Title
Planning with Information Access Costs in Mind

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 28(28)

ISSN
1069-7977

Authors
Waldron, Samuel M.
Patrick, John
Howes, Andrew
et al.

Publication Date
2006

Planning with Information Access Costs in Mind

Samuel M. Waldron (waldronsm1@cardiff.ac.uk)
John Patrick (patrickj@cardiff.ac.uk)
Department of Psychology, Park Place
Cardiff, CF10 3AT UK

Andrew Howes (howesa@manchester.ac.uk)
Geoffrey B. Duggan (geoffrey.duggan@manchester.ac.uk)
Department of Informatics, PO Box 88, Sackville Street
Manchester, M60 1QD UK

Abstract

Least effort tradeoffs concerning the use of memory within the Blocks World Task (Fu & Gray, 2000) are extended to an analogous problem solving task. As the cost of accessing goal-state information from the Blocks Problem Solving Task increased, participants chose to access such information less frequently, and in turn, made more problem solving moves per goal-state visit. This strategy shift led to an increase in the number of moves required to solve each problem, suggesting that effective planning became difficult as the access cost increased. In contrast, increasing the implementation cost, a manipulation known to increase planful behavior (O'Hara & Payne, 1998), revealed quite different problem solving access strategies, and reduced the number of moves required to solve each problem.

Keywords: Interactive behavior; Human problem solving; Information access; Adaptive planning.

Introduction

Problem solving is often conducted within the context of an external display (Larkin, 1989; Payne, 1991). However, the time, physical effort, and mental effort associated with accessing information from such displays may vary. For example, information may be available at the click of a button, or considerable time and effort may be required in order to extract the necessary information from the interface. Such information access costs have been shown to determine what perceptual-motor and memory strategies are used during interactive behavior (Fu & Gray, 2000; Gray & Fu, 2004), but have yet to be extended to problem solving. The current paper aims to test for the presence of shifts in strategy due to information access cost during problem solving, and assess consequences for planning proficiency.

Planning

Research assessing the manner in which humans go about planning has informed us that planning and acting are largely interleaved (Anderson, 1990). That is, we do not plan complete sequences before we perform them. Rather, we evaluate our plans as we progress through the problem space. Typical planning behavior is said to be opportunistic (Hayes-Roth & Hayes-Roth, 1979), in the sense that plans can be revised in the face of new information, or as a consequence of an evaluation of a previous plan. The proposed benefits of planning have also been found to interact with the limited capacity of working memory (Kotovsky, Hayes, & Simon, 1985; Phillips, Wynn, Gilhooly, Della Sala, & Logie, 1999).

Moreover, problem solving within the context of an external display has been shown to be sensitive to costs within the environment. O'Hara & Payne (1998) stated that "problem solving search strategies are chosen so as to optimize performance within the constraints of a particular situation" (p.34). Gray & Fu (2004) recast this assertion by distinguishing between hard and soft constraints. Hard constraints determine behavior that is, or is not possible, whereas soft constraints bias towards certain behavioral strategies.

An example of a soft constraint that promotes planfulness within a problem solving context is the implementation cost. Using the 8-Puzzle, O'Hara & Payne (1998) demonstrated that increasing the cost of making each move from pressing a function key to typing a string increased participant's propensity to plan ahead, and in turn, reduced the number of moves needed to complete the task. Evidence of planfulness came mainly from increased latencies between each move. Verbal protocol analysis also suggested more planful behavior as a function of increased implementation cost. Thus, it appears that the extent to which problem solvers choose to plan can be partially determined by costs within the environment associated with completing the task.

Information Access Cost

The planning research briefly reviewed above assumes that the information needed to solve a problem is either readily available in the world, or is held in memory. For example, the 8-Puzzle goal-state used by O'Hara & Payne (1998) is so easy to maintain in memory that participants were not required to search for additional information in order to solve the problem. For the vast majority of other problem solving tasks (e.g., Tower of Hanoi, Missionaries & Cannibals, Water Jugs), all task-relevant information is provided within the external display with little or no associated access cost.

The converse, however, is often true of problem solving in many applied situations. That is, the information required to solve a problem may not always be readily available within the external display, and may come from a variety of sources. Gathering unavailable or uncertain information routinely forms the first stage of the planning process in many applied situations (Gronlund, Dougherty, Durso, Canning, & Mills, 2005).

While there appears to be little work examining the consequences of information access cost on problem solving, Fu & Gray (2000) have used the Blocks World Task (see Figure 1) to explore the consequences of access costs on routine copying behavior. The aim of the Blocks World Task is to recreate the pattern of blocks dictated by the Target Window in the Workspace Window. This is achieved by clicking on and dragging blocks one at a time from the Resources Window into the Workspace Window. Throughout the Fu & Gray (2000) study, all three windows were covered by grey boxes, of which only one could be uncovered at a time. The Resources and Workspace windows could be uncovered by moving the mouse cursor into each window. Each window was covered again the moment the mouse cursor left the window.

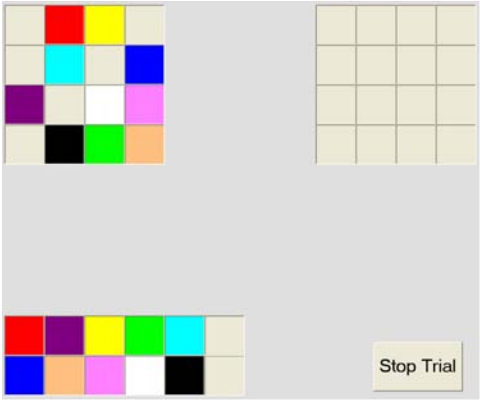

Figure 1: Blocks World Task. Target Window is top-left, Resources Window is bottom-left, and Workspace Window is top-right.

The cost of uncovering the Target Window varied between the three conditions. In the Low-Cost condition participants were required to hold down a function key. The Target Window then remained uncovered until either the function key was released or the mouse cursor entered either of the other two windows. In the Control condition, the Target Window was uncovered during times when the mouse cursor entered the window. This was also the case in the High-Cost condition, with an additional delay of one second per uncovering. Once participants were confident that the Target Pattern had been recreated in the Workspace, they were instructed to click the 'Stop Trial' button. If the two patterns matched, participants were taken to the next trial. If they did not, participants were required to correct their mistakes.

Fu & Gray (2000) found that when compared to the Low-Cost and Control conditions, participants in the High-Cost condition uncovered the Target Window less frequently, spent more time looking at the Target Window, and copied more blocks per visit to the Target Window. These differences indicate that as the information access cost increased, behavior shifted away from reliance upon perceptual-motor strategies, towards reliance upon internal memory. This change in strategy was despite an increase in the number of errors made by participants in the High-Cost condition, and can be explained in terms of Anderson's (1990) rational analysis framework.

The aim of the current study was to assess the presence of such shifts in strategy contingent upon access costs within an analogous problem solving task. In particular, consequences of access cost will be evaluated with regard to participant's ability to plan effectively.

Experiment

The Blocks Problem Solving Task (BPST) was developed, and required participants to recreate the Target Pattern in the Workspace Window by moving one block at a time into an adjacent empty space (see Figure 2). The cost associated with accessing each window and making each move was manipulated.

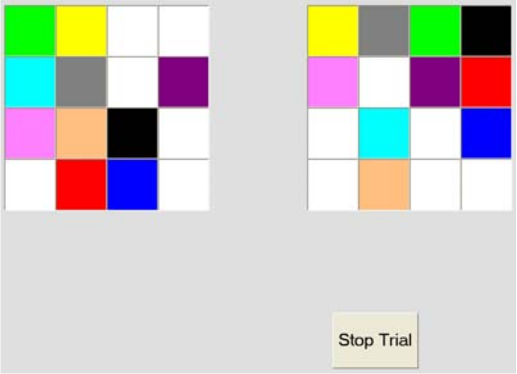

Figure 2: Blocks Problem Solving Task. Target Window is on the left, Workspace Window is on the right.

Based upon Fu & Gray (2000), we predict that as the information access cost (IAC) increases, perceptual-motor strategies will be overlooked in favor of a strategy reliant upon internal memory. Such evidence would indicate that the IAC strategy shift observed during routine copying behavior also extends to problem solving.

The cost associated with making each move, namely, the implementation cost (IC) was also manipulated. We know that such a manipulation increases planful behavior in the 8-Puzzle (O'Hara & Payne, 1998), and should therefore provide a useful comparison to planning behavior in the current task.

Participants may adopt a number of strategies to cope with higher access costs. On the one hand, a higher access cost may prompt them to do more planning in the hope that fewer moves are required and, by implication, fewer visits to the Target Window (O'Hara & Payne, 1998). On the other hand, they may try to remember bigger chunks of the target pattern (Fu & Gray, 2000), and not worry about the number of moves made in the problem space. A between-subjects design will be adopted in order to negate possible contamination via asymmetric transfer (Poulton, 1982).

Method

Participants Seventy-two Cardiff University undergraduate Psychology students participated in the study for course credit and were randomly assigned to one of six conditions.

Apparatus/Materials The experiment was written in Visual Basic 6 and was conducted using a 2Ghz Pentium 4 PC connected to a Tobii 1750 34 x 27cm eye-tracker monitor[1], extended keyboard, and mouse. All eye movements were recorded at a rate of 15 frames per second, with time-stamp accuracy of +/-3ms. Gaze estimation was within 1 degree of accuracy, even across large head movements. Mouse movements and key presses were also recorded and saved.

The Target Window and Workspace Window were the same size and each contained sixteen blocks in a 4 x 4 grid. Ten colored blocks and six empty spaces resided within each window. No colors were used twice, and the empty spaces were white. The rules of the task determined that colored blocks could only be moved into adjacent (horizontal or vertical) empty spaces. Movements were made in two stages. Firstly, using the mouse participants clicked to select the colored block they wished to move. The second stage was dependent upon the level of IC.

Design IAC was manipulated on three levels: both windows were permanently uncovered when IAC was Low. In contrast, both windows were covered by grey masks when IAC was Medium or High, and could only be uncovered by placing the mouse cursor over the window to-be-opened. The grey masks then reappeared the moment the mouse cursor left the respective window. There was an additional 2.5 second lockout associated with uncovering the Target Window when IAC was High.

IC was manipulated on two levels: when IC was Low, participants were required to press the corresponding arrow key depending upon the direction they wished to move the block. When IC was High, participants were required to type a string "move_left/right/up/down_". On average, this resulted in each move taking an extra 2.5 seconds to implement.

Independent manipulation of IAC and IC resulted in six between-subject conditions: low IAC/low IC; low IAC/high IC; medium IAC/low IC; medium IAC/high IC; high IAC/low IC; high IAC/high IC.

Several dependent measures were taken throughout the experiment. The eye-tracker measured the frequency with which participant's eyes visited the Target Window, and the time spent viewing the Target Window. Consecutive fixations within the Target Window were collapsed and counted as one visit. In addition, the Visual Basic program recorded the number of moves made, the time taken to complete each trial, the time between each move (inter-move latencies), and the frequency with which participants clicked the stop button when in fact the two patterns did not match (errors). By dividing the number of moves by the number of Target Visits for each trial, an estimate of the number of moves made per Target Visit was obtained.

Procedure Participants were seated approximately 50cm away from the eye-tracker and handed an instruction sheet. Two practice trials then followed a 16-point eye-tracker calibration. Both practice trials were in the format of the experimental condition. Different block configurations were used for each of the twelve experimental trials. Each participant within each of the six conditions received one of twelve different randomized orders of trials.

Results

A 3 (Low/Medium/High IAC) x 2 (Low/High IC) between-subjects ANOVA conducted upon each of the dependent variables revealed quite different results for IAC and IC (summarized in Tables 1 & 2 respectively). Only one interaction between the two independent variables existed, and will be addressed following separate discussions of the IAC and IC. Target Visit Frequency and moves per Target Visit data was log transformed in order to attain homogeneity of variance. Non-transformed data are presented in tabular and graphical format throughout.

IAC Target Visit Frequency (TVF) decreased as IAC increased, $F(2, 66) = 395.39$, $p < .001$, $MSE = 0.01$, with planned comparisons (Bonferroni corrected) revealing significant differences between all conditions ($ps < .001$). The non-significant trend was for Target Visit Times (TVT) to increase as IAC increased, $F(2, 66) = 2.84$, $p < .07$, $MSE = 27.97$, and planned comparisons (Bonferroni corrected) revealed significant differences between low and high ($p < .05$), medium and high ($p < .05$), but not medium and low conditions ($p > .05$).

Table 1: Effect of IAC

	Low IAC		*Medium IAC*		*High IAC*	
	Mean	SD	Mean	SD	Mean	SD
TVF	49.35	13.10	21.21	12.26	5.25	1.54
TVT	17.11	4.94	17.17	8.01	20.29	7.51
M/TV	0.89	0.35	2.53	1.60	8.70	2.56
IML*	68.71	22.99	72.73	27.20	79.76	26.62
Errors	0.06	0.11	0.13	0.13	0.19	0.15
Moves	38.04	6.06	40.16	5.82	42.36	5.94

Note. Values represent means per trial and exclude delays incurred via IAC.
* Excluding error-correction data.

Increasing the IAC not only affected information seeking strategies, but also had consequences for problem solving search strategies. Firstly, the number of moves participants made per visit to the Target Window (M/TV) increased

---

[1] The Tobii 1750 eye-tracker does not require head mounts.

dramatically as a function of IAC, $F(2, 66) = 349.36$, $p < .001$, $MSE = 0.01$. When IAC was low, participants generally chose to interleave a minimum of one visit to the Target Window per move made in the Workspace Window. When the IAC was high, however, participants were prepared to make up to nine moves in the Workspace Window per visit to the Target Window. Although at first glance inter-move latencies (IML) appeared to increase as IAC increased, $F(2, 66) = 3.71$, $p < .05$, $MSE = 342.14$, removal of error-correction data meant that the marginal rise became non-significant, $F(2, 66) = 2.34$, $p > .05$, $MSE = 320.43$. By error correction data, we mean trial data during which the participant pressed the stop button when in fact the two patterns did not match (and thus had to spent time identifying and correcting the error). Removal of error data did not affect any other variable, and on average, accounted for less than 13% of the data. Finally, the number of errors made increased as IAC increased, $F(2, 66) = 6.03$, $p < .01$, $MSE = 0.02$, as did trial duration, $F(2, 66) = 6.53$, $p < .01$, $MSE = 781.45$, (Low: *Mean* = 125.91, *SD* = 66.79; Medium: *Mean* = 135.79, *SD* = 76.06; High: *Mean* = 154.62, *SD* = 69.63), and the number of moves required to solve each trial, $F(2, 66) = 7.14$, $p < .01$, $MSE = 15.64$.

In sum, increasing the cost associated with accessing goal-state information from the BPST reduced the frequency with which said information was accessed, marginally increased the time spent viewing such information, and increased the number of moves made per access. An increase in IAC did not lead to a reliable rise in inter-move latencies, but did increase the frequency with which errors were made, trial duration, and ultimately, the number of moves required to solve each problem.

IC Increasing the IC induced comparatively different information seeking and problem solving strategy development. Target Visit Frequency (TVF) increased as IC increased, $F(1, 66) = 40.18$, $p < .001$, $MSE = 0.01$, as did time spent viewing the Target Window (TVT), $F(1, 66) = 48.12$, $p < .001$, $MSE = 27.97$.

Table 2: Effect of IC.

	Low IC		*High IC*	
	Mean	SD	Mean	SD
TVF	19.07	15.14	28.86	24.27
TVT	13.86	5.10	22.52	5.94
M/TV	4.84	3.90	3.22	3.60
IML*	55.50	12.87	91.96	22.20
Errors	0.16	0.13	0.09	0.14
Moves	44.58	5.33	35.79	2.78

Note. Values represent means per trial and exclude delays incurred via IAC.
* Excluding error-correction data.

Rather than increasing the moves per Target Visit (M/TV) ratio, the high IC reduced the number of moves participants were prepared to make per visit to the Target Window, $F(1, 66) = 59.24$, $p < .001$, $MSE = 0.01$. In addition, a dramatic rise in inter-move latency (IML) was observed as a function of IC, $F(1, 66) = 38.43$, $p < .001$, $MSE = 354.46$, that was not affected by the removal of error-correction data. Increasing the IC also reduced the frequency with which errors were made, $F(1, 66) = 6.85$, $p < .01$, $MSE = 0.02$, and the number of moves required to solve each problem, $F(1, 66) = 89.00$, $p < .001$, $MSE = 15.64$. As with IAC, trial times increased as a function of IC, $F(1, 66) = 377.61$, $p < .001$, $MSE = 781.45$, (Low: *Mean* = 60.94, *SD* = 12.66; High: *Mean* = 97.14, *SD* = 24.80).

In sum, increasing the cost of making each move from pressing an arrow key to typing a string increased the frequency with which participants chose to access goal-state information, increased the time spent viewing said information, and reduced the number of moves made per Target Visit. An increase in IC also led to reliable increases in the latency between each move, a reduction in the number of errors made, a reduction in the number of moves required to complete each problem, and an increase in trial durations.

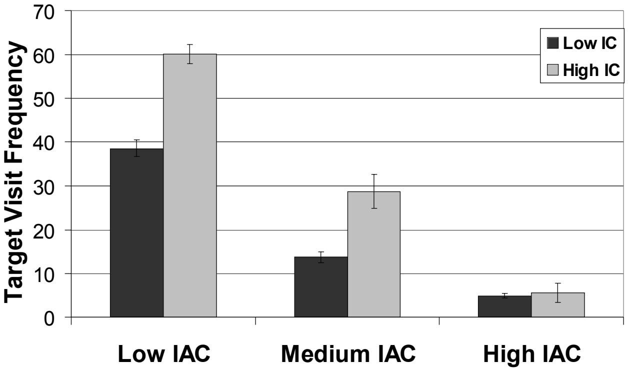

Figure 3: Interaction between IAC and IC on number of Target Visits per trial.

The only interaction between IAC and IC concerned the frequency with which participants visited the Target Window, $F(2, 66) = 8.03$, $p < .001$, $MSE = 0.01$, (see Figure 3). Simple main effects indicated that the effect of IAC was significant at both Low, $F(2, 66) = 72.51$, $p < .001$, $MSE = 50.44$, and High, $F(2, 66) = 178.22$, $p < .001$, $MSE = 50.44$, levels of IC. The effect of IC, although significant at Low, $F(1, 66) = 55.12$, $p < .001$, $MSE = 50.44$, and Medium, $F(1, 66) = 26.74$, $p < .001$, $MSE = 50.44$, levels of IAC, was not apparent at High IAC, $F(1, 66) = 0.51$, $p > .05$, $MSE = 50.44$. This we believe, is likely to be due to a ceiling effect.

Pearson correlational analysis was also conducted in order to assess the inter-relationships between some of the dependent variables. Correlations were computed for measures within each of the six between-subject treatment combinations and also across the complete data-set. The only significant correlation consistently found within each of the six treatment combinations and across the data-set indicated that as Target Visit Frequency increased, the number of moves made per Target Visit decreased (-.78).

Despite similar correlations, this relationship varied with IAC (as illustrated by Figure 4).

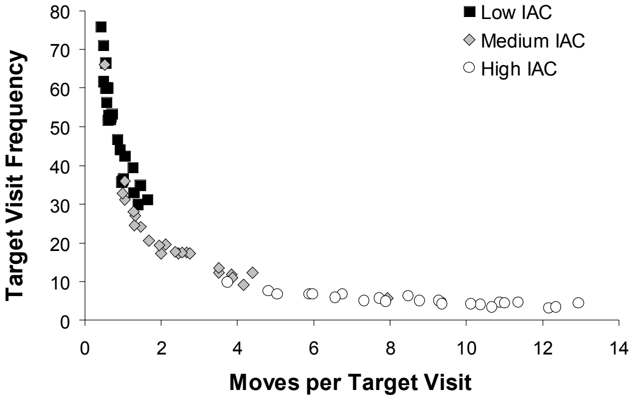

Figure 4: Scatter plot of relationship between Target Visit Frequency and Moves per Target Visit as a function of IAC.

Other significant correlations across the data-set included those found between Target Visit Frequency and moves (-.49), moves per Target Visit and moves (+.47), and inter-move latency and moves (-.49). These relationships between the aforementioned variables suggest that in order to reduce the number of moves required to solve the current set of problems, participants should visit the Target Window frequently, make fewer moves in the Workspace Window per visit, and increase planning time between each move.

Discussion

Firstly, the results indicate that the shifts in strategy contingent upon the IAC observed by Fu & Gray (2000) during routine copying behavior extend to an analogous problem solving task. Secondly, a tentative argument can be put forward that such shifts in information access strategy make efficient problem solving and planning behavior difficult in the current task.

The rational analysis framework (Anderson, 1990) has previously been adapted to explain the tradeoff between perceptual-motor and memory strategies in studies using the Blocks World Task (Fu & Gray, 2000) and simulated VCR programming (Gray & Fu, 2004). Here, we take a similar perspective and argue that as the IAC increases, the cost of accessing goal-state information following each move is considered higher than the cost of using memory (thus, reducing the number of times goal-state information needs to be accessed).

In a related study, Pfeiffer (2004) found similar shifts in problem solving search strategy upon the removal of *current-state* information. When such information was not continually available during the seven balls and boxes problem, participants tended to rely upon memory and make up to nine moves per current-state request. Interestingly, participants were no worse at completing the balls and boxes puzzle when current-state information was only available upon request, compared to when current-state information was continuously available within the external display. Differences in the volume of information to be held in memory may explain why Pfeiffer's study did not reveal any negative consequence of reducing the availability of current-state information, whereas the present study did demonstrate negative consequences associated with increasing the cost of accessing goal-state information.

In order to maintain an internal representation of the entire problem space in Pfeiffer's study, participants would have been required to maintain in memory the positioning of up to seven balls in seven boxes. In the present study, however, participants will have been required to maintain in memory the positioning *and* color of ten blocks within a 4 x 4 grid. It is likely, therefore, that in order to work from memory, participants in the current study will have had to decompose the goal-state into chunks of three/four pieces at a time (Chase & Simon, 1973). Essentially, this would involve problem solving with limited goal-state knowledge in mind.

This argument draws parallels with the finding from Fu & Gray (2006) that reductions in information seeking behavior as a function of increased IAC may lead to poor exploration of the problem space. Fu & Gray (2006) demonstrated that individuals tend to plan less when the cost of planning is high, and stop planning when a reasonable plan is formed. It is suggested that a similar argument may explain the increased number of moves required to solve the BPST in the current study as a function of IAC. Further experimentation is required to evaluate the extent to which this can account for the data. For now, contrasting the effects brought about by manipulating the IAC and IC may provide a better understanding of why raising the IAC increased the number of moves required to solve each problem.

Corroborating the results reported by O'Hara & Payne (1998), the present study found that increasing the IC increased planful behavior, as evidenced by a reduction in the number of moves required to solve each problem, and an increase in inter-move latencies. Interestingly, the information access strategies induced by raising the IC stand in opposition to those associated with an increase in IAC. Most importantly, increasing the IC increased the frequency with which participants chose to access goal-state information throughout each trial, and in doing so, decreased the number of moves made per access. The opposite is true of an increase in IAC.

If we assume that behavior in the current study is representative of efficient problem solving when IC is high, it would follow that regular inspection of the goal-state between moves is also representative of efficient problem solving. Such a proposal would be consistent with opportunistic theories of planning (e.g., Hayes-Roth & Hayes-Roth, 1979), and evidence suggesting that planning and acting are largely interleaved (Anderson, 1990). Such an argument necessitates the assumption that participants in the high IC condition are choosing to shift their information access strategies in said manner in a direct attempt to increase problem solving efficiency. On the other hand, it could be possible that participants in the high IC conditions chose to increase the frequency with which they accessed goal-state information because of an increased likelihood that information would be forgotten while typing move sequences, (compared to pressing arrow keys). If this was true, however, one would expect the number of Target Visits to differ to a greater extent as a function of IC when IAC was high, compared to when it was low. The opposite, in fact, was observed: increased IC led to a significant increase in the number of goal-state inspections when IAC was low, but had no significant effect when IAC was high. Although likely to be the result of a ceiling effect, the pattern of data does not support the 'forgetting hypothesis'.

Strategies that increase the frequency with which goal-state information is accessed in the current task will not only facilitate the interleaving of planning and acting, but will also decrease reliance upon memory, known to improve problem solving efficiency and planning behavior (Kotovsky, Hayes, & Simon, 1985; Phillips et al., 1999). Kotovsky, Hayes, & Simon (1985), for example, found that increasing the load on working memory whilst solving the Tower of Hanoi prevented even minimal planning from occurring. Problem solving efficiency improved as rules were built into the external representation of an isomorphic version of the task. Similarly, Phillips et al., (1999) found reduced levels of planning as secondary tasks, designed to tax working memory, were introduced whilst participants attempted to solve versions of the Tower of London task. Phillips et al., went on to highlight the importance of memory in formulating, retaining, implementing, and revising plans online.

Davies (2003) distinguished between initial and concurrent planning during solutions to well-structured problems (such as the BPST). The increased inter-move latencies and decreased number of moves made per Target Visit associated with an increased IC in the current task are likely to be reflective of *concurrent* planning. The fact that increasing the IAC did not affect inter-move latencies (once error correction data was removed from the analysis), and more notably, increased the number of moves made per Target Visit, suggests that if participants are making use of planning strategies in the high IAC conditions, these will be in the form of *initial* planning, and not concurrent planning. Further verbal protocol studies are required to determine the extent to which each type of planning is represented within each of the conditions.

Implications for display design are twofold. Firstly, if accurate problem solving is of utmost importance, and the speed of problem solving is of less importance, seeding implementation costs into the design of an interactive display may improve problem solving accuracy via planful behavior. Secondly, when solving problems of similar complexity and volume to the BPST employed here, information required to solve the problem should be readily available at minimum cost.

Conclusion

Insight is provided into higher-level cognitive planning using theory originally built upon perceptual processing tasks. In sum, it appears that an increase in the cost associated with accessing goal-state information promotes the use of memory and initial planning during problem solving. In contrast, information access strategies induced by an increase in implementation cost rely less upon memory during problem solving, and more upon concurrent planning.

Acknowledgments

This work was funded by the UK MoD's Defence Technology Centre: Data and Information Fusion (project 4.9).

References

Anderson, J. R. (1990). *The adaptive character of thought.* Hillsdale, NJ: Erlbaum.

Chase, W. G. & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology, 4*, 55-81.

Davies, S. P. (2003). Initial and concurrent planning in solutions to well-structured problems. *The Quarterly Journal of Experimental Psychology, 56A(7)*, 1147-1164.

Fu, W.-T. & Gray, W. D. (2000). Memory versus perceptual-motor tradeoffs in a blocks world task. *Proceedings of the Twenty-second Annual conference of the Cognitive Science Society (pp. 154-159).* Hillsdale, NJ: Erlbaum.

Fu, W.-T. & Gray, W. D. (2006). Suboptimal tradeoffs in information seeking. *Cognitive Psychology, 52(3)*, 195-242.

Gray, W. D. & Fu, W.-T. (2004). Soft constraints in interactive behavior: The case of ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head. *Cognitive Science, 28*, 359-382.

Gronlund, S. D., Dougherty, M. R. P., Durso, F., Canning, J. M., & Mills, S. H. (2005). Planning in Air Traffic Control: Impact of problem type. *The International Journal of Aviation Psychology, 15(3)*, 269-293.

Hayes-Roth, B. & Hayes-Roth, F. (1979). A cognitive model of planning. *Cognitive Science, 3*, 275-310.

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. H. (1989). Display-based problem solving. In D. Klahr & K. Kotovsky (Eds.). *Complex information processing: The impact of Herbert A. Simon*. Hillsdale, NJ: Erlbaum.

O'Hara, K. P. & Payne, S. J. (1998). The effects of operator implementation cost on planfulness of problem solving and learning. Cognitive *Psychology, 35*, 34-70.

Payne, S. J. (1991). Display-based action at the user interface. *International Journal of Man-Machine Studies, 35*, 275-289.

Pfeiffer, T. (2004). Problem solving with a simple transformation problem with and without continuous external support. *European Journal of Cognitive Psychology, 16(4)*, 555-572.

Phillips, L. H., Wynn, V., Gilhooly, K. J., Della Sala, S., & Logie, R.H. (1999). The role of memory in the tower of london task. *Memory, 7(2)*, 209-231.

Poulton, E. C. (1982). Influential companions: Effects of one strategy on another in the within-subjects designs of cognitive psychology. *Psychological Bulletin, 91(3)*, 673-690.