Shedding Light on the Graph Schema

Raj M. Ratwani
(rratwani@gmu.edu)
George Mason University

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

Abstract

The current theories of graph comprehension have posited the graph schema as providing us the necessary knowledge to interpret any graph type. Yet, little is known about the nature of the graph schema, and no empirical data exist showing that there actually is a graph schema. In experiment 1 we show evidence that a graph schema does exist, and that graph schemas are not specific to each and every graph type. In experiments 2 and 3 we show that there is a different graph schema for typical and atypical graphs. We interpret these findings as evidence for a prototypical graph schema.

Introduction

When looking at Figure 1a and attempting to read-off the number of Widgets in Tray B, how does one have the necessary knowledge to be able interpret this specific type of graph? Given the large number of graph types (e.g. bar, line, dot, scatter, box plot, etc.) and the fact that the same symbol can represent completely different information in each of these graph types, how do we activate and use the information specific to each graph type? For example, the "dot" in a scatter plot as compared to the "dot" in a box plot mean very different things and in order to be able to interpret these different graph types, we have to be able to activate the appropriate knowledge.

The current theories of graph comprehension solve this problem by positing a "graph schema." Pinker (1990) suggests it is the graph schema that allows us to recognize specific types of graphs and allows us to find the desired information in a graph. Lohse (1993) suggests that the graph schema contains standard, learned procedures for locating information in the graph. Thus, when reading-off specific information from a graph, the current theories would suggest the following operations: (1) Early visual processes construct all possible relationships among graph elements, (2) Build a propositional representation of the graph, (3) Activate graph schema, (4) Devise the conceptual question, (5) Associate location of bar with each tray, (6) Associate each bar with values for each tray (7) Devise the conceptual message (Carpenter & Shah, 1998; Lohse, 1993; Pinker, 1990; Trickett, Ratwani, & Trafton, under review). The graph schema (step 3) is central to all current theories of graph comprehension. What is interesting is that there has been no empirical work to establish whether a graph schema really exists, and, if so, what its features are.

Our research goal was to use the mixing costs paradigm (Los, 1996, 1999) to investigate the nature of graph schemas. In the mixing costs paradigm there are blocks of pure stimuli, composed of items of the same type (i.e. all bar graphs), and blocks of mixed stimuli, which are composed of items of different types (i.e. bar graphs and line graphs). In the pure blocks, it is thought that because the stimuli are of the same type, each stimulus primes or activates the next and thus people are quick to respond to the stimuli. However, stimuli only prime or activate other stimuli that rely on the same mental representation. Thus, in the mixed blocks, because the stimuli are of different types, they may rely on different representations and result in a slower response as compared to the pure blocks. There are several other interpretations to mixing costs, but this interpretation is very prevalent (Los, 1996, 1999).
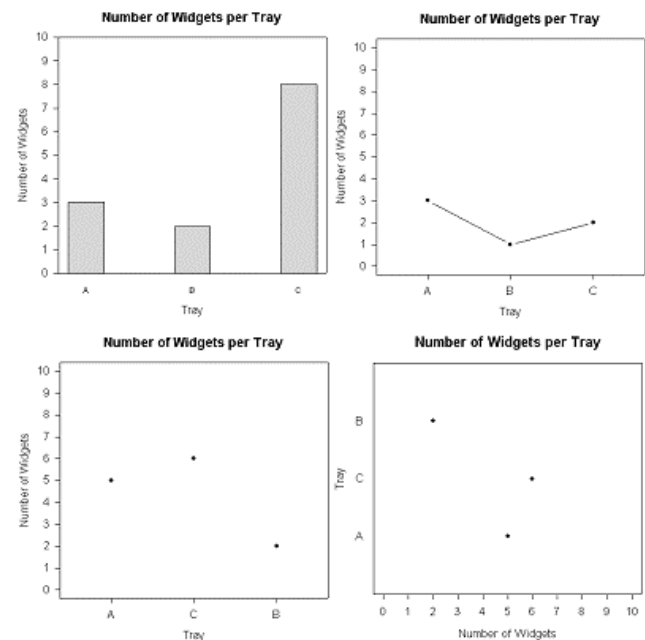
Figures 1a-d. (clockwise from upper left corner) Graphs used in experiments 1-3: bar graph, line graph, dot chart, scatter plot.

By using the mixing costs paradigm, we will be able to show which graphs share a similar mental representation. This internal representation, we believe, is what most graph comprehension theorists call the graph schema. We will describe in some detail exactly what we think a graph schema is in the general discussion. Assuming a graph schema or representation does exist, there appear to be several possibilities as to how the graph schema accounts for our ability to interpret different graph types. First, the schema may be graph specific; each and every graph type may have its own unique mental representation (when we say representation, we mean internal, mental representation). For example, a bar graph may have its own representation and a

line graph may have its own representation. In terms of priming, if each graph type relies on an entirely different representation, one graph type should not prime the other. Switching between graph representations takes time, and since the particular graph representation is not primed, there is a time cost. Thus, graph readers should be slower at responding to a particular graph type in the mixed condition as compared to the pure condition. A second possibility is that there is a general graph schema; we have one single graph representation which is used for each and every graph type. If this is the case, any given graph type should prime any other graph type since they rely on the same graph representation. This means graph readers reaction times to pure conditions of one graph type should be the same as their reaction times to mixed conditions of two graph types, since the same representation is being primed in both conditions. A third possibility, which we believe, is that there are two prototypical graphs – bar and line graphs. These prototypes are activated any time there is a graph, but if the graph type is not a bar or line graph, additional time is needed to interpret the graph and change the mental representation.

In these experiments we examine different graph types which vary in their prototypicality to determine the nature of the graph schema.

Experiment 1

In experiment 1 we used three stimuli types: bar graphs, line graphs, and text. We had pure blocks of each graph type (e.g. pure bar graphs and pure line graphs) and mixed blocks of each stimulus type (e.g. mixed bar graphs and text). First, based on the fact that there has been little research on graph schemas, we wanted to find empirical evidence that a graph schema exists. If a graph schema exists, we would expect that the conditions of pure graphs would be faster than the conditions of graphs and text. In the pure conditions, this graph representation would remain highly activated since it is being primed. However, in the mixed conditions, if there is a graph representation, it would not be primed by the text, resulting in slower reaction times as compared to the pure condition. In the case of no graph schema, there should be no priming in the pure condition or the mixed graph and text condition, resulting in the same reaction times to the graphs in both conditions.

Second, to examine the nature of the graph schema, we began by examining whether the graph schema was graph specific or graph general by using two prototypical graphs. If a specific graph schema exists, we would expect a time cost associated with activating the correct graphical representation in the mixed bar graph and line graph condition as compared to the pure graph conditions. Because each of these graphs relies on a different representation, each time either graph type is viewed, that specific representation for that graph type must be activated; however, in the pure conditions the graphs are of all the same type, so the representation remains activated and thus there is priming and no time cost. For example, because a bar graph may have a specific bar graph schema, the reaction times to the condition of pure bar graphs should be fast since the representation remains highly

activated. The response to bar graphs in the mixed condition of bar graphs and line graphs should be slower since the bar graph representation has to be activated each time a bar graph is viewed.

If these two graph types rely on a general graph schema, we would expect that bar graphs would activate line graphs and line graphs would activate bar graphs. Because each graph type may rely on the same representation, we would not expect to find differences in the pure graph conditions and the mixed bar graph and line graph conditions. Regardless of the graph type, the graph representation will remain activated in both the pure and mixed conditions of graphs resulting in no time costs. The prototypical graph view makes the same predictions as the general graph view– we will explore less prototypical graph types in later experiments.

Method

Participants

Twenty-one George Mason University undergraduate students participated in this experiment for course credit.

Materials

The materials consisted of eighty bar graphs, eighty line graphs and forty text sentences. Each of the graphs depicted the number of Widgets, ranging from 0-9, in three different trays: A, B, and C (see Figure 1a and 1b for examples); each sentence contained a number ranging from 0-9. We chose to use text sentences because we wanted non-graphical and non-spatial stimuli. All of the graphs and text were randomly generated, and the locations of trays A, B, and C were randomly assigned. For each of the graphs in the experiment the participant was asked the same question, "How many Widgets are there in Tray B?", in order to minimize working memory load of remembering the question (Peebles & Cheng, 2003). For each of the text sentences, the participants were asked what number appears in the sentence. For example, the sentence may be "There were two cars in the driveway"; subsequently, the participant would enter "two". While the text sentences are of a different subject matter than the graphs, the primary goal was to examine whether some kind of graph schema exists.

Design

Five different conditions were setup in this experiment, with each condition containing forty stimuli. There were two pure conditions: pure bar graph and pure line graph. Each of these conditions consisted of 40 similar graph types, for example, the pure bar graph condition consisted of 40 bar graphs. There were three mixed conditions: mixed bar graphs and line graphs, mixed bar graphs and text, mixed line graphs and text. Each of these conditions also contained 40 stimuli, 20 of each respective type. The stimuli order in each condition was randomly assigned. Throughout this paper, we refer to the pure conditions as follows: line (pure); this means we are referring to the average reaction time in the pure line graph condition. We refer to the mixed conditions as follows: bar (mixed bar/line); this means we are referring to the average

reaction time for the bar graphs in the mixed bar graph and line graph condition.

Procedure

The order in which the five conditions were presented to each participant was randomly assigned according to a Latin squares design. The stimuli were presented to the participants over the computer. Each participant was instructed to respond to the number of Widgets in Tray B when viewing a graph, and to respond to the number in the sentence when viewing a sentence, by entering the numerical value into the computer by using the keypad on the keyboard. Before each condition, the participant performed three practice trials to ensure they understood the graph type, the interface, and the task. Each participant was instructed to go through each graph or text as quickly and accurately as possible. Once the participant entered the value, the next stimulus automatically appeared for the participant to respond to. After the condition was completed, the experimenter entered the room and loaded the next condition for the participant.

Results and Discussion

The reaction times for incorrect responses and reaction times that were three standard deviations away from the average were removed from all analyses. In the pure conditions, the participant's reaction times were averaged across all stimuli. In the mixed conditions the participant's reaction times were averaged across similar stimuli. For example, in the mixed bar graph and line graph condition, the reaction times of all the bar graphs were averaged, and the reaction times for all the line graphs were averaged. This was done for each of the mixed conditions.

An omnibus ANOVA showed there was a significant difference among the conditions, $F(8, 160) = 6.6$, $p < 0.0001$, $MSE = 33538$. Specific comparisons were conducted using pairwise t-tests with the Tukey HSD adjustment for multiple comparisons. Figure 2 shows the difference scores between conditions based on stimuli type.

The pure conditions serve as a baseline for comparison to the mixed conditions. For example, the first set of three numbers on the far left of the figure represents average reaction times for bar graphs. "Pure" is the condition of bar (pure), "Line" is the condition of bar (mixed bar/line), and 'Text" is the condition of bar (mixed bar/text). The average response time to the condition of bar (pure) was 1791 ms, the average response time to the bar (mixed bar/line) was 1810 ms, and 1964 ms for the bar mixed (bar/text). The bars above the numbers represent the difference in reaction times between the pure conditions and the mixed conditions. The "*" indicates a significant difference between the pure condition and the marked mixed condition via Tukey HSD. Thus, the difference between the bar (pure) (1791 ms) and the bar (mixed bar/text) (1964 ms) is shown graphically as 173 ms, which is a significant difference.
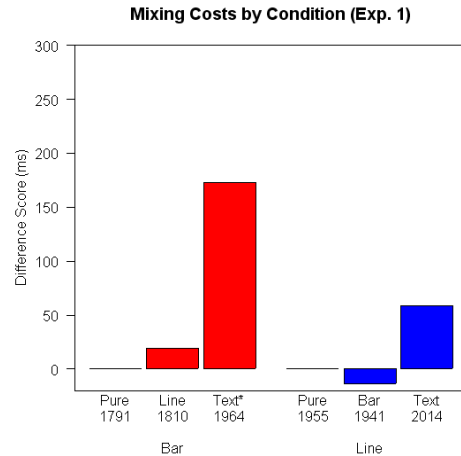

Figure 2. Average reaction times and difference scores.

We first wanted to find evidence of the existence of any kind of graph schema. The existence of a graph schema was evident by the significant difference in the bar (pure) condition and the bar (mixed bar/text) condition as evident in Figure 2, Tukey $p < .05$. This suggests that the text does not activate the bar graph representation since participants are slower at responding in the bar (mixed bar/text) condition, as compared to the bar (pure) condition. In the bar (mixed bar/text) condition, the bar graph representation has to be activated each time a bar graph is viewed, resulting in a time cost as compared to the bar (pure) condition. The line (pure) as compared to the line (mixed line/text) condition trended in the same direction; however, this difference was not significant. These comparisons suggest there is a graph schema, but is the schema graph specific or graph general?

If the graph schema is general, both bar graphs and line graphs should prime the same graph representation, so bars (mixed bar/line) should be as fast as bar (pure). If the graph schema is specific, bar graphs and line graphs should not prime the same representation, so bar (mixed bar/line) should be slower than bar (pure).

The reaction times for the bar (pure) condition were not significantly different from the bars (mixed bars/line) condition, $p = .67$. Likewise, the line (pure) condition was not significantly different from the line (mixed bars/line) condition, $p = .80$ (see Figure 2). Because the line graphs activated the bar graphs and the bar graphs activated the line graphs equally as well as each graph type was activated in their respective pure conditions, this suggests that both bar graphs and line graphs rely on the same graph representation. If the schema was graph specific, we would expect the mixed conditions to be slower than the pure conditions, since a different representation would have to be activated each time a different type of graph appeared.

Experiment 2

In experiment 1, each graph type primed the other, suggesting they rely on the same graph representation. We did not find mixing costs between the conditions of pure graph types and

the conditions of mixed graph types; this is evidence against the specific graph schema view, but does not differentiate between the general or prototypical graph schema views. The general graph schema view predicts that there will never be mixing costs between any different graph types since they all rely on the same mental representation. The prototypical graph view, however, predicts that there will be mixing costs for less prototypical graph types. Thus, we chose both a very typical (line graph; Figure 1a) and a very atypical graph type (Cleveland's dot chart; Figure 1c) for experiment 2. In the dot charts (Cleveland, 1985), the numerical scale appears on the x-axis and the labels appear on the y-axis.

According to the prototypical based graph schema view, since dot charts are very atypical, they should have a different representation than the line graphs. Based on this view, in the dot (mixed dot/line) condition, the dot chart representation must be activated each time the dot chart appears, whereas in the dot (pure) condition this representation should remain activated. Thus, participants should be faster at responding in the dot (pure) condition as compared to the dot (mixed dot/line) condition, indicating that line graphs do not activate dot charts.

The general graph schema view would suggest there should be no mixing costs between these two different graph types; participants should be equally fast in the pure graph conditions as they are in the mixed graph conditions. Since all graph types rely on the same graph representation, it should be equally activated, and there should be no differences between conditions. If no mixing costs are found between the pure graph conditions and the mixed graph conditions, this would be further evidence for the general graph schema view.

Method

Participants

Twenty George Mason University undergraduate students participated in this experiment for course credit.

Materials

The materials were similar to those used in experiment 1, except eighty dot charts replaced the eighty bar graphs (see Figure 1c for an example); the line graphs and text remained the same. The same questions asked in experiment 1 were asked in this experiment as well.

Design

The design was similar to experiment 1. The two pure conditions were: pure dot chart and pure line graph. The three mixed conditions were: dot chart and line graph, dot chart and text, and line graph and text.

Procedure

The procedure was the same as experiment 1.

Results and Discussion

The statistical analyses conducted were the same as in experiment 1. The omnibus ANOVA was significant, $F_{(8, 160)} = 17.5$, $p < 0.0001$, $MSE = 31023$, indicating that there was a significant difference in the conditions. Figure 3 shows the average reaction times by condition and also shows the difference scores. Consistent with experiment 1, the dot (pure) condition was significantly faster than the dot (mixed dot/text) condition, Tukey $p < .05$. The line (pure) condition as compared to the line (mixed line/text) condition trended in this direction, but was not significant. These results were consistent with previous experiments and indicate that there is a graph schema.
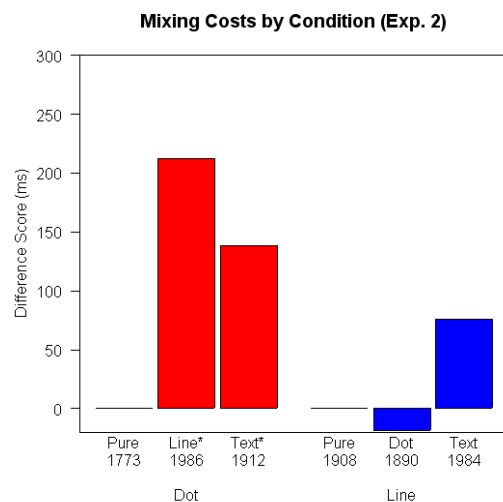

Figure 3. Average reaction times and difference scores.

As the prototypical graph schema view would suggest, the dot charts are not primed by the line graphs. The dot (pure) condition was significantly faster than the dot (mixed dot/line) condition as illustrated in Figure 3, Tukey $p<.05$. In the dot (pure) condition, because the stimuli are all dot charts, this representation remains activated since each dot chart primes the next. However, in the dot (mixed dot/line) condition, because the dot charts and line graphs rely on different representations, the dot chart representation is not activated by the line graphs. Thus, each time a dot chart appears, its representation has to be activated, resulting in a slower reaction time to the dot charts in the mixed condition than in the pure condition.

Because the line graph is a prototypical graph type, there are no mixing costs between the line (mixed dot/line) condition and the line (pure) condition, $p = .65$. Importantly, these asymmetric mixing costs show that the time cost associated with the dot charts was not simply due to a switch cost associated with the different stimuli types. If the mixing costs were due to a switch cost, we would see a similar time cost in the line (pure) as compared to line (mixed dot/line).

These results suggest there is a different graph schema for dot charts and line graphs. Thus, there is not a true general graph schema, and there also seems to be a prototypical graph schema.

While experiment 2 demonstrates that there is not a general graph schema, it could be that the reason that there are mixing costs between dot (pure) and dot (mixed dot/line) is that the dot chart is not only atypical, but also that it has a completely different orientation of axes from the line chart. That is, the dot chart is read in a completely different manner: the participant has to look at the y-axis to find the "B" label. This switching of the labels on the axes between graph types could be responsible for the mixing costs, not the a-prototypicality of the dot chart.

Experiment 3 thus used an atypical graph type (a scatter plot) that had the same axis orientation as the line graphs.

Experiment 3

Experiment 3 compared a prototypical graph type (line graph; Figure 1a) and an atypical graph type (scatter plot; Figure 1d). According to the prototypical graph view, participants should be faster to respond to a scatter (pure) condition than a scatter (mixed scatter/line) condition. This mixing cost would be attributable to the fact that in the mixed condition, the scatter plot representation has to be activated each time a scatter plot is viewed, resulting in an additional time cost, whereas in the pure scatter plot condition, the activation of the scatter plots remains high. The activation of the line graphs, on the other hand, may not be as influenced by the scatter plots since the line graph is a prototypical graph type.

The general graph schema view would suggest there should be no mixing costs between these two different graph types; participants should be equally fast in the pure graph conditions as they are in the mixed graph conditions. Since all graph types rely on the same graph representation, it should be equally activated and there should be no differences between conditions. If no mixing costs are found between the pure graph conditions and the mixed graph conditions, this would be evidence for the general graph schema view; the results of experiment 2 could be attributed to the fact that the orientation of the graphs was different, not the prototypicality of the graphs.

Method

Participants

Twenty-one George Mason University undergraduate students participated in this experiment for course credit.

Materials

The materials were similar to those used in experiment 1; except eighty scatter plots replaced the eighty bar graphs (see Figure 1d for an example); the line graphs and text remained the same. The same questions asked in experiment 1 were asked in this experiment as well.

Design

The design was similar to experiment 1. The two pure conditions were: pure scatter plot and pure line graph. The three mixed conditions were: scatter plot and line graph, scatter plot and text, and line graph and text.

Procedure

The procedure was the same as experiment 1.

Results and Discussion

The statistical analyses conducted were the same as experiment 1. The omnibus ANOVA was significant, $F_{(8, 160)} = 2.2$, $p < 0.05$, $MSE = 64059$, indicating that there was a significant difference in the conditions. We first wanted to replicate the findings in experiments 1 and 2, which suggested that there was some kind of graph schema based on the fact that the text did not activate the graphs. Similar to experiment 1, the scatter (pure) condition was significantly faster than the scatter (mixed scatter/text) condition as illustrated in Figure 4, Tukey $p < .06$. The line (pure) condition as compared with the line (mixed line/text) condition trended in the expected direction, but as in experiments 1 and 2, this relationship was not significant. These findings are consistent with experiment 1 and lend further support to the existence of a graph schema.
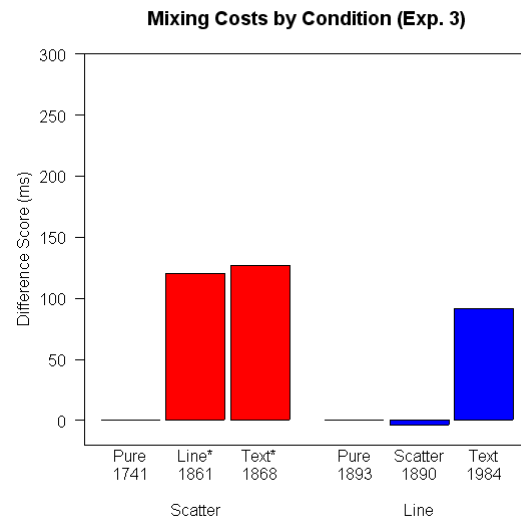

Figure 4. Average reaction times and difference scores.

Next we compared the scatter (pure) condition to the scatter (mixed scatter/line) condition, and also the line (pure) condition to the line (mixed scatter/line) condition. Consistent with the prototypical graph schema view, the scatter (pure) condition was significantly faster than the scatter (mixed scatter/line) condition, as illustrated in Figure 4, Tukey $p < .05$. This time cost in the scatter (mixed scatter/line) condition suggests that the default representation has to be changed to fit the scatter plot representation, resulting in this greater time cost as compared to the scatter (pure) condition. The line graphs apparently did not prime the scatter plots as the general graph schema view would suggest.

Interestingly, the line (pure) condition was not significantly different from the line (mixed scatter/line) condition, $p = .94$. Participants were just as fast at reading line graphs in the line

(pure) condition as compared to the line (mixed scatter/line) condition. These asymmetric mixing costs suggest that our findings are not due to switch costs associated with the differences in the stimuli. The line graphs do not incur a cost, once again suggesting that the prototypical graph schema includes a line graph.

General Discussion

There are many different graph types which use similar symbols in different ways to represent data. The current theories of graphs comprehension (Carpenter & Shah, 1998; Lohse, 1993; Pinker, 1990) rely on the notion of a graph schema to account for how graph readers have the necessary knowledge to be able to interpret any given graph type. We outlined three possibilities for the graph schema: the graph specific view, the graph general view and the prototype view.

Experiment 1 demonstrated, first, that a graph schema does exist, and second, that the graph schema is not graph specific. The bar graphs and line graphs primed each other in the mixed conditions, suggesting that these graph types rely on the same representation.

Experiment 2 sought to examine whether the graph representation was graph general or prototypicality based. Participants were slower in the dot (mixed dot/line) condition than the dot (pure) condition, suggesting that the dot chart relies on a different graph representation. Importantly, there was no difference in the line (pure) condition and the line (mixed dot/line) condition, suggesting that the prototypical graph schema includes a line graph. These asymmetric mixing costs also show that our findings are not due to switch costs from different stimuli types.

Experiment 3 further supported the prototypicality based view. We manipulated prototypicality with the graphical pattern and kept the orientation of the axes the same, which resulted in faster response times in the scatter (pure) condition as compared to the scatter (mixed scatter/line) condition. However, similar to experiment 2, the line graphs did not incur a mixing cost.

How do people use a graph schema? According to our view, any time that someone sees a graph, the prototype graph schema is retrieved. If the graph type being examined is a line or bar graph, the comprehension and usage of that graph proceeds smoothly because the default values already match the graph type. If, however, the graph type being examined is not a line or bar graph – it is a scatter plot or a dot chart or a box plot – the default values of the graph schema must be changed to fit that graph type. Alternatively, a graph specific schema must be activated.

What exactly is the graph schema? We believe the graph schema is our mental representation of how to read a graph type; it is the graph schema that gives us the necessary knowledge to interpret a specific graph. Prototypical graphs like line and bar graphs are activated more easily than atypical graphs. Note that prototypicality does not necessarily mean that it is an easier, better or faster graph to use – it just means that that representation is the default when we see a graph. Prototypicality could vary, as it does in other domains (Medin & Atran, under review). In addition, our participant population was composed of undergraduate students; it is likely that more experienced graph users have a more elaborate schema.

This research does not focus so much on the other default values, or even what the other slots could make up the graph schema; future research will be necessary for that question.

Acknowledgements

References

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. (1985). The elements of graphing data. Monterey, CA: Wadsworth.

Lohse, G. L. (1993). A cognitive model for understanding graphical perception. Human Computer Interaction, 8, 353-388.

Los, S. A. (1996). On the origin of mixing costs: Exploring information processing in pure and mixed blocks of trials. Acta Psychologica, 94(2), 145-188.

Los, S. A. (1999). Identifying stimuli of different perceptual categories in pure and mixed blocks of trials: evidence for stimulus-driven switch costs. Acta Psychologica, 103(1-2), 173-205.

Medin, D. L., & Atran, S. (under review). The native mind: biological categorization, reasoning and decision making in development across cultures. Psychological Review.

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.

Trickett, S. B., Ratwani, R. M., & Trafton, J. G. (under review). Real World Graph Comprehension: High-Level Questions, Complex Graphs, and Spatial Cognition.