

Extracting Explicit and Implicit Information from Complex Visualizations

J. Gregory Trafton¹, Sandra Marshall², Farilee Mintz³, and Susan B. Trickett⁴

¹ Naval Research Laboratory, Washington, DC 20375-5337, USA

`trafton@itd.nrl.navy.mil`

² San Diego State University

`smarshall@sciences.sdsu.edu`

³ ITT Industries

`mintz@aic.nrl.navy.mil`

⁴ George Mason University

`strickett@gmu.edu`

Abstract. How do experienced users extract information from a complex visualization? We examine this question by presenting experienced weather forecasters with visualizations that did not show the needed information explicitly and examining their eye movements. We replicated Carpenter & Shah (1998) when the information was explicitly available on the visualization. However, when the information was not explicitly available, we found that forecasters used spatial reasoning in the form of spatial transformations. We also found a strong imagerial component for constructing meteorological information.

1 Introduction

How do people comprehend and extract information from a complex graph or visualization? Current models of graph comprehension account very well for simple tasks (e.g., “What is the price of graphium in 1982?”). Similarly, most empirical studies have used simple graphs (e.g., few variables and few data points). For example, within the graph comprehension literature, Lohse (1993) studied some of the most complex graphs, and his standard graphs used 6 variables with a total of 72 data points; Carpenter & Shah (1998) provided participants with graphs that had 2 to 4 variables with 4 to 8 data points (containing interactions); Tan & Benbasat (1990) had participants study graphs that displayed 3 variables with 18 total data points.

In contrast, many domains attempt to display tens of variables and tens of thousands of data points (or more). For example, many scientists use extremely complex visualizations with thousands of data points (Trafton, Trickett, & Mintz, in press; Trickett, Fu, Schunn, & Trafton, 2000; Trickett, Trafton, Schunn, & Harrison, 2001) and weather forecasters routinely examine visualizations with 10 or more variables and thousands of data points over multiple spatial and temporal scales (Trafton, Kirschenbaum, Tsui, Miyamoto, Ballas, & Raymond, 2000). Current wisdom suggests that showing a graph well depends

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2002		2. REPORT TYPE		3. DATES COVERED 00-00-2002 to 00-00-2002	
4. TITLE AND SUBTITLE Extracting Explicit and Implicit Information from Complex Visualizations				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory, Navy Center for Applied Research in Artificial Intelligence (NCARAI), 4555 Overlook Avenue SW, Washington, DC, 20375				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT How do experienced users extract information from a complex visualization? We examine this question by presenting experienced weather forecasters with visualizations that did not show the needed information explicitly and examining their eye movements. We replicated Carpenter & Shah (1998) when the information was explicitly available on the visualization. However, when the information was not explicitly available, we found that forecasters used spatial reasoning in the form of spatial transformations. We also found a strong imagerial component for constructing meteorological information.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 15	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

on making the variables and patterns explicit (Gillan, Wickens, Hollands, & Carswell, 1998; Tufte, 1983, 1990, 1997). However, when there are an extremely large number of data points, this explicitness rule may need to be relaxed in order to avoid visual clutter.

We believe that as a direct result of the complexity of the graphs, some data will need to be represented more imprecisely (e.g., pressure on Figure 1) and this imprecision may prevent a user from directly reading off information from the graph. The user may need to use a different set of mental procedures than is needed in simpler graphs.

This paper first reviews current models of graph comprehension, suggests augmenting them in a specific manner when the visualization is complex and the task requires information that is implicitly represented, and presents an initial study to explore these issues.

1.1 Current models of graph comprehension

The most influential research on graph and visualization comprehension is Bertin's (1983) task analysis which suggests three main processes in graph and visualization comprehension:

1. Encode visual elements of the display: For example, identifying lines and axes. This stage is influenced by pre-attentive processes and is affected by the discriminability of shapes.
2. Translate the elements into patterns: For example, notice that one bar is taller than another or the slope of a line. This stage is affected by distortions of perception and limitations of working memory.
3. Map the patterns to the labels to interpret the specific relationships communicated by the graph. For example, determine the value of a bar graph.

Most of the work done on graph comprehension has examined the encoding, perception, and representation of graphs. Cleveland and McGill, for example, have examined the psychophysical aspects of graphical perception (Cleveland & McGill, 1984, 1986). Similarly, Pinker's theory of graph comprehension, while quite broad, focuses on the encoding and understanding of graphs (Pinker, 1990). Kosslyn's work emphasizes the cognitive processes that make a graph more or less difficult to read. Kosslyn's syntactic and semantic (and to a lesser degree pragmatic) level of analysis focuses on encoding, perception, and representation of graphs (Kosslyn, 1989). Tracking users' eye movements, Carpenter & Shah (1998) have shown that people switch between looking at the graph and the axes in order to comprehend the visualization. Similarly, Peebles & Cheng (2001a, 2001b) have suggested that in experimental trials, people cycle between looking at the graph and the question they are trying to answer.

This scheme seems to work very well when the graph explicitly represents the needed information. Thus, when a meteorologist is asked to provide the wind direction at Pittsburgh on Figure 1, the meteorologist searches for Pittsburgh, finds the wind barb over Pittsburgh, and determines the direction it points (280, or westerly, in this case). Note that it is slightly difficult to see the part of the

barb showing wind speed in this figure, which should not affect the perception of wind direction.




Fig. 1. A typical visualization that meteorologists use. Wind speed and wind direction is shown by the wind barbs; temperature is color coded to the legend on the far right, and pressure is shown by lines connecting the same pressure (1008, 1012, etc.). The original is in color.

1.2 When information is not explicitly available

What happens when information is not explicitly available on a graph? For example, in Figure 1, how would the pressure at Pittsburgh be determined? ¹

Current theories either do not deal with how information is extracted when the information is not explicitly available (e.g., Lohse, 1993) or leave the whole process unspecified (e.g., Pinker, 1990; Kosslyn, 1989).

Pinker's (1990) framework, for example, provides the most detailed description of how information is extracted from a graph. He claims that once the graph

¹ If a pressure line goes directly through the location for which a forecaster wants pressure, the forecaster simply reads it off (e.g., St. Louis' pressure is 1012). If, however, a pressure line does not pass directly through the desired location, the forecaster must interpolate between two lines (e.g., Pittsburgh's pressure is 1011 because it is between 1008 and 1012 and closer to 1012 than the middle).

has been encoded, all inferences and reasoning occur propositionally. Pinker’s “conceptual questions” and “conceptual messages” are his primary method of inferencing and reasoning. Thus, Pinker’s theory suggests that people initially encode the aspects of a graph into a propositional representation, and then reason with that propositional representation to answer conceptual questions.²

Trafton & Trickett (2001), in contrast, suggested that a great deal of complex visualization reasoning occurs spatially, especially when the information is not explicitly available on the visualization. Trafton & Trickett (2001) showed that when scientists’ own complex visualizations did not explicitly show the information they needed, scientists created complex mental representations and manipulated them to help them answer the specific questions they had. Trafton & Trickett suggested that the scientists constructed spatial representations.

Trafton & Trickett (2001) suggested that when people who used complex visualizations needed to extract information that was not available, they used spatial information to create internal mental representations to reason with. We have developed a framework for coding and working with these kinds of graphs and visualizations called *Spatial Transformations* that will be used to investigate these issues. We will argue that spatial transformations are a fundamental aspect of complex visualization usage.

Spatial Transformations are cognitive operations that a scientist performs on a visualization. Sample spatial transformations are mental rotation (e.g., Shepard & Metzler, 1971), creating a mental image, modifying that mental image by adding or deleting features to or from it, animating an aspect of a visualization (Hegarty, 1992) time series progression prediction, mentally moving an object, mentally transforming a 2D view into a 3D view (or vice versa), comparisons between different views (Kosslyn, Sukel, & Bly, 1999; Trafton et al., in press), and anything else a scientist mentally does to a visualization in order to understand it or facilitate problem solving. Also note that a spatial transformation can be done on either an internal (i.e., mental) representation or an external image (i.e., a scientific visualization on a computer-generated image). A more complete description of spatial transformations (along with a mini-experiment to teach interface designers when to use a 2D or a 3D representation) can be found at <http://e-lab.hfac.gmu.edu/~trafton/405st.html>

Trafton & Trickett (2001) focused on how scientists created mental representations to answer specific questions. Because that study examined scientists working *in vivo* (Baker & Dunbar, 2000; Dunbar, 1995, 1997; Trickett et al., 2001), it was not possible to control what they saw or what information was needed. In this study, we asked experienced meteorologists to provide specific information from weather visualizations like that shown in Figure 1. We also tracked where they were looking with an eyetracker system.

² There are, of course, a large number of theories that discuss diagrammatic reasoning and suggest that reasoning occurs spatially or via images (Hegarty & Sims, 1994; Larkin & Simon, 1987; Narayanan & Hegarty, 1998; Tabachneck-Schijf, Leonardo, & Simon, 1997). However, most of these theories are concerned with diagram understanding, not graph comprehension.

We had several goals in this study. First, we wanted to provide a baseline measure by showing that when experienced users extracted information explicitly available from a complex visualization, they would extract information based on the canonical graph comprehension model presented earlier; this would also be a partial replication of other graph comprehension studies (Carpenter & Shah, 1998). For a standard meteorological visualization, we predicted that experienced meteorologists would be able to go directly to the desired information and extract it. Second, we wanted to examine how forecasters extract information from the graph that was not explicitly available (i.e., do they use a propositional representation or do they reason with the graph itself). We complete our discussion with some anecdotal evidence of how forecasters remember information that they had seen previously.

2 Method

In order to investigate the issues discussed above, we examined forecasters as they were examining meteorological visualizations.

2.1 Task

Forecasters were presented with a weather visualization (see Figure 1). They were then asked several questions in the following order during the *Graph Comprehension* portion of this study.

- What is the synoptic weather?
- What is the wind speed and wind direction at Pittsburgh, Pennsylvania?
- What is the temperature at Pittsburgh, Pennsylvania?
- What is the pressure at Pittsburgh, Pennsylvania?

Forecasters were then shown a second weather visualization (see Figure 2) and asked two questions:

- What is the wind speed and wind direction at Honolulu?
- What is the relative humidity at Honolulu?

Finally, during the *recall* portion of the study, forecasters were shown a blank screen and asked about wind speed, wind direction, temperature, and pressure at four locations from the first visualization: Pittsburgh, Pennsylvania (Exact), Philadelphia, Pennsylvania (Near), Atlanta, Georgia (Medium), and Houston, Texas (Far). The exact question probed participants' memory for the information they had recently read off. The near, medium, and far conditions allowed us to see how their accuracy changed as the location moved farther away from the area they focused on (Pittsburgh); Philadelphia is very close to Pittsburgh, Georgia is further away, and Houston is quite far from Georgia on the visualization.

It should be noted that all of these question types were based on the types of tasks meteorologists typically do (Lowe, 1999; Trafton et al., 2000).

These questions came in several forms:




Fig. 2. A typical visualization that meteorologists use. The wind barbs represent wind speed and direction; the black lines represent geopotential height, and the dial in the upper right hand corner shows a legend for relative humidity. The original is in color.

- Questions asking for qualitative information, where the users have to integrate across information. The synoptic weather question is an example of this.
- Questions asking for quantitative information, where the answers (numbers) are explicitly represented (as in traditional graph comprehension studies). The question about temperature and the relative humidity question fall into this category because users only have to map a color to the legend and the numbers are printed on the legend. Also, all locations on the map have an instance of the variable. We expect participants to read-off the information as they need it, just as the graph comprehension literature predicts.
- Questions asking for quantitative information that is imprecisely represented in the graph. In other words, the numbers themselves are not represented explicitly, but there is a symbology associated with them that the user must know in order to extract the needed information. Wind direction and wind speed are part of this question type. If the user has the knowledge about how to interpret the symbology, the graph comprehension theories would predict a read-off strategy.
- Questions asking for quantitative information that is entirely implicit in the graph. Numbers are not represented either explicitly or imprecisely to answer that question, but must be inferred. The question about pressure is an instance of this category. Although pressure is represented explicitly

for some locations, it must be inferred for others (like Pittsburgh). Graph comprehension theories do not make good predictions here, but the spatial transformation theory does suggest a framework for which this information could be extracted.

- Questions asking for quantitative information when the information is no longer available externally, but must be retrieved from memory. All the recall questions fall into this category.

2.2 Participants

All forecasters were Naval or Marine Corps forecasters and forecasters-in-training. All had completed at least the first level weather school. They ranged in forecasting experience from 1 to 10 years. All forecasters had significant operational experience.

Four novices and two expert forecasters performed the tasks. The experts had an average of 10 years forecasting experience and the novices had an average of 2.5 years forecasting experience.

2.3 Setting and Apparatus

Meteorological visualizations were collected from <http://www.fnmoc.navy.mil/>. This web site is used a great deal by current Navy forecasters; all forecasters in the study were very familiar with this site and the visualizations that were used. This web site was used extensively by forecasters in a previous study of Navy meteorologists (Trafton et al., 2000)

The experimental sessions took place in a room equipped with a PC and an EyeLink System from SMI. The eyetracker had headmounted optics with three small cameras (left eye, right eye, head compensation). It records eye movements and pupil dilations at 250 Hz (every 4 milliseconds).

2.4 Procedure

Each session began with hooking up the eyetracker and calibrating it to the individual participant. Next, the participant was shown the initial screen. All questions were asked verbally to prevent additional eyetracking to the question (Peebles & Cheng, 2001a, 2001b). After the first set of questions was finished, the next visualization was shown, and the next set of questions was asked. When the participant finished that set, a blank screen was shown for the test questions.

3 Results

How do skilled users extract information from complex visualizations? Is there a difference between extractions that occur when the information is explicitly available on the visualization and when the information is not explicitly available? This results section attempts to answer these questions.

3.1 Overview

All forecasters were able to extract the information from the graphs; they all knew how to read these kinds of meteorological visualizations.

The graph comprehension aspect of this task was not difficult for the participants; experts and novices were very accurate when reading off information from the graph. There were no consistent qualitative or quantitative differences between experts or novices during graph comprehension, so all participants' results were combined.

3.2 Extracting information that is explicitly available

We first examined how forecasters extracted information that is explicitly on a complex visualization. For several of the questions during graph comprehension, the information was explicitly available on the graph. Thus, for wind speed and wind direction, participants simply had to find the location (e.g., Pittsburgh) and read off the wind speed and wind direction. For other variables (temperature and relative humidity) the forecasters had to match a color code with a legend.

Bertin's model presented earlier suggested that after people encoded a graphical element, they would translate the elements into patterns and map the patterns into labels. Bertin (and others) implied that people would perform this operation serially. Carpenter & Shah (1998) however, found that participants interpreting graphs would cycle back and forth between the graph area and the legend, suggesting a much more iterative aspect to graph comprehension than previously known. We explicitly examined if experienced (both expert and novice) forecasters would cycle back and forth between the graph area and the legend, as Carpenter & Shah (1998) found.

During the graph comprehension stage, we asked two questions whose variables had legends: "What is the temperature at Pittsburgh?" and "What is the relative humidity at Honolulu?" We examined the number of times that participants cycled back and forth between the graph area and the legend. Figure 3 shows a typical subject's eye movements as he finds the answer to the question "What is the temperature at Pittsburgh, Pennsylvania?"

Consistent with Carpenter & Shah (1998), we found frequent switching for both questions. Table 1 shows the average number of cycles for each question.

Question	# of glances between graph and legend (avg)
What is the temperature at Pittsburgh?	6.0
What is the relative humidity at Honolulu?	4.5

Table 1. The average number of times that the forecasters looked between the legend and the graph area.




Fig. 3. Eyetracking of a participant after asked the question, “What is the temperature at Pittsburgh, Pennsylvania?” “Dots” are 4 milliseconds apart. There are four glances between the legend and the graph area in this figure.

3.3 Extracting information that is not explicitly available

How do experienced forecasters extract information that is not explicitly available on a visualization? There are some propositional theories (e.g., Lohse, 1993; Peebles & Cheng, 2001a; Pinker, 1990) that do not explain the extraction process well if the information is not explicitly available on the graph. At a minimum, these propositional theories leave this process unspecified. Other theories (like Trafton et al.’s (2001) spatial transformation theory) and many theories on diagram comprehension suggest that people use both propositional and spatial reasoning (e.g., Tabachneck et al.’s (1997) CaMeRa theory).

In order to examine the cognitive process of how experienced graph users extract information that is not explicitly available, we asked experienced forecasters what the pressure was at Pittsburgh. Recall that pressure is not represented explicitly for Pittsburgh; forecasters must interpolate between the two pressure lines to determine the actual pressure of Pittsburgh. Eyetracking results show that all participants went through several steps to extract this information:³

1. Trace the upper pressure line;
2. Find the pressure itself (the number);

³ Note that this was the fourth question asked (the third about Pittsburgh), so participants already knew where Pittsburgh was.

3. Trace the lower pressure line;
4. Find the pressure itself (the number);
5. Bridging (glancing) back and forth between the two pressure lines.

Table 2 shows the average number of times the forecasters performed each tracing task.

Tracing type	# of Traces (avg)
Tracing the upper pressure line	2.4
Tracing the lower pressure line	2.6
Bridging between the two pressure lines	7.4

Table 2. The average number of times that the forecasters traced each pressure line and bridging between the two lines.

Figure 4 shows one of the forecasters' eye movements as he was bridging between the two lines. We interpret these eye tracking movements to suggest that the forecasters were using the diagram to help them calculate the distance between the two lines. We believe that after the forecasters located the pressure lines, they were mentally drawing a line in between the two lines and dividing up the space between Pittsburgh and the pressure lines. This kind of spatial transformation seems to be more effortful than simply line tracing, as shown by the increased number of eye tracking "bridges" between the two pressure lines as compared to the pressure line tracing, $\chi^2(2) = 19.4, p < .001$, Bonferroni adjusted χ^2 s significant at $p < .05$.

3.4 Recalling and constructing information

After the graph comprehension part of the task, forecasters were asked for information that they had just extracted (Pittsburgh) or for information that they had seen but not explicitly extracted. This aspect of the task was seen as quite difficult. Several participants were not able to answer some of the recall questions, especially in the far condition. Experts attempted 88% of the questions, while novices attempted only 56% of them, $\chi^2(1) = 7.1, p < .01$. Figure 5 shows the accuracy of those participants who did complete the questions for pressure (wind speed, wind direction, and temperature all showed a similar pattern). While these results must obviously be interpreted with care because of the low completion rate, it is evident that the experts were far more accurate than the novices, especially at the non-recall questions.

A very small number of forecasters showed an interesting pattern while answering questions about temperature (we are missing some eye movement data because several forecasters closed their eyes or looked off the screen). When asked about the temperature for Pittsburgh (an exact question that was being recalled), most forecasters focused on a specific area of the screen. However, when asked about temperature of other areas (near, medium, and far), 67% of




Fig. 4. A forecaster's eye movements as he was bridging between the two pressure lines. This is only a partial display of his eye movements while answering this question (i.e., the tracing the two pressure lines has been omitted).

the forecasters “looked” to the right to “examine” the legend, even though there was nothing on the blank screen to see. We interpreted this kind of glance as evidence for visual imagery.⁴

Figure 6 shows one of the forecasters' eye movements as he made this kind of glance. We have not run enough participants through this type of experiment or even seen enough evidence of this kind of glance, but we know of no other eye movement study that has demonstrated any kind of “examination” of a mental image by eye movements.

4 General Discussion

In this paper, we explored the process that experienced users went through when they extracted information from a complex visualization when the information was explicitly available and when the information was not explicitly available.

⁴ It is interesting that, while looking at the legend was a frequent glance, there were other glances that did not seem to happen. For example, the participants only rarely “looked” at a location that was geographically correct, though there was some change of glance location when questions about a different location were asked. This suggests that the entire representation is not encoded exclusively as either an image or as a proposition.




Fig. 5. Accuracy of the forecasters that completed the question for the pressure question across all locations.

We replicated the results from Carpenter & Shah (1998) and showed that even experienced users scan back and forth between a legend and its value.

We also presented eye movement data that showed that when users need to interpolate between two lines (like pressure lines or isobars), they do not use a propositional representation; rather, they use some spatial or imagerial process to trace a line that allows them to extract the needed information. This extraction literally occurs as a combination between the internal (visualization) and external representations (spatial transformations). This finding also provides further support for the spatial transformation framework (Trafton & Trickett, 2001; Trafton et al., in press).

This study also provided some very preliminary data on differences between experts and novices. We found that both expert and novice forecasters were able to accurately extract information from complex meteorological visualizations. We also found that experts did a better job during the test phase for non-recall items than novices. This finding is consistent with the finding of Trafton et al. (2000) who suggested that experienced forecasters build a qualitative mental model (QMM) and reason with that (rather than simply using climatological knowledge or memory). This finding is consistent with Trafton et al.'s (2000)




Fig. 6. Eyetracking of a participant after asked the question, “What is the temperature at Philadelphia, Pennsylvania?” This was the “Near” question during the recall stage. The screen itself was blank; it showed nothing but white, so there was literally nothing for the forecaster to look at.

view that expert forecasters are able to generate complex quantitative relationships by extracting primarily qualitative information from a complex weather visualization.

We also suggested that some forecasters use visual imagery to generate information they can not recall. Other researchers have suggested that visual or spatial imagery is used when reasoning with complex diagrams (e.g., Hegarty & Sims, 1994; Tabachneck-Schijf et al., 1997), but most theories of graph comprehension do not. One of our main goals in this paper was to provide some evidence for combining some of the visual and spatial research from diagrammatic reasoning into the graphical reasoning research.

Acknowledgments: This research was supported in part by grants N00014-00-WX-20844 and N00014-00-WX-40002 to the first author from the Office of Naval Research and by grant number N00014-00-1-0331 from the Office of Naval Research to the second author. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U. S. Navy.

Bibliography

- Baker, L. M., & Dunbar, K. (2000). Experimental design heuristics for scientific discovery: the use of "baseline" and "known standard" controls. *International Journal of Human Computer Studies*, 53(3), 335–349.
- Bertin, J. (1983). *Semiology of graphs*. Madison, WI: University of Wisconsin Press.
- Carpenter, P. A., & Shah, P. (1998). A model of the perceptual and conceptual processes in graph comprehension. *Journal of Experimental Psychology: Applied*, 4(2), 75–100.
- Cleveland, W. S., & McGill, R. (1984). Graphical perception: theory, experimentation, and application to the development of graphical method. *Journal of the American Statistical Association*, 79, 531–553.
- Cleveland, W. S., & McGill, R. (1986). An experiment in graphical perception. *International Journal of Man-Machine Studies*, 25, 491–500.
- Dunbar, K. (1995). How scientists really reason: Scientific reasoning in real-world laboratories. In R. J. Sternberg, & J. E. Davidson (Eds.), *The nature of insight*, (pp. 365–395). Cambridge, MA: MIT Press.
- Dunbar, K. (1997). How scientists think: Online creativity and conceptual change in science. In T. B. Ward, S. M. Smith, & S. Vaid (Eds.), *Creative thought: An Investigation of Conceptual Structures and Processes*, (pp. 461–493). Washington, DC: APA Press.
- Gillan, D. J., Wickens, C. D., Hollands, J. G., & Carswell, C. M. (1998). Guidelines for presenting quantitative data in HFES publications. *Human Factors*, 40(1), 28–41.
- 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.
- Hegarty, M., & Sims, V. K. (1994). Individual differences in mental animation during mechanical reasoning. *Memory and Cognition*, 22(4), 411–430.
- Kosslyn, S. M. (1989). Understanding charts and graphs. *Applied Cognitive Psychology*, 3, 185–226.
- 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.
- Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, 11, 65–99.
- Lohse, G. L. (1993). A cognitive model for understanding graphical perception. *Human Computer Interaction*, 8, 353–388.
- Lowe, R. K. (1999). Extracting information from an animation during complex visual learning. *European Journal of Psychology of Education*, 14(2), 225–244.

- Narayanan, N. H., & Hegarty, M. (1998). On designing comprehensible interactive hypermedia manuals. *International Journal of Human-Computer Studies*, 48, 267–301.
- Peebles, D., & Cheng, P. C.-H. (2001a). Extending task analytic models of graph-based reasoning: A cognitive model of problem solving with cartesian graphs in ACT-R/PM. In *Proceedings of the 4th International Conference on Cognitive Modeling*, (pp. 169–174). Mahwah, NJ: Lawrence Erlbaum Association.
- Peebles, D., & Cheng, P. C.-H. (2001b). Graph-based reasoning: From task analysis to cognitive explanation. In *Proceedings of the Twenty Third Annual Conference of the Cognitive Science Society*, (pp. 762–773). Hillsdale, NJ: Lawrence Erlbaum Association.
- 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.
- Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, 171, 701–703.
- 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), 416–438.
- 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., & Mintz, F. E. (in press). Connecting internal and external representations: Spatial Transformations of scientific visualizations. *Foundations of Science*.
- Trafton, J. G., & Trickett, S. B. (2001). A new model of graph and visualization usage. In J. D. Moore, & K. Stenning (Eds.), *Proceedings of the Twenty Third Annual Conference of the Cognitive Science Society*, (pp. 1048–1053). Mahwah, NJ: Lawrence Erlbaum Associates.
- Trickett, S. B., Fu, W., Schunn, C. D., & Trafton, J. G. (2000). From dipsydoodles to streaming motions: Changes in representation in the analysis of visual scientific data. In *Proceedings of the Twenty Second Annual Conference of the Cognitive Science Society*, (pp. 959–964). Hillsdal, NJ: LEA.
- Trickett, S. B., Trafton, J. G., Schunn, C. D., & Harrison, A. (2001). "That's odd!" How scientists respond to anomalous data. In J. D. Moore, & K. Stenning (Eds.), *Proceedings of the Twenty Third Annual Conference of the Cognitive Science Society*, (pp. 1054–1059). Mahwah, NJ: LEA.
- Tufte, E. R. (1983). *The visual display of quantitative information*. Cheshire, CN: Graphics Press.
- Tufte, E. R. (1990). *Envisioning information*. Cheshire, CN: Graphics Press.
- Tufte, E. R. (1997). *Visual Explanations*. Cheshire, CN: Graphics Press.