

The Use of Spatial Cognition in Graph Interpretation

Susan Bell Trickett (trickett@itd.nrl.navy.mil)

J. Gregory Trafton (trafton@itd.nrl.navy.mil)

Naval Research Laboratory, Code 5515
4555 Overlook Drive, Washington DC 20375

Abstract

We conducted an experiment to investigate whether spatial processing is used in graph comprehension tasks. Using an interference paradigm, we demonstrate that a graph task interfered more with performance on a spatial memory task than on a visual (non-spatial) memory task. Reaction times showed there was no speed-accuracy tradeoff. We conclude that it was the spatial nature of the graph task that caused the additional interference in the spatial memory task. We propose that current theories of graph comprehension should be expanded to include a spatial processing component.

Keywords: Graph comprehension, spatial reasoning.

Introduction

The ability to interpret graphs is a crucial skill in today's data-rich world. For example, graph skill has been shown to be predictive of success in learning science (Schunn et al, in press). Not coincidentally, in developing its recent national standards, the National Council of Teachers of Mathematics has emphasized the importance of proactively teaching students of all ages to interpret graphs and use them to make inferences (NCTM 2000). Beyond the classroom, graphs also occupy an important part in people's lives, from daily encounters in the popular press to the tools used by highly-skilled scientists, as well as a range of usages in between.

Understanding the processes by which people comprehend, interpret, and use graphs is an important first step in both improving graph interpretation skills and informing effective graphic design (e.g., Tufte, 2001; Tan & Benbasat, 1990). We have argued elsewhere (Trickett & Trafton, 2006) that current theories of graph interpretation are inadequate on two counts. First, most current theories address only very simple graph interpretation tasks performed on very simple graphs, such as reading a value off the y-axis of a bar graph (e.g., Cleveland, 1985) or identifying a trend from a line graph (e.g., Zacks & Tversky, 1997). "Real-world" graph interpretation, on the other hand, involves a range of complex tasks, such as integrating information and making predictions. Moreover, it uses highly complex graphic representations. Second, current theories either assume people reason from graphs using exclusively propositional representations and processes (Lohse, 1993; Peebles & Cheng, 2003; Pinker, 1990), or are non-committal, failing to specify what kind of representation and processes are involved (Freedman & Shah, 2002; Roth & Bowen, 2003).

In contrast to these theories, we have evidence that, in complex tasks and domains at least, people use a great deal of *spatial processing* when interacting with graphs (Trafton et al, 2000; Trickett et al., in press). Furthermore, in Trickett

and Trafton (2006), we have argued, based on task analysis, that spatial processing is involved in a range of graph interpretation tasks, regardless of the graph's complexity—in fact, in most graph tasks that move beyond either reading values from an axis or making purely perceptual judgments (such as comparing two adjacent bar heights). In this paper, we present an experiment that provides empirical support for our claim that graph interpretation involves spatial processing, even when the task is simple and the graph itself is uncomplicated.

What do we mean by spatial processing? Baddeley was instrumental in establishing the distinction between verbal and spatial processing (Baddeley & Lieberman, 1980). Spatial processing involves "the internalized reflection and reconstruction of space in thought" (Hart & Moore, 1973).

Operationally, we define spatial processing in two ways. Spatial processing involves maintaining spatial information (e.g., the relative locations of objects) in working memory (so-called spatial working memory). Instances of spatial processing can therefore be identified by means of task analysis (Gray et al., 1993). Spatial processing can also be identified via the use of mental spatial transformations, which occur when a spatial object is transformed from one mental state or location into another mental state or location. Mental spatial transformations—which we refer to simply as spatial transformations—occur in a mental representation that is an analog of physical space and are frequently part of a problem-solving process. There are many types of spatial transformations: creating a mental image, modifying that mental image by adding or deleting features, mental rotation (Shepard & Metzler, 1971), mentally moving an object, animating a static image (Bogacz & Trafton, 2005; Hegarty, 1992), making comparisons between different views (Kosslyn et al., 1999; Trafton et al., 2005), and any other mental operation which transforms a spatial object from one state or location into another.

In general, our research on spatial processing in graph interpretation has focused on complex graphs in complex domains (e.g., fMRI, target motion analysis, meteorology). Our results have been consistent: verbal protocols show that in these domains, at least, people use a great deal of spatial processing to extract and use information from data visualizations (Trafton et al., 2000; Trafton & Trickett, 2001). Further evidence of spatial processing is found in meteorologists' gestures when they talk about how they performed the task (Trafton et al., 2006). Additionally, in keeping with the important role of domain knowledge in graph comprehension (Freedman & Shah, 2002; Roth & Bowen, 2003; Tabachneck-Schijf et al., 1997), experts use far more spatial processing than journeymen (Trafton et al.,

2006). This general result—that spatial processing is prevalent in complex graph comprehension—has been replicated in studies of astronomy and computational fluid dynamics (Trickett & Trafton, in press).

Based on these data, we have become convinced that spatial processing is an important component of a comprehensive model of graph comprehension, i.e., a model that moves beyond the laboratory into the “real world” of practice. Yet, curiously, spatial processing is not *explicitly* included in any of the current models of graph comprehension (e.g., Freedman & Shah, 2002; Lohse, 1993; Peebles & Cheng, 2003; Pinker, 1990). In this paper, we address this criticism of these theories, namely, their lack of an explicit spatial component.

In order to test our hypothesis that spatial processes are involved in many graph interpretation tasks, we used an interference paradigm adapted from (Oh & Kim, 2004) in establishing the role of spatial working memory in visual search. Briefly, Oh and Kim developed both spatial and non-spatial versions of a memory task, and participants performed one of these tasks in tandem with a visual search task. Performance on the spatial working memory task was impaired by the search task, and performance on the search task was impaired by the spatial task. In contrast, the non-spatial memory task did *not* interfere with the visual search task. Oh and Kim concluded, “the visual search process and spatial working memory storage require the same limited-capacity mechanisms.” (See Baddeley 1980 for more on the distinction between spatial and visual processing.)

In our experiment, we substituted a graph interpretation task for Oh and Kim’s visual search task, but the logic of the experimental design remains the same. Dual task performance is almost always relatively worse than single task performance, regardless of the respective natures of the primary and secondary tasks. However, tasks that are similar in nature (i.e., draw on the same processing mechanisms) are likely to interfere more than dissimilar tasks. The key to interpretation, then, is the extent to which one type of secondary task interferes *more* than another. We hypothesized that if the graph task involved spatial processing, it would interfere more with the *spatial* memory task than with the *non-spatial* memory task. If the graph task did not involve spatial processing, there would be no difference in the amount of interference caused by either the spatial or the non-spatial memory task.

This interpretation rests on the assumption that both memory tasks were equally difficult; in order to make sure that this was the case, and in order to establish baseline performance for these two tasks, we also examined performance on both memory tasks carried out alone. Thus our design includes both single and dual task conditions.

Method

Participants

Twenty-six George Mason University undergraduates participated for course credit. Participants were randomly

assigned to one of two experimental conditions (described below): spatial (14 participants) or visual (12 participants).

Stimuli and Tasks

Two tasks were developed, a memory task (primary task) and a graph task (secondary task). We created two versions of the memory task, one spatial and one visual (non-spatial). The memory and graph tasks were presented both alone and in a dual-task combination.

Graph Task Participants were presented with a column graph representing two variables (the graph stimulus is illustrated in Figure 1). One variable was colored grey and the other black. Participants were asked to judge which set of bars, grey or black, was larger overall. They responded by pressing the *g* or *b* keys on the keyboard, respectively. The graphs were constructed such that a) there was always an interaction between the variables, b) there was some constraint on the size difference between the bars, so that the difference could not be ascertained by a purely perceptual strategy, and c) *grey* and *black* were each the correct answer on 50% of the trials.

We deliberately chose a simple task performed on a simple graph, in part to avoid the suggestion that spatial processing in graph interpretation occurs only during highly complex graphical reasoning tasks and/or when the graphs themselves involve a complex spatial layout. In addition, college students are familiar with column graphs, and consequently did not need any special training on how to interpret such graphs; nor were there likely to be large individual differences in their performance on this task.

The graph task could be performed in a number of different ways (Simkin & Hastie, 1987; Trickett & Trafton, 2006). For example, one could mentally stack the grey bars and the black bars, and compare their respective heights, one could mentally find and compare the respective mid-points between the grey and black bars, or one could gauge the difference between the leftmost black and grey bars and compare it with the difference between the rightmost bars. All these operations involve spatial transformations (Trafton et al., 2005) and therefore involve spatial processing.

Spatial Working Memory Task The stimuli for the spatial working memory task (hereafter referred to as the spatial task) consisted of four solid black squares, randomly chosen from among eight possible locations (illustrated in Figure 2). Participants were instructed to remember the location of the four squares until the end of the trial, at which time they were shown one empty black square, again chosen from among the eight possible positions. The location of the empty square matched the location of one of the original four squares in half the cases; in the other half, it was in a different location. Participants had to judge whether the empty square was in a location that had previously been occupied by one of the original squares. They answered by pressing the *y* and *n* keys for *yes* and *no*, respectively. In this task, participants had to remember *location*, a spatial task.

Visual Working Memory Task The stimuli for the visual working memory task (hereafter referred to as the visual task) consisted of four colored squares, chosen at random from among eight easily distinguishable colors (dark green, lime green, yellow, teal, dark blue, red, magenta, and purple). The squares were placed symmetrically around a central fixation point (Figure 3 shows an example). Participants were instructed to remember the color of the four squares until the end of the trial, at which time they were shown one colored square in the center of the screen. The color matched the color of one of the original four squares in half the cases, and in the other half, it was different. Participants had to judge whether the single square was the same color as one of the four original squares. As in the spatial task, they answered by pressing the *y* and *n* keys for *yes* and *no*, respectively. In this task, participants had to remember *color*, a visual task. Note that because the squares were always arranged in the same pattern, and the test square was always in the center of the screen, no spatial memory was required for this task.

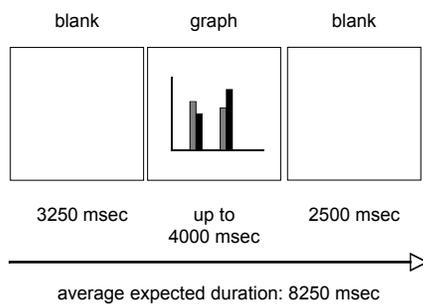


Figure 1: Procedure for graph-only condition

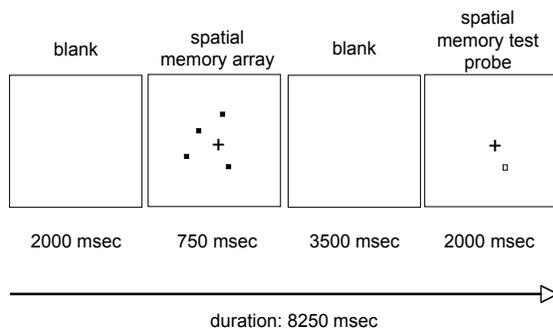


Figure 2: Procedure for memory-only condition (spatial task is illustrated; not drawn to scale).

Experimental Design and Procedure

The design was mixed, with one between-subjects variable (memory task: spatial or visual) and one within-subjects variable (task set: graph-only, memory-only, or dual-task).

Conceivably, participants could “translate” the spatial memory stimulus into a propositional representation and rehearse it. For example, in the stimulus shown in figure 2, they might encode the four squares as an upside-down anchor. Similarly, they might group and rehearse the colors in the visual memory stimulus. To prevent this kind of rehearsal, following Oh and Kim’s methodology (2004), participants were required to continuously recite the first four letters of the alphabet in the memory-only condition.

The procedure for trials in each task-set condition is illustrated in figures 1 to 3. We could not control the amount of time people would take to answer the graph task. However, pilot-testing showed a mean of 2500 msec.; thus, in the memory-only condition, we substituted a blank screen displayed for 2500 msec. for the graph stimulus. Oh and Kim (2004) showed the memory array for 500 msec., but in pilot tests our participants performed the task at chance levels at this duration, so we extended the time to 750 msec. For all tasks, we collected reaction time and accuracy data.

The stimuli were presented using the E-Prime software. Ninety-six unique pairs of memory array/memory test probe combinations were used, randomly selected for each participant. Forty-eight unique graphs were used, with each graph presented twice, using E-Prime’s random selection procedure. Each condition (graph-only, memory-only, and dual-task) was tested in blocks, randomly ordered across participants. A block comprised 10 training trials and 96 experimental trials.

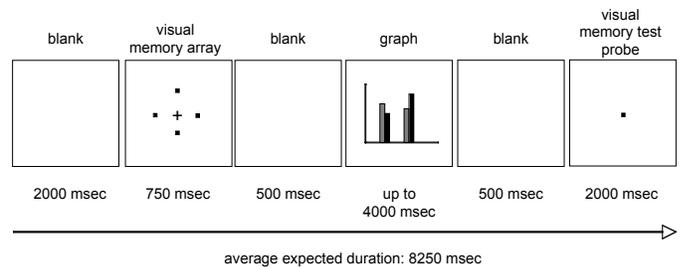


Figure 3: Procedure for dual-task condition (visual task is illustrated; not drawn to scale). Stimuli were in color.

Results and Discussion

Our main goal was to test the hypothesis that the graph task involved spatial processing, and that therefore it would interfere more with the spatial task than with the visual task. In order to test this hypothesis, we conducted a repeated measures ANOVA with one between-subjects variable (memory task: spatial vs. visual) and one within-subjects variable (task set for the memory task: memory-only vs. dual-task). We anticipated a main effect of task set, in that accuracy on the memory task would be significantly worse in the dual task condition than the memory-only condition. Critically, for our hypothesis, we also predicted a significant interaction, such that in the dual-task condition, accuracy on the spatial memory task would suffer a *larger decrement* than performance on the visual memory task.

As predicted, the analysis of variance showed a significant main effect of task set, $F(1,24) = 45.193, p < .01$. This result is not surprising; in general, performance worsens when a task is performed in conjunction with a second task. There was no main effect of memory task, $F(1,24) = 1.877, p = .18$, so that *overall* there was no difference in performance on the spatial or visual tasks. Importantly, there was a significant interaction between task set and memory task, $F(1,24) = 5.76, p < .05$. Figure 6 illustrates these results.

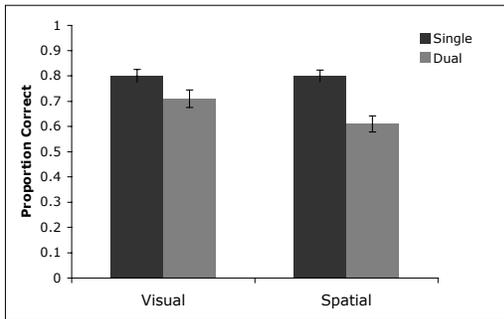


Figure 6: Accuracy for both visual and spatial memory tasks, in single- and dual-task conditions

There are several important points to be made about Figure 6. First, it shows that participants performed equally well on both memory tasks in the single task condition, indicating that the tasks were of equal difficulty. Thus the increased decrement in performance on the spatial task in the dual-task condition cannot be attributed to the spatial task's being more difficult than the visual task. The fact that the graph task interfered significantly more with the spatial memory task than the visual memory task provides support for our hypothesis that the graph task involved spatial processing, and that it was the spatial nature of the graph task that caused the additional interference. In addition, Figure 6 shows that participants' accuracy on the spatial memory task was a full ten percent worse than on the visual task; thus, the difference is not only statistically significant but represents a meaningful drop in performance.

It is possible that in order to offset the effects of the interference, participants engaged in some kind of speed-accuracy tradeoff, such that in the dual-task condition, although they performed more *accurately* in the visual memory task, they performed more *slowly*. If this were the case, the interference (evidenced by reaction time) would actually be *greater* in the visual memory task than in the spatial memory task. In order to test this possibility, we examined the reaction time data. We performed a repeated measures ANOVA similar to our analysis of the accuracy data, with one between-subjects variable (memory task: spatial vs. visual) and one within-subject variable (task set for the memory task: memory-only vs. dual-task). As Figure 7 shows, there was no main effect of either memory task or task set, and no interaction (all F -values < 1). Participants responded equally fast to the memory task, regardless of the interference caused by the secondary graph task.

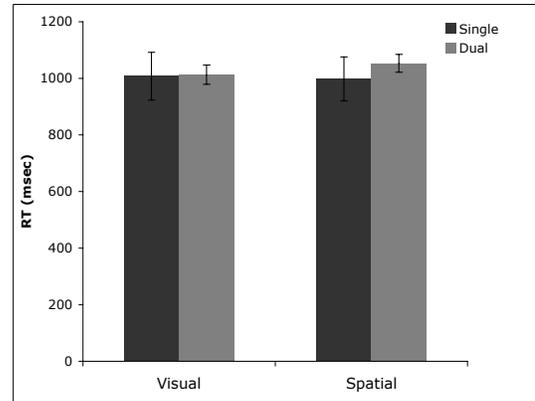


Figure 7: Reaction time for the memory tasks

Finally, we examined the graph data, to determine whether either of the memory tasks interfered with the graph task. Figures 8 and 9 show, respectively, the accuracy and reaction time data for the graph task. As Figure 8 shows, participants were remarkably consistent in their accuracy on the graph task, answering correctly approximately 75% of the time, regardless of memory condition (spatial vs. visual) and task set (single vs. dual). Note that although the graph task was not complex, neither was it trivial. Although participants performed well above chance, they were nowhere near ceiling performance. A repeated measures ANOVA, conducted as for the memory task data, showed no significant effects nor any interaction (all F -values < 1).

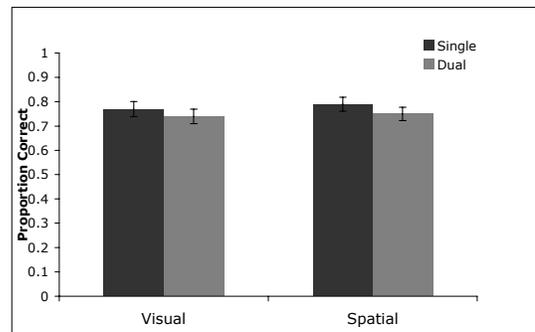


Figure 8: Accuracy for the graph task

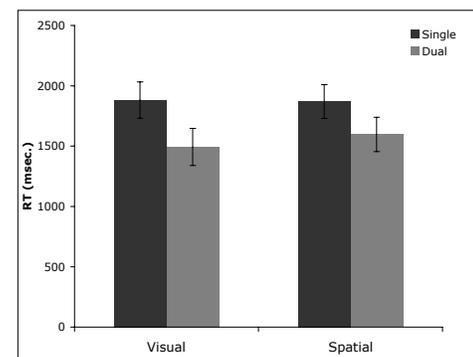


Figure 9: Reaction time for the graph task

The reaction time data showed a similar pattern, with the exception that there was a significant effect of task set, $F(1,24) = 10.78$, $p < .01$. Interestingly, participants were slightly *faster* to answer the graph question in both the dual task conditions; critically, however, there was not a significant interaction between memory task condition (spatial vs. visual) and task set (single vs. dual), $F(1,24) = 1.05$, $p = .32$. Note also that all participants were able to answer the graph question well within the 4000 msec. for which the stimulus was displayed.

In summary, accuracy on the graph task was not affected by the additional load of the memory task; nor was it affected by the task type in the dual-task condition. Reaction time on the graph task was not affected by the type of task in the dual-task condition, but was slightly faster when participants had to perform the memory task as well. It appears that in both spatial and visual memory conditions, participants were trying to hasten their performance on the graph task, in order to prevent further decay of the memory stimulus. Critically, the speed-up was approximately the same in both memory conditions.

General Discussion and Conclusion

We conducted an experiment using an established interference paradigm to test our hypothesis that graph interpretation involves spatial processing. Our results showed that the graph task interfered with performance on a spatial memory task, but not a visual (non-spatial) memory task. This finding suggests that the graph task and the spatial memory task tapped the same underlying cognitive processes, whereas the graph task and the visual task drew on different processes. Given that the spatial memory task has been established as tapping spatial working memory, we conclude that the graph task involved spatial processing.

Our experimental design does not completely rule out the possibility that some other aspect of the graph task (apart from spatial cognition) caused the different levels of interference in the memory tasks. If so, one would expect the same pattern of results using a different, non-spatial graph task. We are currently running an experiment to rule out this explanation, using a graph task that involves purely reading off values, without the use of spatial cognition. We anticipate that the interference caused by this secondary task will be the same in both the spatial and visual primary tasks.

For the experiment described above, we chose a simple graph task, not because such a task is fully representative of “real-world” graph use, but because it is *sufficiently* representative to serve as a proof of concept. This task is of similar complexity to many tasks used in most prior studies of graphing that have been used to construct current theories of graph comprehension—like them, it is context-lean (or context-free) and demands no domain knowledge and only minimal graph skill. Yet even this simple task-graph combination appears to involve spatial processing.

In a recent review of the graph comprehension literature, (Shah et al., 2005) draw a distinction between perceptual and conceptual processes. In their interpretation, perceptual

processes are “bottom-up encoding mechanisms,” which focus on the visual features of the display. Conceptual processes equate to “top-down encoding processes,” which influence interpretation. They propose that when these perceptual processes are not sufficient, “information must be retrieved by complex inferential processes.” Although several models agree that these “complex inferential processes” are an essential part of the graph comprehension process under some circumstances, they remain largely unspecified. We propose that spatial processing comprises a substantial part of these complex inferential processes

It appears that it is not the *complexity* of the graph or task *per se* that determines whether spatial processing is used, but rather whether the information can be directly extracted from the display (e.g., reading off specific, marked values) or inferred using direct perceptual processes (e.g., determining whether a line slopes up or down). In accordance with current models, we propose that in such cases, the task can be accomplished without spatial processing. However, we suggest that when these conditions are not met—for example, when users must integrate disparate visual chunks, make inferences, or otherwise go beyond the data—spatial processing will likely be used. We propose that, although this may be more likely in complex graph/task combinations, it can nonetheless be the case regardless of the complexity of the graph.

Quite possibly, the reason spatial processing has not been part of graph comprehension models is that the focus on simple tasks and graphs has made it unnecessary, since in general, simple perceptual processes are sufficient to account for performance in these circumstances. In addition, the strength of graphs as a form of representation is that they can make implicit things explicit (at least, good graphs do), so graphs are designed and selected so that it *is* possible to make direct comparisons between visual chunks. However, as tasks, domains, and visualizations become more complex, this transparency may not always be possible or even desirable (meteorologists, for example, don’t want simpler graphs; they want many variables represented). As graph comprehension research moves out of the laboratory into the “real world” of practice, it will be more important for graph comprehension models to address this reality.

Without incorporating spatial processing, it appears that current models of graph comprehension will be incomplete. Including spatial processing in these models will help us to understand why some representations might be better than others at a cognitive level, by shedding light on processes that underlie different graph/task interactions. It can help identify situations in which spatial processing is unavoidable and can help us make predictions about performance using these graphs. In some situations, it can help us design better graphs, by developing creative ways to reduce the amount of spatial processing required (for example, by facilitating direct comparisons that can be performed perceptually). In sum, we propose that including spatial processing is an important step in building a comprehensive model of graph comprehension.

Acknowledgments

Work on this project was supported by ONR grants M12439 and N0001403WX30001 to Greg Trafton. Thanks to David Cades, Melanie Diez, David Kidd, Sarah Kriz, Patrick McKnight, and Raj Ratwani for their helpful comments, and to Kim Hilton for assistance with data collection.

References

- Baddeley, A. D., & Lieberman, K. (1980). Spatial working memory. In R. S. Nickerson (Ed.), *Attention and performance viii* (pp. 521-539). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Bogacz, S., & Trafton, J. G. (2005). Using dynamic and static displays: Using images to reason dynamically. *Cognitive Systems Research*, 6(4), 312-319.
- Cleveland, W. S. (1985). *The elements of graphing data*. Monterey, CA: Wadsworth.
- Freedman, E. G., & Shah, P. (2002). *Toward a model of knowledge-based graph comprehension*. Paper presented at the Diagrammatic Representation and Inference, Second International Conference, Diagrams 2002, Callaway Gardens, GA, USA.
- Gray, W. D., John, B. E., & Atwood, M. E. (1993). Project Ernestine: Validating a gOMS analysis for predicting and explaining real-world performance. *Human-Computer Interaction*, 8(3), 237-309.
- Hart, R. A., & Moore, G. T. (1973). The development of spatial cognition: A review. In R. M. Downs & D. Stea (Eds.), *Image and environment: Cognitive mapping and spatial behavior* (pp. 246-295). Chicago: Aldine Publishing Company.
- Hegarty, M. (1992). Mental animation: Inferring motion from static displays of mechanical systems. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 18(5), 1084-1102.
- Kosslyn, S. M., Sukel, K. E., & Bly, B. M. (1999). Squinting with the mind's eye: Effects of stimulus resolution on imaginal and perceptual comparisons. *Memory and Cognition*, 27(2), 276-287.
- Lohse, G. L. (1993). A cognitive model for understanding graphical perception. *Human Computer Interaction*, 8, 353-388.
- NCTM. (2000). *Principles and Standards for School Mathematics*. Reston, VA: National Council of Teachers of Mathematics.
- Oh, S., & Kim, M. (2004). The role of spatial working memory in visual search efficiency. *Psychonomic Bulletin & Review*, 11(2), 275-281.
- Peebles, D., & Cheng, P. C.-H. (2003). Modeling the effect of task and graphical representation on response latency in a graph reading task. *Human Factors*, 45(1), 28-46.
- Pinker, S. (1990). A theory of graph comprehension. In R. Freedle (Ed.), *Artificial intelligence and the future of testing* (pp. 73-126). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Roth, W.-M., & Bowen, G. M. (2003). When are graphs worth ten thousand words? An expert-expert study. *Journal of the Learning Sciences*, 21(4), 429-473.
- Schunn, C. D., Cho, K., & Raghavan, K. (in press). Precursors for improved learning in a model-based middle-school science curriculum. *Applied Developmental Psychology*.
- Shah, P., Freedman, E. G., & Vekiri, I. (2005). The comprehension of quantitative information in graphical displays. In P. Shah & A. Miyake (Eds.), *Cambridge handbook of visuospatial thinking* (pp. 426-476). New York, NY: Cambridge University Press.
- Shepard, R., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, 171, 701-703.
- Simkin, D., & Hastie, R. (1987). An information-processing analysis of graph perception. *Journal of the American Statistical Association*, 82(398), 454-465.
- Tabachneck-Schijf, H. J. M., Leonardo, A. M., & Simon, H. A. (1997). Camera: A computational model of multiple representations. *Cognitive Science*, 21(3), 305-350.
- Tan, J. K. H., & Benbasat, I. (1990). Processing of graphical information: A decomposition taxonomy to match data extraction tasks and graphical representations. *Information Systems Research*, 1(4), 416-439.
- Trafton, J. G., Kirschenbaum, S. S., Tsui, T. L., Miyamoto, R. T., Ballas, J. A., & Raymond, P. D. (2000). Turning pictures into numbers: Extracting and generating information from complex visualizations. *International Journal of Human Computer Studies*, 53(5), 827-850.
- Trafton, J. G., & Trickett, S. B. (2001). A new model of graph and visualization usage. In J. D. Moore & K. Stenning (Eds.), *The proceedings of the twenty third annual conference of the cognitive science society* (pp. 1048-1053). Mahwah, NJ: Erlbaum.
- Trafton, J. G., Trickett, S. B., & Mintz, F. E. (2005). Connecting internal and external representations: Spatial transformations of scientific visualizations. *Foundations of Science*, 10, 89-106.
- Trafton, J. G., Trickett, S. B., Stitzlein, C. A., Saner, L., Schunn, C. D., & Kirschenbaum, S. S. (2006). The relationship between spatial transformations and iconic gestures. *Spatial cognition and computation*, 6(1), 1-29.
- Trickett, S. B., & Trafton, J. G. (2006). Toward a comprehensive model of graph comprehension: Making the case for spatial cognition.
- Trickett, S. B., & Trafton, J. G. (in press). "What if...": The use of conceptual simulations in scientific reasoning. *Cognitive Science*.
- Trickett, S. B., Trafton, J. G., Saner, L., & Schunn, C. D. (In press). "I don't know what's going on there": The use of spatial transformations to deal with and resolve uncertainty in complex visualizations.
- Tufte, E. (2001). *The visual display of quantitative information*. (2nd ed.). Cheshire, CT: Graphics Press.
- Zacks, J., & Tversky, B. (1997). Bars and lines: A study of graphic communication. Paper presented at the 2nd AAAI workshop on Reasoning and Diagrams, Cambridge, MA.