

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Adaptive Selection of Problem Solving Strategies

Permalink

<https://escholarship.org/uc/item/7mx0j1n1>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 23(23)

ISSN

1069-7977

Authors

Fum, Danilo
Del Missier, Fabio

Publication Date

2001

Peer reviewed

Adaptive Selection of Problem Solving Strategies

Daniilo Fum* (fum@univ.trieste.it)
Department of Psychology, via S. Anastasio 12
Trieste, I-34134 Italy

Fabio DelMissier (delmissier@univ.trieste.it)
Department of Psychology, via S. Anastasio 12
Trieste, I-34134 Italy

Abstract

The issue of strategy selection in solving the Tower of Hanoi (TOH) problem is investigated by focusing on the critical issues of whether the selection process is contingent and adaptive. The results of an experiment in which participants solved a series of different four-disk TOH problems under instructions requiring accuracy maximization vs. effort minimization are presented. A computer simulation, comparing a number of known strategies to the experimental data, has been carried out to try to identify the strategies used by the participants. The findings support the hypothesis of adaptive and contingent strategy selection in the TOH domain.

Introduction

Much work in the problem solving arena has dealt with the Tower of Hanoi (TOH)— considered as a typical well-structured problem— producing important theoretical and empirical results. Researchers have discovered interesting phenomena and tried to provide explanations for them. Several solution strategies have been described (Simon, 1975), and various models have been proposed to simulate human performance on this task (Karat, 1982; Ruiz & Newell, 1989; Anderson, Kushmerick & Lebiere, 1993; Anderson & Lebiere, 1998; Altmann & Trafton, 2000). Detailed accounts of learning how to solve the TOH on a trial-by-trial basis (Anzai & Simon, 1979) have been put forward together with hypotheses concerning the strategies and the heuristics people seem to learn in successive attempts to solve the problem (VanLehn, 1991).

Despite these achievements, many issues are still unresolved and many topics are currently investigated. Two recent examples involve the role of goal encoding and retrieving in memory (Altmann & Trafton, 2000), and the possible use of active planning to avoid previously visited states (Davies, 2000).

Given that different models and strategies have been proposed in different experimental settings, it seems important to try to identify the factors affecting the selection of solution strategies in this domain.

We propose the hypothesis that strategy selection in the TOH is a contingent process, i.e., it is sensitive to task and contextual factors. Following a widely accepted idea about human problem solving (Simon, 1975; Anderson, 1990; Christensen-Szalanski, 1998) and decision-making (Payne, Bettman & Johnson, 1993), it is further hypothesized that strategy selection is adaptive. Given a specific task and context, it is functional to the achievement of a good trade-off between accuracy and cognitive effort (Christensen-Szalanski, 1998; Fum & DelMissier, 2000).

These two strategy-related questions (i.e., is the process of strategy selection contingent? is it adaptive?) are the main topics of this work. In the paper we briefly discuss some issues concerning research on the TOH strategies. Then we present the results of an experiment in which participants solved a series of different four-disk TOH problems under instructions requiring accuracy maximization vs. effort minimization. A computer simulation, comparing several solution strategies to the experimental data, has been carried out to try to identify the strategies used by participants in the two instruction groups.

Issues on Strategy Research

Research on strategies in TOH, and related problem solving tasks, must deal with several theoretical and empirical issues.

A first issue concerns identifiability (Anderson, 1990): patterns of behavioral data are used as a trace to induce the existence of a given strategy, but in many cases the data do not allow discriminating among distinct models of strategic behavior. In our specific domain, however, very few attempts (an exception being represented by Altmann & Trafton, 2000) of directly comparing different models on the same data set have been done.

Other theoretical problems deal with the underspecification and the low generalizability of some of the proposed strategies. With underspecification we mean

*The order of authorship is arbitrary; each author contributed equally to all phases of this project.

the fact that the description of a strategy does not allow a unique identification of the move to be done for every problem state. With low generalizability we mean the fact that the proposed strategy results ad hoc and cannot be extended to deal with some classes of TOH problems people are able to solve.

A further theoretical limitation is constituted by the fact that some strategies are willfully optimal (Anderson & Lebiere, 1998), while people seldom achieve such a brilliant performance (Goel & Graffan, 1995; Miyake et al., 2000; Karat, 1982).

On the empirical side, there is the problem of the intrusiveness of the methods utilized to identify the existence of a given strategy. Verbal protocols, for instance, (Anzai & Simon, 1979; Van Lehn, 1991) have proved to be a useful exploratory tool, but there is evidence (Stinnesen, 1985; Ahlum-Heath & DiVesta, 1986) that participants verbalizing during the task perform differently from participants that do not verbalize. The very use of verbal protocols could prompt the adoption of different solution strategies.

A related issue deals with the suggestiveness of the experimental instructions. For instance, Anderson, Kushmerick & Lebiere (1993) gave hints that deliberately encouraged the adoption of a particular strategy. The generalizability of their model is, therefore, directly related to the way the same strategy is spontaneously adopted by the participants when no hints are given.

Another concern is constituted by the fact that strategy selection in the TOH has often been studied by having people perform many trials over the same problem. In this way it cannot be excluded that the improvement in the participants performance could be attribute to rote memorization instead of genuine learning. To control for this factor, Anderson, Kushmerick & Lebiere (1993) presented a wider range of problems to their participants preventing them from evolving special-case strategies.

In our experiment we investigated a factor that could possibly affect the adoption of different solution strategies, and we ran a simulation study to try to identify them. To do this, we had to make some underspecified strategies computationally workable by postulating a few additional assumptions. We concentrated our attention on general strategies— i.e. on strategies capable of solving problems put not only in their standard (i.e., tower-to-tower) form— and on strategies that do not prescribe an optimal solution. Furthermore, we refrained to force participants to justify and comment on their moves, and carefully avoided suggesting any specific solution procedure. Finally, we utilized a set of different problem types.

The Experiment

The main goal of the experiment was to test the hypotheses of contingent and adaptive strategy selection. We manipulated the experimental instructions to modify the importance participants gave to the distinct goals of accuracy maximization vs. effort minimization.

According to the contingent and adaptive hypothesis, we expected to find a rational use of different strategies in different experimental groups. The strategies used by participants in the accuracy group should increase the accuracy of the solutions by paying a higher temporal cost. The strategies used in the effort group should yield effort savings but less accurate solutions.

Method

Participants The participants were 34 undergraduates students, aged between 18 and 24. None of them was suffering from any perceptual, cognitive or motor deficiency. The sample was balanced for gender. All the participants had a basic familiarity with computers and were able to use the mouse.

Procedure Participants read an instruction document that explained the basic rules of the TOH, showed the interface used by the computer program, and described how to use it. The instructions required the participants to solve the problem "in the fewest possible number of moves" or "in the shortest possible time", depending on the group (accuracy vs. effort, respectively) to which they were randomly assigned. The experimenter (always one of the authors) asked the participants about their knowledge of the task and was willing to answer possible questions about the procedure. After going through a short training session, participants started to solve the series of test problems.

Materials A number of different three- and four-disk TOH problems were randomly generated for the experiment. The problems comprised four possible configurations of disks obtained by combining a flat vs. tower disposition in the start state with a flat vs. tower disposition in the goal state.

Two randomly generated three-disk problems, with an optimal solution path of seven moves and with a flat-to-flat configuration, were used for training and presented to the participants in casual order.

The test set comprised eight randomly generated four-disk problems, two for each possible configuration. Each problem had an optimal solution path of 15 moves. The test set was delivered using block randomization.

Apparatus A PowerMacintosh 9500 computer was used for the experiment. A program implementing the

TOH task was written using MCL 4.3 and CLIM 2. The program recorded each participant move (including the moves violating the TOH rules) with the associated time.

The interface was composed by two identical windows, vertically stacked and centered. The upper window showed the initial state of the problem and could be acted upon by the participants. The lower window, which showed the goal state, presented a fixed display. The participants had to perform a drag-and-drop operation with the mouse to move disks from peg to peg in the upper window. In case of an illegal move, an auditory warning was delivered, and the dragged disk was forced back to its source peg.

Experimental Design Two independent variables— one between-subjects (instruction type) and one within-subjects (trial number)— were manipulated in a 2x8 mixed design. The number of trials in the test session (eight) was chosen to obtain an acceptable balance between the possibility of obtaining learning effects and that of inducing fatigue effects. The basic dependent variables were the number of errors (i.e. legal moves in addition to minimum path length), the number of attempted illegal moves, the total time to solve the problem, the mean move latency (excluding the first move), and the time necessary to execute the first move.

Results

All the data analyses were performed on 31 cases¹ (15 in the accuracy, 16 in the effort group) either on transformed and untransformed variables². Given the absence of any difference, we will present only the results obtained using the untransformed variables.

Errors A 2x8 analysis of variance (ANOVA) on the number of errors (Figure 1) showed the significant main effects of instruction type ($F(1,29)=6.57, M SE=173.53, p<.05$), and trial ($F(7,203)=4.95, M SE=69.33, p<.001$). The interaction was not significant. The participants in the accuracy group made fewer errors than those in the effort group ($M=6.73$ for accuracy; $M=11.02$ for effort). In both groups the number of errors decreased from the first block of four trials to the second block ($M=11.38$ for the first block, $M=6.38$ for the second one). A post hoc analysis carried out with the Tukey HSD test

¹Two cases were excluded because the participants needed more than the maximum allowed time (45 min) to complete the first two problems in the test session. One case was excluded because the participant said, only at the end of the session, that she had previously written a program capable of solving this kind of task.

²A logarithmic transformation was performed on all the variables measuring time, while a square-root transformation was applied to all the variables recording the number of moves.

showed significant differences between the following pairs of trials: 1-5 ($p<.05$), 1-6 ($p<.01$), 1-8 ($p<.05$), 2-5 ($p<.01$), 2-6 ($p<.001$) and 2-8 ($p<.01$). The Bonferroni procedure confirmed the results.

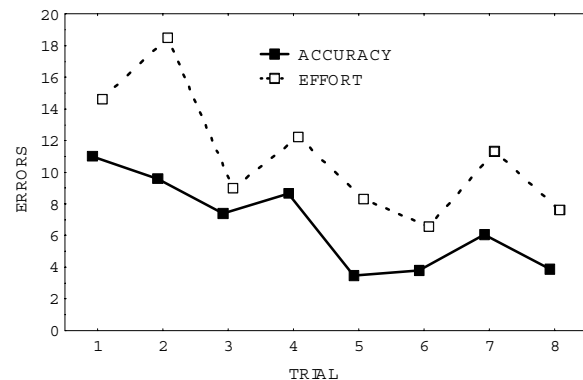


Figure 1: Number of errors for each trial in the accuracy and effort conditions.

Illegal Moves Participants attempted to execute very few illegal moves. The number of such moves was however lower in the accuracy group than in the effort group ($M=0.77$ for accuracy, $M=1.87$ for effort), and decreased from the first to the second block of trials ($M=2.00$ for the first; $M=0.65$ for the second block). Both the effects, but not the interaction, were statistically significant ($F(1,29)=6.71, M SE=11.17, p<.05$ and $F(7,203)=6.37, M SE=3.76, p<.001$, respectively).




Figure 2: Solution times for each trial in the accuracy and effort conditions.

Solution Times A 2x8 ANOVA on the solution times (Figure 2) revealed the significant main effects of the instruction type ($F(1,29)=7.83, M SE=7947.67, p<.01$) and of the trial ($F(7,203)=8.68, M SE=2259.27, p<.001$), while their interaction was not significant. The participants needed more time to solve the problems in the accuracy than in the effort group ($M=100$ s for accuracy, $M=69$ s for effort). The time necessary to complete the task decreased from the first to the second block of trials ($M=103$ s for the first, $M=66$ s for the

second). The Tukey test and the Bonferroni procedure highlighted significant differences between the first trial and the last six and between the second trial and the last four (with the exception of the pair 2-7).

Move Latencies A 2x8 ANOVA on move latency times showed a significant interaction ($F(7,203)=3.10$, $MSE=946783$, $p<.01$) between instructions and trial.

The main effects were also significant: $F(1,29)=14.85$, $MSE=9131212$, $p<.001$ for instructions, $F(7,203)=8.00$, $MSE=946783$, $p<.001$ for the trial. A 2x7 ANOVA with the exclusion of the first trial confirmed the main effects but not the interaction ($F(6,174)=1.29$, $MSE=745583$, $p=.26$). This result suggests that the interaction could be attributed to the extremely high latencies of the participants in the accuracy group on the first trial. This was confirmed by the post hoc tests on the first ANOVA. The move latency was higher in the accuracy group than in the effort group ($M=3.82$ s for accuracy, $M=2.34$ s for effort), and decreased from the first to the second ($M=3.46$ s first block; $M=2.71$ s second block). The Tukey post hoc analysis on the second ANOVA showed significant differences between the pairs 2-5 ($p<.05$), 2-6 ($p<.05$) and 2-8 ($p<.05$). The Bonferroni procedure confirmed only the difference between the trials 2 and 8 ($p<.05$).

First Move Latency A 2x8 ANOVA on the first move latency showed only the significant main effect of the instruction type ($F(1,29)=13.18$, $MSE=583.18$, $p<.01$) with latency higher for participants in the accuracy group ($M=14.78$ s for accuracy, $M=3.64$ s for effort).

Cluster Analysis of Move Latencies We performed also a k-means cluster analysis to determine whether the means of the move latencies and the mean percentages of moves within given latency boundaries were different between the two instruction groups. The cluster analysis was performed on all the moves that required less than 4 s to be executed³. For each subject a solution with 2 clusters (moves having an almost exclusive motor component vs. moves requiring more relevant cognitive processes) was looked for.

A 2x2 (Move x Instruction) ANOVA on the cluster means showed a significant interaction ($F(1,29)=9.46$, $MSE=80919$, $p<.01$) and significant main effects of the move kind ($F(1,29)=2499.81$, $MSE=10229$, $p<.001$) and of the instruction type ($F(1,29)=10.31$, $MSE=10229$, $p<.01$). The interaction is explained by the fact that the

difference of 140 ms between participants in the two instruction groups for the "cognitive" moves ($M=3.08$ s for accuracy and $M=2.94$ s for effort) was significantly smaller than the difference of 305 ms found between the groups for the simplest "execution" moves ($M=1.87$ s for accuracy vs. $M=1.57$ s for effort). These results confirm the indications obtained from the previous move latency analysis, but suggest also a potential execution speed-up for the participants in the effort group.

A further analysis was focused on the mean percentages of cases belonging to the two move clusters and to the moves requiring 4 s or more (the third cluster of "long" moves) in both instruction groups. The results showed significant differences between the accuracy and effort groups for the execution moves ($Mann-Whitney U$ test, $U=57$, $z=2.49$, $p<.05$) and long ones ($U=32$, $z=3.47851$, $p<.001$). In particular, the mean percentage of cases belonging to execution moves was greater in the effort group ($M=61.77$, $SD=8.35$) than in the accuracy group ($M=50.94$, $SD=12.11$). The reverse was true for the long moves (accuracy: $M=24.82$, $SD=10.26$; effort: $M=12.08$, $SD=5.98$). This could mean that participants in the effort group made a higher percentage of execution moves and a lower percentage of cognitive moves in comparison with the moves made by the participants in the accuracy group.

Discussion

There is clear evidence that the experimental manipulation has been very effective in changing the way the TOH problems are solved. As expected, participants are able to achieve their respective goals of minimizing effort and maximizing accuracy, and they are forced by the instructions to trade a lower number of moves with a higher solution time.

There is also clear evidence of the existence of a learning effect. Participants in both groups learn to perform better in successive trials, making fewer errors and using less time. The learning profiles for the two groups remain however distinct across all the trials. The difference concerns not only the errors made and the times needed for solution, but extends to all the dependent variables suggesting that participants in the two groups were selecting and using different solution strategies.

The Simulation

The goal of the simulation was to try to identify the strategies used in each trial by participants in the two instruction groups by comparing several known TOH solution strategies on their capacity to fit the data.

³ Given an independent estimate of 2.15 s for the time needed to move a disk using a TOH program with a direct-manipulation user interface (Anderson, & Lebiere, 1998), we assume that moves requiring 4 s or more are also affected by some kind of higher-order cognitive operation.

The Implemented Strategies

For the simulation we developed a series of ACT-R (Anderson & Lebiere, 1998) models implementing the following solution strategies:

SS1 The selective search strategy described by Anzai & Simon (1979), and subsequently studied by Van Lehn (1991). At each step only disks that are free to move in the current state are considered. The choice of which disk to move and where is guided by two heuristics: "(1) do not move the same disk on consecutive moves, and (2) do not move the smallest disk back to the peg it was on just before it was moved to its current peg" (Van Lehn, 1991, p. 6). Because the strategy is under-specified, an additional assumption has been made: "(3) whenever possible, choose the move which has the effect to put the largest out of place disk (the LOOP disk) into the target peg", which gives the strategy a more goal-oriented attitude. Because the participants did not always follow the directives of the don't-move-twice and don't-undo-move heuristics, the model employs them probabilistically according to two empirically-derived parameters (93% of the cases in which they could be applied when modeling the participants in the accuracy condition, and 90% of the times for the effort condition). Finally, whenever there is still uncertainty about which move to make, the model chooses randomly.

SS2 The selective search strategy previously described augmented with the new one-follows-two heuristics that states that if you have just moved the disk of dimension two, you should now put the smallest disk on top of it.

SP The (simple) perceptual strategy described in Simon (1975) and rephrased as follows: "(1) if all n disks are placed on the target peg, stop; else (2) find the next disk (i) to be placed on the target peg (3) if there are smaller disks on top of disk i , clear them (4) clear disks smaller than i off the target peg (5) move disk i to the target peg (6) go to 1." (Goel & Graffman, 1995, p. 633). In order to avoid being stuck into an infinite loop, because clearing the source peg to move disk i will block the target peg and vice versa, a stack of subgoals is maintained which allows the strategy to be rescued.

KR The strategy described in Karat (1982) which combines elements of domain-specific knowledge into a general problem-solving framework. The strategy adopts a limited look-ahead: if the movement of the LOOP disk from its source to the target peg is blocked by only the smallest two disks, the task of moving the smallest disks on the third peg is considered as trivial, and the moves are immediately executed.

AT In addition to implementing the above mentioned strategies, we utilized also the activation-based model of memory for goals (Aliman & Trafton, 2000)⁴. The

model adopts the strategy of Anderson & Lebiere (1998), but stores goals as ordinary declarative memory elements instead of caching them in the architectural goal stack, and uses a strengthening process for encoding and priming from cues for retrieval.

As previously mentioned, all the strategies are sub-optimal, i.e. they do not generally reach the solution with the minimum number of moves, a performance that also our participants were seldom (i.e., 12% of the times in the accuracy, and 5% in the effort condition) able to make.

Procedure and Results

We executed a simulation of all the strategies on the TOH problem used in the experiment.

We decided to compare the strategies only on their capacity to predict the number of errors made by the participants. Additional assumptions and parameter tuning would be required to model also the times. Therefore, we preferred to stick to a very conservative simulation policy.

The trial-by-trial results of the simulation are presented in Table 1. The table shows the strategies that, in each trial, predicted a number of errors falling into the 99% confidence intervals (CI) computed from the experimental data.

Table 1: Trial-by-trial simulation results.

Group	Trial							
	1	2	3	4	5	6	7	8
Accuracy	KR	KR		KR			KR	KR
	SP	SP	SP	SP	SP	SP	SP	SP
	AT		AT			AT		AT
Effort	S2							
	KR	KR	KR	KR	KR	KR		KR
	SP	SP	SP					SP

The global fit of the three best strategies (SP, KR and AT)—measured using the mean absolute difference (MAD), the root mean square error (RMSE) and the percentage of trials in which the prediction of the model is within the 99% CI (P99CI)—is presented in Table 2.

Table 2: Simulation results for the best strategies.

Strategy	Group	MAD	RMSE	P99CI
KR	accuracy	3.098	3.615	62.5%
SP	accuracy	2.899	3.239	100%
AT	accuracy	5.093	5.763	50%
KR	effort	2.264	3.074	87.5%
SP	effort	5.992	7.111	50%
AT	effort	9.382	10.112	0%

The best fitting strategies are SP in the accuracy condition and KR in the effort condition. The AT

⁴We thank Erik Aliman for making the model available and allowing us to use it in the simulation study

strategy yields good results on half of the trials in the accuracy condition. The selective search strategies are not able to achieve a good fit: only the use of SS2 in the first trial of the effort condition cannot be excluded.

Discussion

The basic conclusion that can be drawn from the simulation is that the results are mainly in compliance with the contingent and adaptive selection hypotheses.

The perceptual strategy is actually more accurate but probably more effortful than the Karat's strategy (that does not require expensive recursive operations). The Altmann & Trafton's model is more accurate than the other two strategies, but probably more expensive than the Karat's model.

Further simulations, using model-tracing and time data, should provide additional supporting evidence.

Conclusions

A preliminary support has been gained for the contingent and adaptive nature of strategy selection in the TOH. On this basis, we suggest that it is important to pay attention to the problem solving factors affecting the accuracy vs. effort trade-off, due to their influence on the strategy selection.

Many other issues must be cleared to obtain a deeper understanding of the selection processes in the TOH and in similar well-structured problems. In this context, we regard as especially important the transition towards more detailed, cognitively grounded strategies to further constrain and specify the existing models, and to allow more detailed comparisons.

This process could yield both the redesign of old strategies and the definition of new ones. Altmann & Trafton (2000) offered a first important contribution with their memory-based model of the Anderson & Lebiere (1998) strategy. We think that a closer analysis and experimental investigation of the attentional and perceptual processes in the TOH could produce significant advances in our understanding of the cognitive processes underlying the solution strategies.

References

- Ahlum Heath, M. E. & DiVesta, F. J. (1986). The effect of conscious controlled verbalization of a cognitive strategy on transfer in problem solving. *Memory & Cognition*, 14, 281-285.
- Altmann, E. M. & Trafton, J. G. (2000). An activation-based model of memory for goals. Manuscript submitted for publication.
- Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Erlbaum.
- Anderson, J. R., Kushmerick, N. & Lebiere, C. (1993). The tower of Hanoi and goal structures. In J. R. Anderson (Ed.), *Rules of the mind*. Hillsdale, NJ: Erlbaum.
- Anderson, J. R. & Lebiere, C. (1998). *The atomic components of thought*. Hillsdale, NJ: Erlbaum.
- Anzai, Y. & Simon, H. A. (1979). The theory of learning by doing. *Psychological Review*, 86, 124-140.
- Christensen-Szalanski, J. J. J. (1998). Problem-solving strategies: a selection mechanism, some implications, and some data. In L. R. Beach (Ed.), *Image Theory*. Hillsdale, NJ: Erlbaum.
- Davies, S. P. (2000). Memory and planning processes in solutions to well-structured problems. *The Quarterly Journal of Experimental Psychology*, 53A, 896-927.
- Fum, D. & DelMisser, F. (2000). Adaptive spatial planning. *Proceedings of the Seventh Annual ACT-R Workshop*. Pittsburgh: Carnegie Mellon University.
- Goel, V. & Grafman, J. (1995). Are the frontal lobes implicated in "planning" functions? Interpreting data from the tower of Hanoi. *Neuropsychologia*, 33, 623-642.
- Karat, J. (1982). A model of problem solving with incomplete constraint knowledge. *Cognitive Psychology*, 14, 538-559.
- Miyake, A., Friedman, N. P., Emerson, M. J., Witzki, A. H., Howerter, A. & Wager, T. D. (2000). The unity and diversity of executive functions and their contribution to complex "frontal lobe" tasks: a latent variable analysis. *Cognitive Psychology*, 41, 49-100.
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. New York: Cambridge University Press.
- Ruiz, D. & Newell, A. (1989). Tower-noticing triggers strategy-change in the Tower of Hanoi: a Soar model. *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society* (pp. 522-529). Hillsdale, NJ: Erlbaum.
- Simon, H. A. (1975). The functional equivalence of problem solving skills. *Cognitive Psychology*, 7, 268-288.
- Stinesen, L. (1985). The influence of verbalization on problem solving. *Scandinavian Journal of Psychology*, 26, 342-347.
- VanLehn, K. (1991). Rule acquisition events in the discovery of problem solving strategies. *Cognitive Science*, 15, 1-47.