

Extending Task Analytic Models of Graph-based Reasoning: A Cognitive Model of Problem Solving with Cartesian Graphs in ACT-R/PM

David Peebles (djp@psychology.nottingham.ac.uk)

Peter C.-H. Cheng (pcc@psychology.nottingham.ac.uk)

ESRC Centre for Research in Development, Instruction and Training,
Department of Psychology, University of Nottingham, Nottingham, NG7 2RD, U.K.

Abstract

Models of graph-based reasoning have typically accounted for the variation in problem solving performance with different graph types in terms of a task analysis of the problem relative to the particular visual properties of each graph type (e.g. Lohse, 1993; Peebles, Cheng & Shadbolt 1999, submitted). This approach has been used to explain response time and accuracy differences in experimental situations where data are averaged over experimental conditions. An experiment is reported in which participants' eye movements were recorded while they were solving various problems with different graph types. The eye movement data revealed fine grained fixation patterns that are not captured by current analyses based on optimal fixation sequences. It is argued that these patterns reveal the effects of working memory limitations during the time course of problem solving. An ACT-R/PM model of the experiment is described in which a similar pattern of eye fixations is produced as a natural consequence of the decay in activation of perceptual chunks over time.

Introduction

A recent development in the field of cognitive modelling is the proposal of frameworks to understand interactive behaviour with external representations and artifacts. Gray (2000; Gray & Altmann, 2000), for example, has proposed the *Cognition-Task-Artifact triad* within which to characterise behaviour in human-computer interaction tasks in terms of the complex interaction of three primary elements: the cognitive abilities of the user, the representational and physical properties of the artifact, and the specific requirements of the task. This framework has recently been developed by Byrne (in press) to encompass the perceptual and motor capabilities of the user. Similarly, in the area of graph-based reasoning, Peebles, Cheng and Shadbolt (1999, submitted) have proposed the *Graph-Based Reasoning* (GBR) model incorporating these three elements to account for the ability of users to retrieve and reason about information in different types of Cartesian co-ordinate (x-y) graph.

The primary purpose of these frameworks is to inform the development of detailed cognitive models of the cognitive, perceptual and motor processes involved in the tasks under study. In contrast with *cognitive task analysis* (Gray & Altmann, 2000) which simply specifies the cognitive steps required to perform the task, the construction of cognitive process models that are grounded in cognitive theory allows the incorporation and testing not only

of relevant cognitive factors such as the required declarative and procedural knowledge, the strategies adopted, and the limitations of working memory but also perceptual/motor factors such as mouse movements and shifts in visual attention.

One such model in the area of graph-based reasoning is UCIE (Lohse, 1993). By adding cognitive parameters to the GOMS class of task analysis techniques (Card, Moran, & Newell, 1983; Olson & Olson, 1990; John & Kieras, 1994), Lohse produced a model which simulated certain question answering procedures using line graphs, bar graphs and tables and predicted question answering times by assuming an optimal sequence of eye movements to scan the graphical representation that minimised the number of saccades and fixations to reach the target location.

More recently, the GBR model (Peebles et al., 1999, submitted), employed a similar set of assumptions to account for data from experiments investigating the various interacting factors affecting reasoning with different types of *informationally equivalent* (Larkin & Simon, 1987) Cartesian graph. Figure 1 shows examples of the types of graph used. Both graphs encode the same two functions between time and the variables A and B. The *Function* graph in Figure 1a represents time on the x axis and the A and B variables on the y axis whereas the *Parametric* graph in Figure 1b represents the A and B variables on the x and y axes respectively while time is plotted as a parameterising variable along the curve.

Our experiments have revealed significant differences in both response time and error rates between users of the two graph types for a wide range of questions (Peebles et al., 1999, submitted). The GBR model has been successful in explaining why such differences occur with these graph types despite their numerous visual and conceptual similarities. Using the graphs in Figure 1 as an example, we found that when participants were asked to retrieve the value of A when B equals 1, parametric graph users' responses were significantly more rapid and accurate than those of function graph users. The GBR model accounts for these differences in terms of the optimal visual scan path the users follow through the graph. The variability in responses is apparent from the sequence of hypothesised saccades in the two graphs. In Figure 1a, the sequence of saccades is m, n, o , whereas in Figure 1b the process requires just two saccades, as shown by the

line sequence *a*, *b*. The higher probability of an erroneous response using the function graph was explained by the additional number of possible incorrect saccades that the function graph users may make.

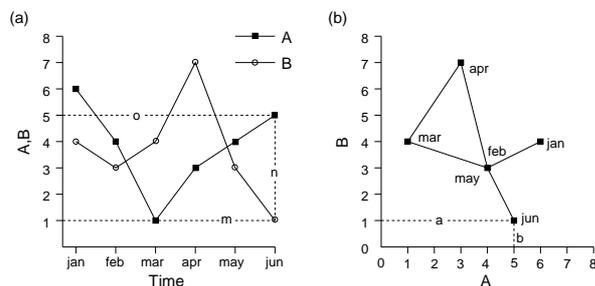


Figure 1: Informationally equivalent function and parametric graphs

The optimality assumptions incorporated into the UCIE and GBR models are useful as they allow the prediction of response times and provide an explanation of variations in mean RT and error data for different graph types. It is clear, however, that such assumptions do not take important cognitive factors such as working memory limitations or strategic decisions into account. For example, it is likely that, during the time course of a complex graph-based reasoning problem, certain information may be forgotten and have to be rescanned. In addition, given that graph users are aware that information is available for rescanning at all times, it is possible that they may trade off additional saccades for a reduction in working memory load. If this is the case, then the current analyses may miss out an important level of detail which sheds light on the cognitive load that these tasks are imposing and the strategies by which graph users optimise their retrieval procedures. Furthermore, if the goal is to produce detailed cognitive models of these tasks, then information at this level of detail will provide valuable constraints on such models.

In this article we report the results of a graph-based reasoning experiment designed to address these issues. In the experiment, participants were asked to solve simple tasks using function and parametric graphs which, based on the optimality assumptions described above, would be predicted to produce varying response patterns by requiring different optimal fixation sequences. To determine whether these optimality assumptions are justified, some of the participants' eye movements were recorded as they solved the problems. We show that, although the RT and error data are in line with the GBR model's predictions, certain patterns in the eye movement data do not follow the optimal sequence predicted by the model which may be interpreted as indicating the effects of working memory limitations. We then describe an ACT-R/PM model of the experiment in which a similar pattern of eye fixations is produced as a natural consequence of the decay in activation of perceptual chunks over time.

Experiment

One of the most common tasks carried out when using a graph is to elicit the value of one variable corresponding to a given value of another. This task was chosen for the experiment as it is so widely performed and because the procedures involved are relatively simple. The knowledge required to carry out these tasks is primarily the sequence of fixations required to reach the *given location* in the graph representing the given value of the given variable and then from there to the *target location* representing the corresponding value of the required variable. In previous research, however, we have discovered that the effectiveness of a particular graphical representation for retrieving the required information depends on the details of the task, i.e. which variable is given and which is sought (Peebles et al., 1999, submitted).

Method

Participants and materials Forty-nine undergraduate and postgraduate psychology students from the University of Nottingham were paid £3 to take part in the experiment. Of these, four were paid an additional £2 to have their eye movements recorded while they carried out the experiment. The experiment was carried out using PC computers with 17 in displays. The eye tracker employed in the experiment was an SMI iView system using a RED II desktop pupil/corneal reflectance tracker with a 50 Hz sampling rate recording eye movements at 20 ms intervals remotely from a position in front of the computer display. In addition to the system's own automatic head movement compensation mechanism, participant's heads were restrained in a frame fixed to the table to reduce recording error due to head movement.

The stimuli used in the experiment were four graphs, shown in Figure 2, depicting the amount of UK offshore oil and gas production between two decades. Participants were seated approximately 80 cm from the 72 ppi computer display. The graphs were 15.5 cm square (including axis labels), corresponding to approximately 11.1° of visual angle. The characters representing variable values were 0.4 cm high (approximately .21° of visual angle) while those for the axis labels and questions were 0.4 cm and 0.5 cm high (approximately .29° and .36° of visual angle) respectively. Axis ticks were spaced 1.5 cm (approximately 1.1° of visual angle) apart.

The graphs and data sets were designed so that the independent variable (IV—year) and the two dependent variables (DVs—oil and gas) all had ten values ranging from 0 to 9 and that the full range of these values was represented by the data points for oil and gas in both decades. A set of 120 questions was produced using all of the values for the three variables in both decades. The questions had the same basic structure, giving a variable's value and requiring a corresponding variable value.

Design and Procedure The experiment was a mixed design with one between-subjects variable, (graph type) and two within-subjects variables (question type and

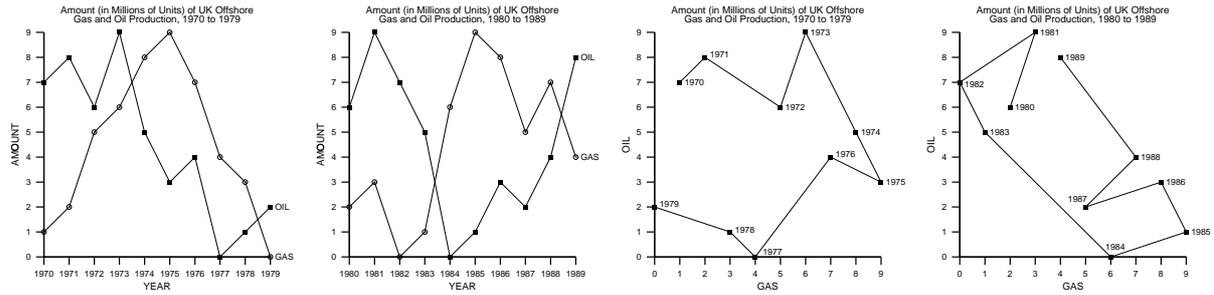


Figure 2: Function and Parametric Graphs Used in the Experiment

graph number). Participants were randomly allocated to one of the two graph type conditions. On each trial, a graph would be presented with a question above it. For example, the question GAS = 2, OIL = ? required the value of oil when gas is equal to 2 to be found. Participants were instructed to enter only the final number of the target year when a year value was required. Each element of the question was centered on a co-ordinate point which remained invariant throughout the experiment with approximately 3.5 cm (approximately 2.5° of visual angle) between the centres of adjacent text items. Together with the graph and question, a button labelled *Answer* appeared in the top right corner of the window. Participants were instructed to click on this answer button as soon as they had obtained the answer to the question. Response times were recorded from the onset of a question to the mouse click on the answer button. When this button was clicked upon, the button, graph and question were removed from the screen and a circle of buttons labelled clockwise from 0 to 9 appeared centered on the answer button. Participants entered their answers by clicking the appropriate number button. When the number button was clicked, the next graph, question, and answer button appeared on the screen. This method was devised so that participants in the eye movement study would not have to take their eyes away from the screen to enter answers, as would be the case if using the keyboard. Before starting the experiment, participants were asked to answer the questions as rapidly and as accurately as possible and were given time to become familiar with the graphs and practice entering numbers using the circle of number buttons and mouse.

Results

Response accuracy and latency data The proportions of correct responses and mean response times (RTs) for each of the question types for the two graphs in each condition are presented in Table . The data reveal high levels of accuracy for all three question types in both graph conditions. An ANOVA on the response accuracy data, however, revealed a significant effect of question type $F(2, 239) = 28.187, p < 0.01, MSE = 0.123$. Although there is little variability in response accuracy between conditions, RTs vary significantly both between condi-

tions and within each condition according to the type of question being attempted. An ANOVA on the RT data revealed significant effects of question type $F(2, 239) = 18.447, p < 0.01, MSE = 4974038$, and graph number $F(1, 239) = 5.76, p < 0.05, MSE = 1223302$ and significant interactions between graph type and question type $F(2, 239) = 36.314, p < 0.01, MSE = 9791754$ and between graph type, question type and graph number $F(2, 239) = 3.913, p < 0.05, MSE = 466423$.

The results of this experiment are in line with predictions of the GBR model which explains these differences in terms of a detailed task analysis and the assumption of different optimal scan paths through the graphs to the target location. However, as the main focus of this article is the eye movement data and the ACT-R/PM model, no analysis of these data will be provided here. A full description of the GBR model, its predictions and analyses of these and similar tasks can be found in previous articles (Peebles et al., 1999, submitted).

Eye movement data To analyse the eye movement data, the display was divided into three regions in a manner similar to that employed by Carpenter and Shah (1998). The regions, shown in Figure 3, were the same for all four graphs and define the relevant units of the display for the fixation analysis: *question*, *graph pattern*, and *answer* buttons. Dividing the display in this manner allows an analysis of the frequency and duration of fixations on the question and the graph and also the pattern of transitions between these regions during the time course of an individual trial. For the analysis, we adopt Carpenter and Shah's (1998) term *gaze* to refer to a sequence of consecutive fixations on a display region that is unbroken by fixations in other regions. The raw x and y coordinate data from the eye tracker were aggregated into gazes, the minimum duration of which, based on a preliminary study of the data, was defined as 100 ms. This value was sufficiently large to eliminate most saccades, short fixations and noise in the data while still capturing all the relevant fixations. The data from each participant were analysed so that fixations of 100 ms or more in each region were recorded and a scan path consisting of the sequence of gazes from question to graph to answer button regions for each trial was produced. From a total

Table 1: Correct response proportions and mean RTs for each question type in Function and Parametric conditions

Question Type	Correct Response				Response Time (s)			
	Function		Parametric		Function		Parametric	
	Graph 1	Graph 2	Graph 1	Graph 2	Graph 1	Graph 2	Graph 1	Graph 2
DV-DV	.911	.898	.925	.891	5.93	6.18	4.89	4.83
IV-DV	.971	.982	.993	.989	5.01	5.07	4.88	5.17
DV-IV	.943	.930	.939	.925	5.51	5.80	4.38	4.41

of 480 trials in the eye movement study, 28 were removed due to the analysis producing an unusable scan path (e.g. containing only one gaze recorded before reaching the answer region). The rest of the trials were placed into four categories according to the number of transitions from question to graph regions. Of these trials, 37.1% involved only one transition from the question to the graph, 48% involved two such transitions, 11.9% involved three transitions and 2.8% involved four or more transitions. An analysis of the data showed that these categories were not related to specific graph type or question type conditions. Figure 4 shows the average gaze duration on the question and graph regions for the first three categories, (the fourth was removed from the analysis due to its relative rarity).

Figure 4 shows that participants took on average just over 400 ms to read the three elements of the question. The fact that this time is consistent across all three transition groups is strong evidence that the categories do not indicate different problem solving strategies. If the transition categories simply reflected the use of different strategies, (e.g. to switch between the question and graph, reading individual question elements and then identifying their locations in turn), then it is likely that the first gaze duration on the question would be different for each category. In the 1 transition trials, participants took approximately 2.28 s to scan the graph before entering an answer whereas in the other trials, participants looked at the graph for a shorter time before looking again at the question for approximately 300 ms. In the 2 and 3 transition trials, participants then looked at the graph again for over a second before entering an answer or looking for a third time at the question, respectively. Participants in the 3 transition trials then scanned the graph for a third time before entering an answer.

These gaze patterns become clearer when interpreted in the light of the GBR model task analysis of the problem solving procedure. After reading the three items of information in the question, (the *given* variable, *given* value and *required* variable), in order to answer the questions, the user must carry out several graph location tasks in a specific order: (1) find the *start location* determined by the given variable, (2) find the *given location* representing the given value, (3) find the *required location* representing the required variable, (4) read off the *required value* which is the answer to the question. These four steps require the three items of information

in the question to be utilised at different stages of the problem solving process. It is likely therefore, that once the question has been read, the probability that each item can be recalled when required decreases as processing proceeds and that the item most likely to be forgotten is at the final stage of the process.

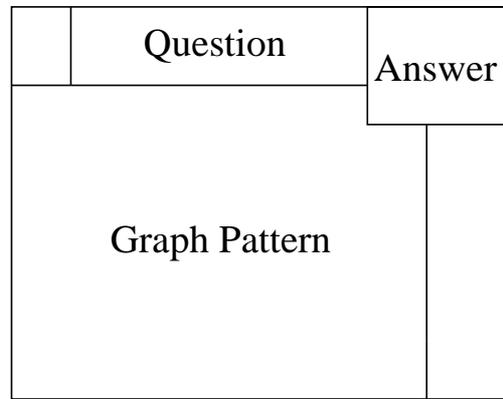


Figure 3: Three regions of the display defined for the fixation analysis

According to this analysis, therefore, the 37.1% of trials requiring only 1 transition suggests that participants were able to retain all three question items while solving the problem in these trials, whereas in the 48% of trials requiring 2 transitions and the 11.9% of trials requiring 3 transitions, participants had failed to retrieve one and two question items respectively. A detailed analysis of the raw eye movement data for individual trials and the duration of the first gaze on the graph for each transition group in Figure 4 support this account. In the 2 transition trials, participants took 1.71 s to identify the start location, find the given location before having to look again at the required variable in the question to find the required location and solve the problem. In the 3 transition trials, however, participants were looking at the graph for only 860 ms in order to identify the start location before having to look again at the given value in the question to find the given location and then having to look again at the required variable in the question to find the required location and solve the problem. It follows from this analysis that the probability of recalling each question item can be computed from the data. The probability of re-

calling the given variable is .972, that for the given value is .852 while that for the required variable falls sharply to .372.

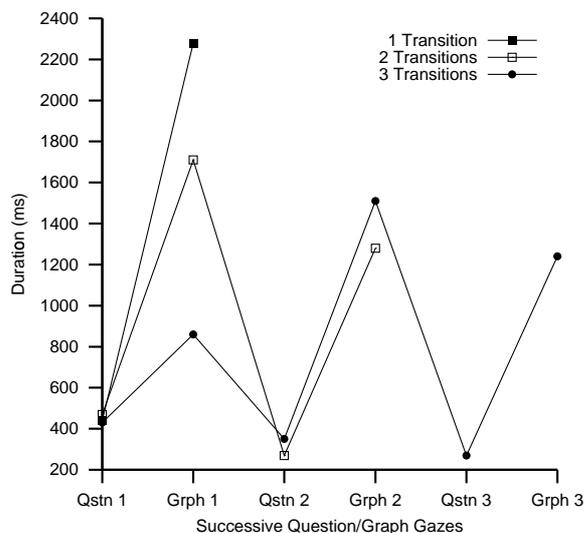


Figure 4: Mean gaze duration on Question (Qstn) and Graph (Grph) regions as a function of the number of transitions required to answer the question

The ACT-R/PM Model

One of the main aims of this research is to construct models of graph-based reasoning that are grounded in cognitive theory and incorporate cognitive factors such as memory decay and interference together with perceptual/motor components that provide realistic interactive behaviour. With these additional factors, it is expected that the optimal analyses provided by the GBR model may be extended to account for detailed experimental data such as eye movements and gaze durations. ACT-R/PM has the required cognitive and perceptual mechanisms with which to develop such models. More specifically, ACT-R's declarative memory has an activation-based retrieval process and includes a mechanism by which the base-level activation of a chunk decreases over time. In addition, the vision module of ACT-R/PM allows the activation of a chunk to be increased when visual attention is refocused upon the visual object that it represents. Space limitations preclude further elaboration of these mechanisms here but detailed discussions can be found in Anderson & Lebière (1998) and Byrne & Anderson (1998).

An ACT-R/PM model of the experiment was constructed which was able to interact with an exact replica of the software used to run the experiment. The model consists of productions to read the three question components and a set of productions to locate the given variable, given value and required variable in the correct order. A further set of productions was created to allow the model to enter an answer using the mouse and answer buttons once the required value had been obtained.

Given a top level goal to do a trial of the experiment, the model sets a subgoal to read the question. When each question component has been read, a chunk representing that component is created in declarative memory. When all three elements of the question have been read, the model carries out a sequence of three main operations: (1) to identify the starting location (x axis, y axis, or plot region) corresponding to the given variable, (2) to find the given location (x axis tick label, y axis tick label, or plot symbol label) corresponding to the given value, (3) to find the required location (x axis tick label, y axis tick label, or plot symbol label) representing the required variable's value. The sequence of these operations is determined by the state of the problem represented by the current location of visual attention in the graph which is a slot in the top level goal. At each of these main steps in the problem solving procedure, the production initiating the step must retrieve the declarative chunk representing the relevant question element created when the question was read. As the base-level activation of these chunks is decreasing over time, however, the probability that the chunk will be retrieved reduces as problem solving proceeds. If the production fails to retrieve the chunk, a second production fires which stores the current location of attention, sets a subgoal to reread the appropriate question element, thereby increasing the activation of the associated chunk, and then returns the focus of attention to its previous location. With this chunk now sufficiently active, the first production is able to fire and the problem solving process continues. When the model locates the required value, a subgoal is set to enter the answer using the mouse and the answer buttons. When the answer has been entered, the next trial begins.

Simulation An initial test of the model in the parametric graph condition has been conducted. In this test, the model was run through the entire experiment five times and the number of times the question elements were reread was recorded for each run. Of the total of 600 trials, the given variable was never reread as the activation of the given variable chunk was always highly active at the start of a trial, while the given value had to be reread 59 times and the required variable 314 times. As with the eye movement data, the probability of recalling each question item can be computed from these scores. The probability of recalling the given variable is 1, whereas that for the given value is .9 while that for the required variable is .48. Table 2 shows these probabilities against those computed from the data and reveals a close similarity between the patterns of recall probabilities.

An analysis of the model's behaviour revealed that the large decrease in the recall probability between the given value and the required variable is due not only to the time interval between the recall of the two question elements, but also because when the given value was not recalled, the additional time required to reread the question element ensured that the required variable was also not recalled.

Table 2: Probability of recall for the three question elements computed from the eye movement data and the ACT-R/PM model

Question Element	Data	Model
Given Variable	.972	1.00
Given Value	.852	.902
Required Variable	.372	.477

Discussion

Reasoning with Cartesian graphs involves a complex interaction between the perceptual and cognitive abilities of the reasoner, the visual properties of the graph, and specific task requirements. Models of graph-based reasoning have largely focussed on providing detailed task analyses in relation to the the visual properties of the graph and explaining differences in performance in terms of the interaction of these two elements. By incorporating cognitive factors as the user's knowledge, strategies and working memory capacity and perceptual/motor components into graph-based reasoning models, the explanatory and predictive power of these models can be increased and greater insights into the processes and factors affecting these complex interactions can be obtained.

In such a visual domain as graph-based reasoning, eye movements are an important source of information regarding how people acquire and process graphical information during problem solving. The experiment and the eye movement study reported here show how eye movement data can be used to make hypotheses about effects of working memory limitations on problem solving with graphs. The scan paths revealed by the eye movement study show that the optimality assumptions of current models serve as an approximation that obscures the complex sequences of saccades made by individuals.

In contrast, the ACT-R/PM model of the experiment provides a detailed explanation of these scan paths in terms of the decay of base-level activation of perceptual chunks during the time course of problem solving. According to the model, participants initially encode all three elements of the question but are required to re-encode parts of it as the problem progresses, with the probability of re-encoding increasing over time.

This research shows that eye movement data can provide important information concerning the cognitive factors underlying reasoning with external representations that were commonly overlooked by optimal task analytic models and also how a cognitive model such as ACT-R/PM can account for these data.

Acknowledgements

This research is funded by the UK Economic and Social Research Council through the Centre for Research in Development, Instruction and Training.

References

- Anderson, J. R., & Lebière C. (Eds.) (1998) *The Atomic components of thought*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Byrne, M. D., (in press). ACT-R/PM and menu selection: Applying a cognitive architecture to HCI. *International Journal of Human-Computer Studies*.
- Byrne, M. D., & Anderson, J. R. (1998). Perception and action. In J. R. Anderson & C. Lebière (Eds.) *The Atomic components of thought*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Card, S. K., Moran, T. P., & Newell, A., (1983). *The psychology of human-computer interaction*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Carpenter, P. A., & Shah, P. (1998). A model of the perceptual and conceptual processes in graph comprehension. *Journal of Experimental Psychology: Applied*, 4, 75–100.
- Gattis, M., & Holyoak, K. (1996). Mapping conceptual to spatial relations in visual reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 231–239.
- Gray, W. D. (2000). The nature and processing of errors in interactive behaviour. *Cognitive Science*, 11, 205–248.
- Gray, W. D., & Altmann, E. M. (2000). Cognitive modeling and human-computer interaction. In W. Karwowski, (Ed.), *International encyclopedia of ergonomics and human factors*. New York: Taylor & Francis, Ltd.
- John, B. E., & Kieras, D. E., (1994). *The GOMS family of analysis techniques: Tools for design and evaluation*. (Tech. Rep. CMU-HCII-94-106). Pittsburgh, PA: Carnegie Mellon University, Human-Computer Interaction Institute.
- Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, 11, 65–100.
- Lohse, G. L. (1993). A cognitive model for understanding graphical perception. *Human-Computer Interaction*, 8, 353–388.
- Olson, J. R., & Olson, G. M., (1990). The growth of cognitive modeling in human-computer interaction since GOMS. *Human-Computer Interaction*, 5, 221–265.
- Peebles, D., Cheng, P. C.-H., & Shadbolt, N. (1999). Multiple processes in graph-based reasoning. In *Proceedings of the Twenty First Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum.
- Peebles, D., Cheng, P. C.-H., & Shadbolt, N. (submitted). A model of graph-based reasoning: Integrating the role of visual features, knowledge and search.
- Shah, P., & Carpenter, P. A. (1995). Conceptual limitations in comprehending line graphs. *Journal of Experimental Psychology: General*, 124, 43–62.