Why Diagrams Are (Sometimes) Six Times Easier than Words: Benefits beyond Locational Indexing

Peter C-H. Cheng

Department of Informatics, University of Sussex
Falmer, Brighton, BN1 9QH, UK
p.c.h.cheng@sussex.ac.uk

Abstract. By building computational models, Larkin and Simon (1987) showed that the effects of locational indexing give an explanation of 'Why a diagam is (sometimes) worth ten thousand words', to quote the title of their seminal paper. This paper reports an experiment in which participants solved three versions of Larkin and Simon's simple pulley system problem with varying complexity. Participants used a diagrammatic, tabular or sentential representation, which had different degrees of spatial indexing of information. Solutions with the diagrams were up to six times easier than informationally equivalent sentential representations. Contrary to predictions derived from the idea of locational indexing, the tabular representation was not better overall than sentential representation and the proportional advantage of the diagrammatic representation over the others did not increase with problem complexity. This suggests that the advantage of diagrams goes beyond the effects that locational indexing has on the processes of searching for items of information and the recognition of applicable rules. A possible explanation resides in the specific problem solving strategies that the participants may have been using, which depended on the structure of the representations and the extent to which they supported solution path recognition and planning.

A. Blackwell et al. (Eds.): Diagrams 2004, LNAI 2980, pp. 242-260, 2004.

1 Introduction

The nature of the representations used for any non-trivial task will greatly impact upon, or even substantially determine, the overall difficulty of doing task. This is now a well-established and accepted finding in Cognitive Science and allied areas, and building upon this, our understanding of the nature of representations is growing. In superficially different problems with the same underlying structure, the form representation used can make finding a solution take up to 16 times as long [6]. External representations provide benefits such as being able to off-load computation on to the representation [7], [9]. Using orthogonal visual dimensions for distinct dimensions of information in a graphical display can make problems solving easier by supporting the separation of those information dimensions in the mind [10]. Scientific discoveries are often made possible through the invention of new representations [1], [3]. Conceptual learning can be enhanced by alternative representations that more clearly reveal the underlying structure of scientific and mathematic topics, by deliberately using representational schemes that encode the underlying structure in a transparent fashion [2].

1.1 Why a Diagram Is (Sometimes) Worth 10,000 Words'

Larkin and Simon's 1987 seminal paper [8], which has the above title, is arguably the seed around which much of this research area in cognitive science has crystallized. That paper provided critical insights into the potential benefits of diagrammatic representations over propositional or sentential representations. Larkin and Simon built computational (production system) models for two domains, simple pulley problems and geometry, which explained why diagrams are often computationally better than informationally equivalent sentential representations. In diagrams items of information that are likely to be processed at the same time in a problem solution are often found together spatially. This locational indexing yields benefits both in the search for information and the recognition of what rules to apply. Further, this reduces the need to laboriously match symbolic labels.

The starting point for the Larkin and Simon paper is the assumption that diagrams are often better than sentential representations and the relative advantage of diagrams needs to be explained. Given the importance of this paper, the consequences of that assumption being wrong would be substantial. It would undermine Larkin and Simon's claims and could draw into question the subsequent work built upon it. However, there has apparently been no direct empirical evaluation to demonstrate that the diagrams in the paper are, in reality, found by people to better for problem solving than sentential representations. This paper presents an experiment using simple pulley system problems, like those used in the Larkin and Simon paper, to investigate the relative difficulty of solution within different representations.

Reassuringly, the experiment demonstrated that humans find using diagrams of the simple pulley systems problem substantially easier to solve than equivalent sentential representations. Of course, this finding is not really a surprise. Rather, the main aims of the experiment were actually to: (1) quantify the benefits of the diagram for the pulley problem over problems of varying complexity; and, (2) explore the role of locational indexing in more detail, by comparing the diagrammatic and sentential representations with a tabular representation that is mid way between the others in the extent to which it uses spatial location as a means to index information.

1.2 Isomorphic Problems with or without Equivalent Rules

Larkin and Simon define a representation as consisting of a format and a set of operators and heuristics [8]. The format is the graphical (or notational) structure of the display (or expressions). The operators are applied to change the components of the display or to change mental states encoding the display. Heuristics guide the selection of suitable operators to use. All three, format, operators and heuristics, are essential for a representation since together they determine the nature of the problem space experienced by the problem solver. For instance, given a set of numbered concentric circles, the representation and problem is very different if the user brings 'Venn diagram' set theoretic operators or arithmetic operators for ‘target shooting' to the interpretation of the display.

It has been shown that isomorphic problems can be over an order of magnitude more difficult under alternative representations. For example, Kotovsky, Hayes and Simon [6] used variants of the Tower of Hanoi problem, such as the *monster-move* and *monster-change* problems, which had different operators or rules. In the monster move problem globes were passed back and forth between the monsters in a manner similar to the transfer of disks between pegs in the original Tower of Hanoi problem. In the monster change problem the change in size of globes held by each monster was equivalent to the transfer of disks. Although the underlying logical structure of the problems was always the same, with equivalent size and structure of problem state spaces, the different rules, productions, required for the alternative versions of the problem imposed substantially different loads on working memory. This impacted substantially on the overall difficulty of the tasks. In the move version of problem the image of the location of the globes/disks provided strong visual cues about what moves are legal to make, whereas in the change version of the problem logical mental inferences were required to work out permitted legal states.

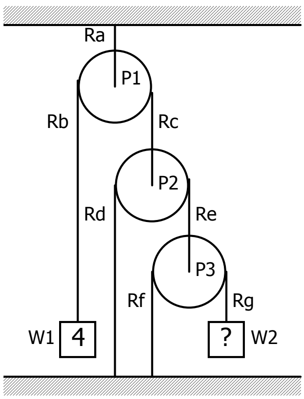

Fig. 1. Diagrammatic representation for the medium complexity pulley system problem used in the experiment. The alphanumeric symbols are not present in the experimental stimulus (except ‘4' and ‘?') but are present to enable readers to cross reference components to the tabular and sentential representations Figs. 2 and 3

In the Tower of Hanoi studies the problems were isomorphic but the representations were different as the rules were experimentally manipulated. This raises the question of how much the difficulty of problems may differ when format of the display varies but the rules remain the same. The present experiment examines alternative representations of simple pulley problems to quantify the relative impact of different representations with equivalent rules but under alternative formats. The formats are diagrams, tables and sentential representations. Problems of three complexity levels were designed. Figs. 1, 2 and 3 show examples of a diagram, a tabular and a sentential representation as used in the experiment. In the problems, the participants are given the value of one of the weights and they compute the value of another unknown weight, assuming that the system is in static equilibrium.

	W1	Rb	P1	Rc	Ra	C	Rd	P2	Re	F	Rf	P3	Rg	W2
Weight 1	x													
Rope b		x												
Pulley 1			x											
Rope c				x										
Hangs	[ ]	[ ]												
Pulley system		[ ]	[ ]	[ ]										
Rope a					x									
Ceiling						x								
Hangs			[ ]		[ ]									
Hangs					[ ]	[ ]								
Rope d							x							
Pulley 3								x						
Rope e									x					
Hangs				[ ]				[ ]						
Pulley system							[ ]	[ ]	[ ]					
Floor										x				
Anchors							[ ]			[ ]				
Rope f											x			
Pulley 3												x		
Rope c													x	
Hangs									[ ]			[ ]		
Pulley system											[ ]	[ ]	[ ]	
Floor										x				
Anchors										[ ]	[ ]			
Weight 2														x
Hangs													[ ]	[ ]
Value	4													?

Fig. 2. Table representation (medium complexity) as used in the experiment. Columns co-reference components in the rows, as indicated by 'x' symbols. Cells with '[ ]' indicate components are part of the same assembly and the value is to be completed by problem solver. Final value to be found indicated by '?'. Arrows are explained in the discussion

The centre of pulley P1 is suspended from the ceiling by rope Ra.

Rope Rb hangs from the left of the pulley.

Rb supports a weight, W1, of value 4.

On the other side of the pulley P1 hangs rope **Rc**.

Attached to rope **Rc** is the centre of another pulley, **P2**.

On the left of P2 hangs rope **Rd**.

Rd is anchored to the floor.

On the other side of this pulley, **P2** is the rope **Re**.

Attached to Re is the centre of another pulley, P3.

On the left of this pulley hangs rope Rf.

Rf is anchored to the floor.

On the otherside of the pulley P3 the rope, Rg, support a weight W2.

Fig. 3. Sentential representation for the medium complexity pulley system. Emboldened items and lines are not shown in the experimental stimulus and are explained in the text

Fig. 4. lists the common rules for solving the problems. The four rules are used to compute incrementally the tensions in the strings connecting the weights, pulleys, ceiling and floor. For the problem in Figs. 1, 2 and 3, the minimum set of rule applications is: (1) PR1 to find the tension in rope Rb; (2) PR2 to find the tension in rope Rc, (3) PR3 to find the tension in rope Re; (4) PR3 to find the tension in rope Rg; (5) PR1 to find the value of weight W2. Production rule PR4 is not needed in this problem, as there is no weight hanging from two ropes.

PR1. A single rope supporting a weight will take the value of that weight as long as no other rope is supporting it.
PR2. In a pulley system the two ropes over, or under, the same pulley must have the same value.
PR3. The pulley takes the value of the sum of the two ropes over or under it. This same value is then taken by a rope supporting it or that hangs from it.
PR4. When a weight is supported by two ropes, its value is equal to the sum of the value of those two ropes.

Fig. 4. Common set of production rules for solving pulley problems

Consider step 3 in more detail to contrast how the rule PR3 is interpreted in each representation. The tension in rope Rc is known, either because it has just been computed or because a value is written for it in the representation ('4' if the solution is correct so far). With the diagram, Fig. 1., we see that Rc is connected to the centre of pulley P2 and see there are two ropes Rd and Re hanging from the sides of P2, hence the tension in each rope is 2, half that of Rc. With the table, Fig. 2, row 14 would have '4' written in the Hangs-Rc cell so one would repeat this value for pulley P2 and for P2 in row 15. Then '2' would be entered in each of the cells for Rd and Re, and this value repeated in row 21 for Re. In the sentential representation, Fig. 3, it is a matter of finding the statement about what hangs from Rc, Pulley P3, then searching for the statements about which ropes hang from P3, and then writing the value beside each one.

1.3 Locational Indexing in Diagrams, Tables and Sentential Representations

In addition to the diagrammatic and sentential representations that Larkin and Simon considered, the experiment uses a tabular representation in order to investigate why the problem difficulty varies with alternative representations. According to Larkin and Simon, diagrams are often better than sentential representations because items of information that are needed for inferences are often spatially co-located. Such location indexing facilitates the finding of relevant information and the picking of appropriate rules to apply. It is instructive to try comparing how easy (a) it is to find all the propositions in Fig. 3 that refer to components that are part of pulley system P2 with (b) identifying the same components in Fig. 1. Similarly, if one knows only the value of rope Rc, deciding which rule to apply given the sentential representation is less obvious than in the diagram.

The tabular representation was designed to exploit some of the benefits of location indexing, but rather than using co-location as in the diagrams, table rows and columns groups items of related information together more systematically than in the sentential representations, but without necessarily having them in close proximity to each other, nor necessarily separating them from unrelated items. Presumably, searching for information is easier in the table than the sentential representation and knowing which rules are applicable is facilitated by the patterns of cells so far completed. Hence, it may be predicted that the ease of problem solving with the tabular representation should be more similar to the diagram than the sentential representation.

Knowing which rules to apply is an important aspect of this task, so different problems were designed for the experiment that required different numbers of rule types, but about the same minimum number of rule applications to solve. Fig. 5a and 5b show, respectively, the simple and complex forms of the problem as diagrams. Fig. 1. shows the medium complexity version. Note that the number of pulleys, ropes and weights is the same for each problem. The complex problem requires all four rules in Fig. 4, the medium complexity problem does not require rule P4, and the simplest problem does not require either P3 or P4. It might be expected that the greater the complexity of a problem the harder it would be to solve and that the relative proportional benefit of the diagrams and the tables over the sentential representations will increase with increasing complexity.

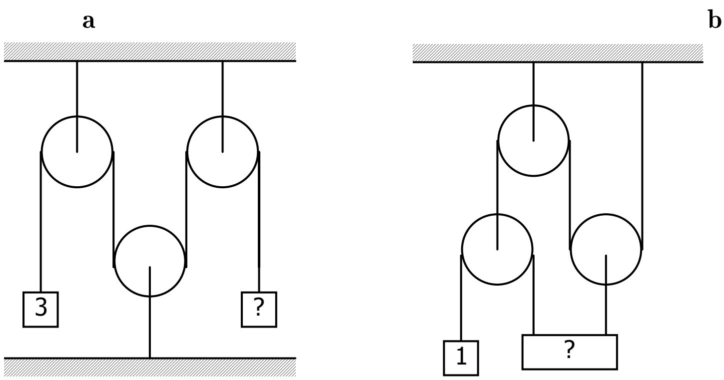

Fig. 5. Diagrammatic representations of (a) the simple and (b) the complex variants of the pulley system problem. See ***Fig. 1.*** for the medium complexity problem

2 Experiment

The first aim of the experiment was to determine how much harder people find solving pulley system problems with the sentential representation compared to the diagrammatic representation. The second aim was to test the hypotheses derived from explanations of the differences between alternative representations based on locational indexing. It was predicted that: (1) the tabular representation would be more like the diagrams than the sentential representations; (2) with increasing complexity of the problem the benefits of the diagrams and table over the sentential representation would interact with complexity, showing a magnified proportional advantage.

2.1 Design and Participants

The independent variables of the experiment were representation format and the problem complexity. A mixed 3X3 design was used with representation format (diagram, table, sentence) being a between participants factor and complexity (simple, medium, complex) serving as a within participants factor. The independent variable was solution time. Success rate was not used as a dependent variable, because performance was at ceiling in all nine conditions, as expected.

The 36 participants were volunteer non-engineering undergraduate students from the University of Nottingham. They were randomly allocated to one of the three representational format factors, with the proviso that the group sizes were made equal.

2.2 Materials and Procedure

The experimental session for each participant lasted no more than an hour and consisted of three phases. The first involved familiarization with the general nature of the task and the production rules in particular, as shown in Fig. 4. The second phase included training on each of the productions using mini-problems that comprised a small number of components. For each production there was a worked example and two problems that were completed by the participant, with assistant from the experimenter if necessary. The training examples were printed on sheets of paper. In order to make the training materials easily comprehensible the examples were presented using written descriptions and a diagram. Table participants were additionally shown tables for the mini-problems. However, the participants were only permitted to write on the particular representation that they would be using in the test phase. None of the participant had any difficulty learning how the productions worked.

In the main test phase the three problems of varying complexity were presented on printed sheets to the participants in a random order. The stimuli had only the one representation and the participants were instructed to use only that representation to solve the problem. The solution times were recorded.

3 Results

All the participants completed all the problems. Fig. 6 shows the median times for problem completion for each condition, with bars for quartiles giving information about the distributions of values. As can be seen, data for some conditions are highly skewed and in opposite directions. Thus, non-parametric statistics have been chosen for the analysis of the data and a conservative level of significance set, α=.01 (all tests are $p<.01$, unless otherwise stated). The reality of the impression given by Fig. 6 that for each level of complexity the significant order of difficulty of solution is sentence–table–diagram is supported by Jonckheere tests for order alternatives (3 conditions, 12 participants per group: low complexity, S=304; medium, S=248; high, S=372). Similarly, for the diagram and sentential representations the expected order of increasing difficult with greater complexity is present, as confirmed by Page tests for ordered alternatives (diagram, L=164; sentence, L=162.5). However, the increase in solutions times for the tabular representation is only marginally significant (L=155.5, $p<.05$).

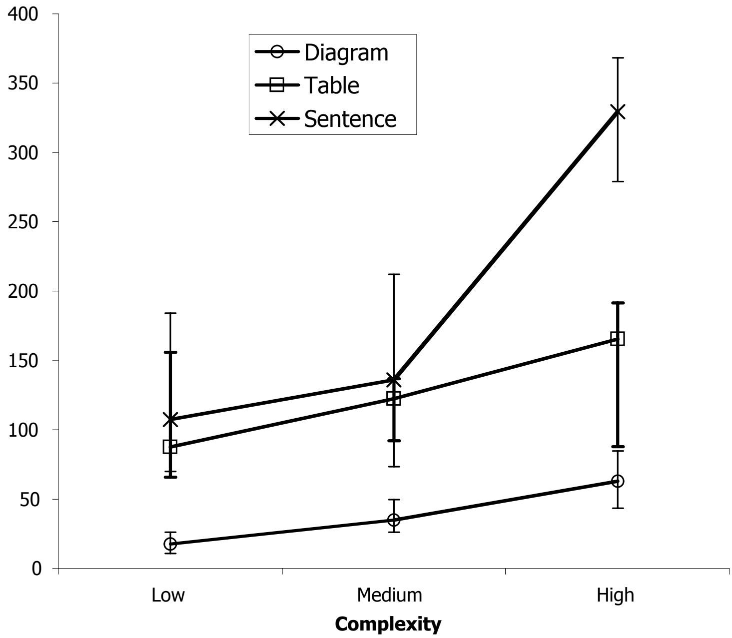

Fig. 6. Median solution times for the three representation types by the three levels of complexity. The bars on for each point show first and third quartiles. For clarity the bars for the table data points are thicker lines

As both the Jonckheere and Page tests merely determine whether at least one of the medians is greater than the next in order, and to investigate the possibility of an interaction of the representation and complexity factors, consider comparisons of the solution times between pairs of conditions. Tables 1 and 2 summarize pair wise Wilcoxon and Mann Whitney tests on combinations of complexity and representations conditions. For the table representation there are no significant differences between adjacent levels of complexities (at α=.01). For the sentence representation there is a significant difference in difficulty between the medium and high complexities only. The reverse is true of the diagram representation, with a significant difference occuring between the low and medium complexities. This is suggestive of an interaction between complexity and representation.

Table 1. Summary of Wilcoxon tests on pairs of levels of complexity under each of the representations

Comparison	Sentence	Table	Diagram
Low-medium	T=28.0, n.s.	T=23.0, n.s.	T=0.0, $p<.01$
Medium-high	T=0.0, $p<.01$	T=13.0, ($p<.05$)	T=22.0, n.s.

Table 2. Summary of Mann-Whitney tests on pairs of representations for each level of complexity

Comparison	Low	Medium	High
Sentence-table	U=63.5, n.s.	U=64, n.s.	U=14.0, $p<.01$
Table-diagram	U=0.0, $p<.01$	U=13.0, $p<.01$	U=15.5, $p<.01$

Table 2 shows that there are significant differences between the solution times of the diagrams users and the table users at all three levels of problem complexity. The difference between table users and the sentence users is only significant at the high complexity level. This implies that the tables are more similar to the sentential representations than they are to the diagrams.

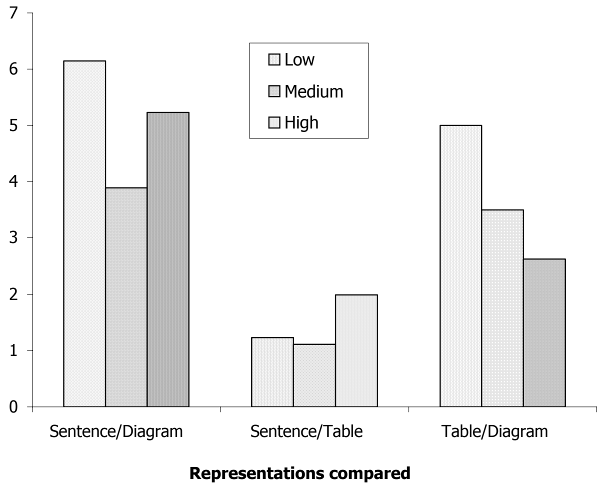

Fig. 7. Comparison of relative difficult of solutions with different pairs of representation across the level of complexity

Fig. 7. shows the ratios of the solutions times for each pair of representations for the three problem complexity levels, which illustrates the relative proportional difficulty of solution. The sentential representation is between four times and up to six times harder, approximately, than the diagrammatic representation. The tables are between three and five times harder than the diagrams. The greatest advantage of the diagram representation over the others is on the problem with the lowest complexity and not as predicted on the problem with the greatest complexity. The relative advantage of the tables over the sentential representations is no more than a factor of two.

4 Discussion

In the classic studies on isomorphs of the Tower of Hanoi the problem difficulty varied by up to 16 times. Unlike those studies, the different representations in this experiment varied only in the format of the representation and not in the set of rules used by the participants. Nevertheless, solving the pulley problem with the tabular and sentential representations was, respectively, up to five or six times that of using the diagrammatic representation. Just changing the format of a representation produces a substantial effect, which is something that needs to be explained.

Whereas the overall pattern of problem difficulties was as expected, the detailed comparisons of performance challenges the predictions derived from considerations of locational indexing. On this basis, it was expected that the table representations would be superior to the sentential representation and comparable to the diagrammatic representation, because the structure of the table provides a strong spatial indexing of the information through the use of rows and columns, even though the Euclidean distance between related items is not necessarily small. The difference between the table and diagram cannot simply be explained by a failure of the table participants to use the indexing system of the table, because the success rate was high and finding information in tables is a well-practiced skill for university undergraduates. Thus, it is inferred that the effect of locational indexing on the search for items of information does not provide a full account of the benefit of diagrams.

Similarly, the prediction that the benefits of locational indexing for diagrams would be greater the more complex the problem was not supported. There is about an extra one-fold decrease in the proportional difficulty of the sentential representation compare to the diagram with each increase in level of complexity. The complexity of the problems varied most substantially in terms of the number of types of rules that were needed to solve the problems. The easier recognition of rules due to location indexing does not seem to have been a major factor on the relative problem difficult between representations.

How can this state of affairs be explained? A possible account for the failure of both predictions is to consider the specific nature of the representations and the problem solving strategies that are being used with each of them. A clue comes from the lack of a significant increase in the solution time with the table representation with increasing problem complexity, although the medians do have an increasing trend. This implies that the table participants are using an approach that is relatively independent of the complexity of the problems. Examining the tables for the three problems, it is oblivious that they are all more similar to each other than the diagrams are to each other (which is why only one table is reproduced here). The minimal solution path for the medium complexity table is show by the arrows in Fig. 2. Note that there are alternative paths branching off the minimal path that do not contribute to the required answer. It seems that the participants may have been progressing thorough the table completing all the empty cells without regard for the minimal path. Inspection of the solution sheets used by the participants confirms this was the case. In 34 out of the 36 solutions produced all the cells were completed. In the two cases where some cells were not completed the participants appeared occasionally to be following the minimal path, but in so doing they will have needed to have searched for the dead end branches and reasoned about avoiding them. Thus, it is clear that all the table users executed exhaustive searches and nearly all completed all the information in the tables. This would have added substantially to the computation time needed for solution, even though information required for any particular inference would have been easy to locate using the indexing of rows or columns.

In contrast the diagram participants seem to have been able to follow the minimal path through their representation. Examining Fig. 1, and 5a and 5b, the overall pattern connectivity of ropes and pulleys throughout the system is easily comprehended, so it is possible to readily spot paths that do not progress towards the unknown weight. Hence, a more efficient strategy that follows the minimal path could be executed, unlike the tabular representation.

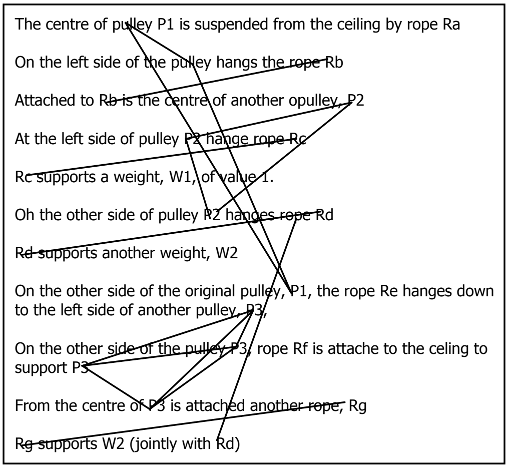

Fig. 8. Complex problem in the sentential representation. Lines show co-references to the same components of the pulley system

The ease of comprehension of the structure of the pulley system also provides an explanation of why the relative proportional benefit of the diagram representation was greatest with the simplest problem, rather than the most complex, contrary to what was predicted by a straightforward locational indexing explanation. In addition to requiring only two rules to solve, the simple problem is symmetric and a single string connects the given and unknown weights. Both of these features are obvious in Fig. 5.a but are not apparent in either the tabular or sentential representation for the same problem. Spotting that the problem is symmetrical, a diagram participant may simply have inferred that the given and unknown would be identical. Spotting that the same string directly connects both weights, a participant may simply infer that the constant tension means the weights will be equal. These short cut inferences are not possible with the diagrams for the more complex problems. Hence, the greatest relative benefit of the diagram, about six times faster than the sentential representation and five times faster than the table, was on the simplest problem.

Locational indexing does provide an explanation of the jump in difficult of the complex problem in the sentential representation compared to the other problems under the same representation. The lines in Figs. 3 and 8 show the distribution of labels for the components that are involved in the minimal solution to the medium and complex problems, respectively. Clearly, search for information in the simple problem will be easier than in the complex one, because items of information about individual components in the simple problem are separated into different regions of the representation, whereas they are mingled amongst each other and more widely dispersed in the complex representation. As expected this has a more than proportional effect on the difficult of the complex problem.

Locational indexing does provide partial explanation of why some types of representation are better than others and the impact this has on solutions to problems with different complexities. However, to adequately account for the pattern of results obtained in this experiment it was necessary to consider the strategies that can be used with each representation in general and the specific characteristics of representations of particular problems. The nature of the representations impacts on how their users solve the problems, with the diagram supporting the selection of information pertaining to the most direct path between initial state and goal, whereas the table encourages a less direct exhaustive inferences over all of the information, because the overall structure of the system and problem is hidden, so apparently preventing quick look ahead and planning. The opacity of the table, in this respect, is similar to that of the sentential representation. The similarity of performance on the tabular and sentential representations, plus the contrast with the diagrammatic representation, suggests the benefit to problem solving of a representation that reveals the overall structure of a problem may be at least as important as the use of locational index of information.

This conclusion is consistent with previous claims about representational systems and problem solving. Koedinger and Anderson [5] have shown empirically and computationally the importance of *diagrammatic configuration schemas* (DCS) as memory structures that expert geometry problem solvers use to encode their knowledge. A DCS has a diagram (e.g., weight suspended from two ropes) around which information is stored about the relations (e.g., weight is sum of the rope tensions) and necessary conditions for inferences to be made using the rules (e.g., both tensions given). By visually matching DCS diagrams with parts of the problem diagram, valid inferences can quickly be identified, which supports planning and a working forward strategy. Koedinger [4] suggests that DCSs may have a similar role in diagrams for pulley system problems. The outcome of the present experiment supports this claim and it is noteworthy as participants were not experts in mechanics or the domain, although they did receive training. It is possible that the table participants had also acquired and used something equivalent to DCSs but for cells configurations in the tables. However, the potential benefit that table users could gain from table configuration schemas is likely to be lower than DCSs, because patterns of cells in a table are less easy to discriminate than images of pulleys, ropes and weights.

Finally, there are implications for the design of effective representational systems to support problem solving and learning. First, although Larkin and Simon's [8] explanation of the value of diagrams must be taken as part of a larger and more complex account, the design of a good representation should use location indexing as a means to coordinate information that will be useful to problem solving. It is unlikely that a representation without such a scheme will be effective. Second, the design of the representation should attempt to make the underlying structure of the problem readily apparent to the user. Such a representation is likely to support planning and rule selection well. This is consistent with Cheng's [2] claims about the nature of effective representational systems. Specifically, an effective representation should attempt to encode the underlying relations, or meaning, of a domain directly in the structure of its representational schemes. Representations, having such semantic transparency, have been shown to substantially improve for conceptual learning in a number of knowledge rich topics in science and mathematics [2].

Acknowledgements

Thanks must go to Anna Mallalieu who prepared the materials and collected the data for this experiment. This research was largely conducted in the School of Psychology at the University of Nottingham.

References

[1] Cheng, P.C.-H., *Scientific discovery with law encoding diagrams.* Creativity Research Journal, 1996. **9**(2&3): p. 145-162.

[2] Cheng, P.C.-H., *Electrifying diagrams for learning: principles for effective representational systems.* Cognitive Science, 2002. **26**(6): p. 685-736.

[3] Cheng, P.C.-H. and H.A. Simon, *Scientific discovery and creative reasoning with diagrams.*, in *The Creative Cognition Approach*, S. Smith, T. Ward, and R. Finke, Editors. 1995, MIT Press: Cambridge, MA. p. 205-228.

[4] Koedinger, K.R., *Emergent properties and structural constraints : Advantages of diagrammatic representations for reasoning and learning*, in *AAAI Technical Report on Reasoning with Diagrammatic Representations (SS-92-02)*, N.H. Narayanan, Editor. 1992, AAAI: Menlo Park, CA.

[5] Koedinger, K.R. and J.R. Anderson, *Abstract planning and perceptual chunks: Elements of expertise in geometry.* Cognitive Science, 1990. **14**: p. 511-550.

[6] Kotovsky, K., J.R. Hayes, and H.A. Simon, *Why are some problems hard?* Cognitive Psychology, 1985. **17**: p. 248-294.

[7] Larkin, J.H., *Display-based Problem Solving*, in *Complex Information Processing: The Impact of Herbert A. Simon*, D. Klahr and K. Kotovsky, Editors. 1989, Lawrence Erlbaum Associates: Hillsdale, New Jersey. p. 319-341.

[8] Larkin, J.H. and H.A. Simon, *Why a diagram is (sometimes) worth ten thousand words.* Cognitive Science, 1987. **11**: p. 65-99.

[9] Scaife, M. and Y. Rogers, *External cognition: how do graphical representations work?* International Journal of Human-Computer Studies, 1996. **45**: p. 185-213.

[10] Zhang, J. and D.A. Norman, *The representation of relation information*, in *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society.*, A. Ram and K. Eiselt, Editors. 1994, Lawrence Erlbaum: Hillsdale, NJ. p. 952-957.