really hard: explorations in the problem space of difficulty. *Cognitive Psychology*, 22, 143-183.
- Kotovsky, K. & Kushmerick, N. (1991). Processing constraints and problem difficulty: a model. *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society*. pp 790-795.
- Newell, A. & Simon, H. A. (1972). *Human Problem Solving*. Englewood Cliffs, NJ: Prentice Hall.
- Reber, P.J. & Kotovsky K. (in preparation). Implicit and explicit mechanisms of learning and problem solving.
- Rescorla, R. A. & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current theory and research* (pp 64-99). New York: Appleton-Century-Crofts.