A cognitive architecture-based model of graph comprehension

David Peebles (d.peebles@hud.ac.uk)
Department of Behavioural and Social Sciences,
University of Huddersfield,
Queensgate, Huddersfield, HD1 3DH, UK

Abstract

I present a model of expert comprehension performance for 2×2 "interaction" graphs typically used to present data from two-way factorial research designs. Developed using the ACT-R cognitive architecture, the model simulates the cognitive and perceptual operations involved in interpreting interaction graphs and provides a detailed characterisation of the information extracted from the diagram, the prior knowledge required to interpret interaction graphs, and the knowledge generated during the comprehension process. The model produces a scan path of attention fixations and a symbolic description of the interpretation which can be compared to human eye movement and verbal protocol data respectively, provides an account of the strategic processes that control comprehension, and makes explicit what underlies the differences between expert and novice performance.

Keywords: Graph comprehension, ACT-R

Introduction

Working with graphs is a complex skill that requires specific knowledge of the representational system being used together with a set of procedures to map spatially represented information in the graph with a set of propositions that specify quantitative and qualitative relationships between the entities represented. Providing a detailed account of this skill therefore requires one to specify a number of core assumptions including: what and how information is encoded in the diagram, what and when information is obtained from the diagram by the user during a task, what and how prior graph knowledge is stored and utilised, and what new knowledge is created during the process. In addition, one must also specify the strategies people employ to carry out different tasks and how much these strategies use information in the diagram and in stored internal representations.

There have been several attempts to provide detailed process models of different aspects of graph use. Models are constructed from sets of perceptual and cognitive operators (e.g., encode the value of an indicator, make a spatial comparison between indicators (Gillan, 1994), compare two digits in working memory, or make a saccade (Lohse, 1993)), obtained either from task or verbal protocol analyses. Lohse (1993) and Gillan (1994) have produced models of question answering with several different graph types (including line graphs, bar charts and scatter plots) by constructing sequences of operators (each of which has an associated execution time) to generate predicted scan paths across the graph and total task completion times which can be compared to human data.

Other researchers have procedurally analysed graph use for different purposes. For example, Casner (1991) identified a set of perceptual and cognitive operators to construct models of several graph-based tasks which informed an automated system that generated graphical representations most suited to the tasks commonly undertaken with them. A similar method was adopted by Tabachneck-Schijf, Leonardo, and Simon (1997) in their analysis of an economics expert's construction of a graph while explaining the principle of supply and demand which they then used to develop a computational model incorporating both diagrammatic and propositional representations.

More recently, the cognitive modelling of reasoning with information displays has been advanced by the development of *cognitive architectures*; computational theories of the large-scale structure of the mind providing accounts of how cognition is controlled and how knowledge is encoded, stored, retrieved and utilised (e.g., ACT-R (Anderson, 2007), EPIC (Meyer & Kieras, 1997), and Soar (Laird, Newell, & Rosenbloom, 1987)).

The first two of these architectures incorporate theories of visual processing and motor control which allows modellers to produce more detailed accounts of the information obtained from the display during the task. For example Peebles and Cheng (2003) used ACT-R to produce a computational model of question answering using two different types of line graph. Their model generated saccades and fixations as it answered each question which, together with task completion times, were compared to human data. In addition, the model was able to account for human scan paths in terms of the varying demands on memory imposed by different questions.

The Peebles and Cheng study, as did those by Lohse (1993) and Gillan (1994), investigated question answering in which participants were given items of information and were required to produce associated information using different processes, including identification (e.g., "In 1997, what was the value of gas?" (Peebles & Cheng, 2003)), comparison (e.g., "In 1977 did tin cost less than sulphur?" (Lohse, 1993)), and arithmetic computation (e.g., "What is the sum of A, B, and C?" (Gillan, 1994)).

While these are important tasks, particularly for investigating sequences of elementary processes, it could be argued that they do not necessarily reflect how many people normally work with graphs and that they do not address the important prior comprehension stage where labels and graphical features are encoded, associated, and interpreted (Carpenter & Shah, 1998).

Comprehension requires knowledge of the conventions used in the graph to represent data and other facts such as how labels are to be interpreted based on their location. The output of the process is assumed to be a set of knowledge structures that represent the variables and graphical features together with structures that encode knowledge about the quantitative

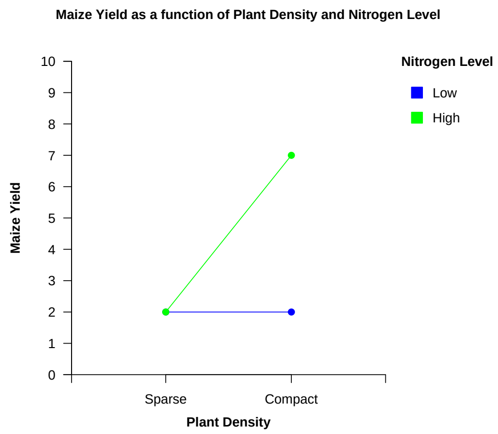

Figure 1: One of the eight line graphs used in the expert study (Peebles & Ali, in preparation).

or qualitative relationships between the variables depicted.

A prime example of a scenario where people encounter a graph with the sole aim of comprehending the relationships between variables (as opposed to identifying trends or individual values for example) is the analysis of data from factorial experiments. The simplest form of factorial design is the *two-way factorial design*, containing two factors, each with two levels, and one DV. Statistical analysis of these designs typically results in a 2×2 matrix of mean values of the DV corresponding to the pairwise combination of the two levels of each IV. Interpreting the results of even these simplest of designs accurately and thoroughly is often not straightforward however, but requires a significant amount of conceptual understanding—for example the concepts of simple, main, and interaction effects. As with most other statistical analyses however, interpretation can be eased considerably by representing the data in diagrammatic form.

Data from two-way factorial designs are most often presented as either line or bar graphs—variously called *interaction* or *ANOVA* graphs. An examples line graph is shown in Figure 1. Because the data come from pair-wise combinations of the IV levels, the rules for interpreting interaction graphs are quite specific however and sufficiently different from other more frequently encountered line graphs that simply applying general interpretive rules will not prove particularly helpful (other than for obtaining the DV values of specific conditions etc.). The key elements of knowledge to be obtained from interaction graphs are the simple, main and interaction effects of the IVs and these have to be identified in specific features of the graph.

In a series of studies, Peebles and Ali have observed and recorded novices (undergraduate psychology students) and experts (cognitive science professors and postgraduate researchers) interpreting interaction graphs like the one in Figure 1 (Peebles & Ali, 2009; Ali & Peebles, in press; Peebles & Ali, in preparation). These studies have shown that without knowledge of the appropriate interpretive rules, novices' interpretations are often limited to qualitative descriptions of differences between conditions and can be skewed by the different Gestalt principles of perceptual organisation (Wertheimer, 1938) operating in the graph. In contrast, expert users are able to employ their knowledge of which graphical features represent which effects to identify relationships between variables much more rapidly and accurately with no prior knowledge of the domain variables being represented in the graph.

The purpose of the research reported here is to develop a computational model of graph comprehension that specifies the processes underlying both expert and novice behaviour with sufficient detail and comprehensiveness to satisfy all of the criteria outlined at the beginning of this paper. Specifically, the model aims to provide a precise account of the minimum information required to interpret interaction graphs appropriately together with a hypothesis as to the nature of the processes involved in representing and interpreting that information. The model is developed within the ACT-R cognitive architecture and therefore embodies assumptions about the nature of the mental representations and the computations that form the strategies used to generate new representations. Finally, the model provides an explanation for the differences between expert and novice interpretations.

A model of graph comprehension

Space limitations preclude a detailed description of ACT-R here. However a comprehensive account of the cognitive architecture can be found in Anderson (2007). In summary, ACT-R consists of a set of modules that acquire information from the environment, process information, and execute motor actions to achieve goals. ACT-R has memory stores for declarative and procedural knowledge. The former consists of a network of knowledge chunks while the latter is a set of production rules. Cognition proceeds via a pattern matching process that attempts to find production rules with conditions that match the current state of the system and tasks are performed through the successive actions of production rules.

ACT-R also incorporates a subsymbolic level of computations that govern memory retrieval and production rule selection and which allow models to account for widely observed recency and frequency effects on retrieval and forgetting. Subsymbolic computations also underlie ACT-R's different learning mechanisms.

For tasks involving displays and other devices, task environments can be defined to be acted upon by the model. The graphs used in this study are defined as sets of visual objects (lines, circles, rectangles, and text) with certain features (size, colour) at specific x-y coordinates on a 2D window.

The graph comprehension model is based on verbal protocol data from novice and expert users (Peebles & Ali, 2009; Ali & Peebles, in press; Peebles & Ali, in preparation). In these studies, verbal statements recorded during the comprehension task were coded and categorised in terms of their functional role and content (e.g., "an association between a level and its identifier"; "a comparison between the two legend variable levels for one of the levels of the x axis variable") to produce a set of common interpretive operations.

The verbal protocols indicate that comprehension is typically carried out in two main phases: (a) a variable identification stage followed by (b) a pattern recognition and description stage. The protocols also reveal that experts and a large proportion of novices rarely report specific DV values, but typically produce qualitative descriptions of the differences between conditions.

In the first stage, the three variables are identified, categorised as dependent or independent according to location, and the latter associated with their levels, which in turn are associated with identifiers (left or right position for the x axis variable and colour for the legend variable).

In the second stage, distances between plot points are observed and compared, with the results being used to probe long-term declarative memory for interpretive knowledge. If this is successful, the retrieved knowledge is used to provide an interpretation. If there is no interpretation available however, the model will simply describe the identification or comparison process being carried out. Interpretive operations are carried out until either a full interpretation is produced or until no other operations are available or identified.

Representing and encoding information in the graph

The key information that the model must encode and utilise from the graph representation is a set of four x-y coordinate points and the spatial distances between them. Although the model processes symbolic representations, it assumes that spatial information is initially encoded quantitatively and subsequently categorised into qualitative descriptions. The perceptual processes by which the information is obtained or represented are not specified in detail, although it is assumed that it is via a set of prior *elementary perceptual tasks* (Cleveland & McGill, 1984). Cleveland and McGill (1984) identified ten such tasks (e.g., length, direction, area, and position on a common scale) as the "perceptual building blocks" of graph comprehension that encode quantitative information from graphical elements.

At least two such elementary perceptual tasks are assumed to be required for these graphs. The first is *position on a common scale* and this is the primary comparison that takes place. It is assumed that readers initially encode the spatial distance between plot points into a quantitative representation (the proportion, p, of the distance to the overall length of the y axis) and then categorise this ratio according to size. For this model six categories were assumed: "no" ($p = 0$), "very small" ($0 < p < 0.2$), "small" ($0.2 =< p < 0.4$), "moderate" ($0.4 =< p < 0.6$), "large" ($0.6 =< p < 0.8$), and "very large" ($0.8 =< p =< 1.0$). Although it is an assumption of the model that distances are categorised in this way, the exact processes by which these final categories are produced are not specified in detail.

The second process that is assumed readers can perform is to compare the magnitude of two distances and produce a symbolic description of the difference. The elements formed for this comparison are assumed to be the result of Gestalt processes of perceptual organisation (Ali & Peebles, in press; Kosslyn, 1989; Pinker, 1990) which allow users to group objects by colour or proximity. This comparison also allows users to perceive and compare the directions of the two differences (i.e., the relative sizes of the variables' level values). These relative values produce the various patterns such as crossed, parallel and diverging lines which are recognised and interpreted by expert users.

Prior graph knowledge

Two forms of declarative knowledge are used during the task: prior knowledge relating to how the graph represents information and knowledge of the variables and their relationships generated during the comprehension process itself.

There are three core items of knowledge required to interpret interaction graphs. Two are common to many Cartesian graphs and concern (a) the typical allocation of the dependent and independent variables to the graph axes and legend and (b) the principle that the distance between two graphical elements is directly related to the magnitude of the relationship between the conceptual entities that the elements represent.

The third set of facts required are specific to the graph type and concern the graphical and spatial indicators of the three key important interpretive facts; simple effects, main effects, and interactions. The three indicators are (a) the distance between two plot points which indicates the size of the *simple effect* of the level jointly represented by those points, (b) differences in the y-axis location of the midpoints between two pairs of plot points which indicate the size of the *main effect* of the variable, and (c) differences in the inter-point distances between levels, combined with information about their point ordering, which indicates the size, and type of any interactions that may exist.

This knowledge is represented as symbolic structures in the model's long term memory and is currently the minimum required to indicate that the interpretative process has succeeded. It is possible however to add further causal knowledge relating to the various effects to allow the model to provide more detailed explanations of the relationships identified.

Generated knowledge

Several declarative knowledge structures are also generated during comprehension. The first is a set of related chunks that represent each variable, the levels associated with it, and the identifiers of each level. Three other knowledge structures are generated to accumulate and associate items of graph and interpretive information during a specific sub-task. In the expert model all knowledge retrieval requests will succeed, resulting in knowledge structures that associate qualitative descriptions of differences and their interpretation. These structures could then be used to produce verbal explanations.

For example one structure records the processing of an individual level which could produce the explanation: "The difference between the two values for high plant density is very large so there's a very large simple effect of high plant density" while another records the information accumulated when comparing the average values of two levels of one variable (e.g., "There is a large difference between the fasting levels; high fasting generally resulted in greater glucose uptake than low fasting, which indicates a large main effect of fasting"). The third stores the results of comparing the lengths and point ordering of two levels (e.g., "Although the effect size of the cement type levels is the same, the direction of their effects is different so that means there is an interaction between the two independent variables".)

Finally, a representation is produced when a simple comparison between points is made which does not associate an interpretation (e.g., "When the nitrogen level is high, maize yield is much greater for compact plants than for sparse plants".).

The comprehension process

In the Appendix is an output trace produced by the model as it carries out the comprehension task using the graph in Figure 1, with each line in the trace representing one or more steps in the process (variable names have been shortened to allow the lines to fit the format of this paper). The text in square brackets is information currently being processed that has either been obtained from the graph or retrieved from declarative memory[1].

In the trace, numbers in square brackets represent the perceptual difference between two objects on the screen. These are subsequently translated into qualitative size judgements according to the categories described above. Other text in the output is simply to indicate other events (e.g., goal setting or memory retrieval failures) or to clarify what a particular knowledge element represents.

As previously intimated, the model assumes that comprehension proceeds after an initial phase of variable identification, a process usually initiated by reading the title. Currently when the model reads the title the three words that name variables are identified by retrieving previously defined word category information from declarative memory. This mechanism is undoubtedly simplistic and currently substitutes for a more complex knowledge retrieval process that is assumed to take place.

The model then seeks items of text at the left, right and lower regions of the display. When each variable label is identified, the model identifies it as a particular type according to its location and then, associates the independent variables with their level labels by identifying nearby text. The model also associates each of the four levels with its physical attribute; left, right, blue and green and uses these labels when processing the graph. This is consistent with verbal protocol and eye movement data from our studies showing that graph readers often produce an interpretation and then must re-read the appropriate label in order to identify which particular level is being processed.

When the three variables have been processed, the model then attends to the pattern produced by the four coordinate points in the plot region and then selects a particular feature or pair of features to process. The probability of selecting a particular feature to process may depend on a number of factors, including visual salience and pattern familiarity. For example, a large difference between objects, or parallel or crossing lines may draw the user's attention and lead them to attempt to interpret the feature first. Although it is possible to incorporate these processes for the model to select features in any order, for simplicity, the current model selects features in the order: simple, followed by main, and finally interaction effects.

These three effects are identified by different indicators in the graph. The size of the simple effect of a level is indicated by the distance between the level's two plot points while the main effect of a variable is indicated by the difference in the y-axis location of the midpoints between the variable's two pairs of plot points. Finally, the nature and size of interaction effects are indicated by differences in the inter-point distances between levels, combined with information about their point ordering.

The model represents the interpretation process by a set of production rules for each indicator type. When the appropriate condition occurs (i.e., the model is directing attention to the plot region), individual production rules fire to draw attention to specific indicators. The indicator (a spatial distance, difference or order comparison), is extracted from the pattern and (together with information about what the indicator is) used to probe declarative memory for an interpretation consisting of the name and size of the effect. For example on line 29 of the trace the model identifies that there is no difference between the plot points on the left of the display and then retrieves the knowledge that this indicates that there is no simple effect of sparse plant density (these labels being obtained by seeking the text below the points being observed).

For each indicator, if the memory retrieval attempt fails, the model simply describes the difference being attended to. This is demonstrated in lines 37 and 38 of the trace which compare the levels of the legend variable for each of the x axis variable levels and which correspond to the statement "when plant density is sparse, low and high nitrogen levels are the same but when plant density is compact, the high nitrogen level is greater than the low nitrogen level". This form of statement is very common in novice graph users.

Once a recognition production rule fires to initiate the process, a chain of subsequent productions is triggered which obtains further information from the graph and declarative memory until an interpretation is produced. The current production set is sufficient to process any 2×2 data set of three variables to produce an appropriate interpretation similar to

[1]A video of the model interpreting all eight graphs from the expert study (Peebles & Ali, in preparation) can be viewed at `http://youtu.be/z2kAwrOrjIM`

the trace in the Appendix.

Discussion

Comprehending and reasoning with graphs requires a wide range of perceptual and cognitive operations sequenced together in various combinations to perform specific tasks. The type and sequence of operators involved in a task may differ depending on a number of factors, including the graph or domain knowledge of the user, the type of graph being used, or individual cognitive factors such as working memory capacity (which may determine the relative frequency of memory retrieval requests and saccades to graph labels etc.).

Graph comprehension is an important area to study therefore because it provides an opportunity to investigate how environmental and internal factors interact to produce behaviour. In addition, graph-based tasks can be analysed using behavioural measures such as eye movements and concurrent verbal protocols to provide insights into what and when information is being processed during the course of the activity.

Computational modelling is a valuable tool for developing and testing hypotheses about the representations and mechanisms necessary for cognitive tasks as it provides a formalism for characterising them, requires one to be explicit about the boundaries of one's model in terms of which processes are being defined precisely and which are not, and allows one to explore the consequences of particular assumptions (McClelland, 2009).

Developing models within a cognitive architecture such as ACT-R provides the additional benefit of allowing the model to incorporate a large number of assumptions regarding issues such as knowledge representation, cognitive control, visual attention, learning and forgetting etc., all of which are supported by previous empirical research. In addition, ACT-R's vision module includes mechanisms that allow models to simulate certain Gestalt principles of perceptual organisation, which are regarded as playing a crucial role in the visual processing of graphical representations (Kosslyn, 1989; Pinker, 1990). Specifically, the comprehension model associates variables and their levels, and levels with their colour identifiers using mechanisms that are functionally equivalent to the Gestalt laws of *proximity* and *similarity* respectively.

The model described above represents an initial attempt to specify at a detailed algorithmic level the representations, cognitive processes, and strategies involved in comprehending interaction graphs. It provides a precise account of the graph knowledge required and the spatial information necessary to interpret the graph accurately and specifies a control structure that determines the flow of information during the task to generate a set of knowledge representations, saccades and fixations over the graph, and a sequence of output statements which are largely consistent in terms of order, function and content with verbal protocols produced by expert users.

The assumptions of the model imply that to interpret interaction graphs accurately, novices must acquire three forms of graph-specific knowledge: an understanding of what effects the different distances and spatial differences in the graph indicate, the relationship between distance and effect size, and how the various combinations of distance differences and point orders can be interpreted in terms of the interactions between the IVs. The model provides a precise specification of the relatively small amount of knowledge required and a clear demonstration of its sufficiency to interpret the graphs.

The current model can be considered a first approximation to a more detailed model that incorporates additional factors to broaden the scope of behaviour accounted for. Previous studies have shown that comprehension performance varies quite widely, even between experienced users (Peebles & Ali, 2009; Ali & Peebles, in press; Peebles & Ali, in preparation). For example, the order in which effects were identified varied, often as a result of the relative visual salience of the graphical features being displayed (e.g., very large main effects were often identified rapidly). Also, explicitly identifying simple effects was uncommon and other effects were sometimes overlooked by experienced users.

This variation in performance is no doubt due to a number of factors including the different effects of visual salience and Gestalt principles of perceptual organisation operating (Ali & Peebles, in press), and varying levels of graph knowledge and working memory capacity etc. In addition, previous studies compared expert and novice performance on both bar and line graph formats and showed that the interpretations of all users (but novices in particular) were affected by the format used. Specifically, line graphs users are influenced to attend to the legend variable while bar graph users attend to the two IVs more equally (Peebles & Ali, 2009; Ali & Peebles, in press; Peebles & Ali, in preparation). Broadening the scope of the model further, other factors such as domain knowledge and the number of variable levels (Shah & Freedman, 2011) should also be addressed.

The current model provides a solid basis from which to explore hypotheses concerning the mechanisms underlying this broader range of behaviour. These hypotheses will take the form of enhanced or reduced declarative graph or domain knowledge, additional recognition productions, and mechanisms to represent visual salience. A more comprehensive model must also bring ACT-R's subsymbolic mechanisms that govern memory retention, retrieval, and learning processes into play as these no doubt have a significant effect on strategy choice and eye movement patterns (Peebles & Cheng, 2003).

Finally, the current model does not attempt to provide a detailed account of the perceptual processes by which spatial information is encoded or represented during the execution of elementary perceptual tasks. There are currently several attempts to develop mechanisms for spatial representation and processing within cognitive architectures—including ACT-R—however (a number of which are presented in (Gunzelmann, 2011)) and it may be possible for the current functions to be replaced in a future model by ones more conforming with theory and empirical evidence.

Beyond the goal of developing the model to account for the full range of observed behaviour with an increasing number of interaction graph formats, lies the larger aim of constructing a model of comprehension for a broader class of graphs. As discussed earlier, interaction graphs embody a specific set of interpretive rules that are not shared by other graphs. The current model clearly identifies and characterises these rules and distinguishes them from the knowledge and procedures that can be applied to other graphs. It is hoped that in so doing, the model will simplify the task of identifying graph-specific operators and form a basis upon which to develop and explore a range of graph comprehension models for other graphical formats. As it stands however, the model provides a valuable demonstration that the assumptions it currently embodies are sufficient to produce an expert interpretation of the relationships depicted in 2×2 interaction graphs.

Appendix: Model output for the graph in Figure 1

1 seek text at top of display...
2 [m-yield] = [variable]
3 [as] [a] [function] [of] [p-density] = [variable]
4 [and] [n-level] = [variable]
5 seek text at far right of display...
6 [n-level] at [far right] = [independent] variable
7 look to nearest text...
8 [low] = level of [n-level]
9 [high] = level of [n-level]
10 seek objects in plot region...
11 [blue] [line]
12 no memory for [blue] look to legend
13 [blue] [rectangle]. looking for nearest text...[blue] = [low]
14 [green] [rectangle]. looking for nearest text...[green] = [high]
15 seek text at far left of display...
16 [m-yield] at [far left] = [dependent] variable
17 seek text at bottom of display...
18 [p-density] at [bottom] = [independent] variable
19 look to nearest text...
20 [compact] = level of [p-density]. [compact] = [right]
21 [sparse] = level of [p-density]. [sparse] = [left]
22 identify legend levels...
23 [0.0] diff [blue] so [no] [simple] effect [low] [n-level]
24 [0.5] diff [green] so [moderate] [simple] effect [high] [n-level]
25 compare [blue] & [green] levels...
26 [small] diff. [high] [n-level] > [low] [n-level]
27 [small] [main] effect [n-level]
28 identify x axis levels...
29 [0.0] diff [left] so [no] [simple] effect [sparse] [p-density]
30 [0.5] diff [right] so [moderate] [simple] effect [compact] [p-density]
31 compare [left] & [right] levels...
32 [small] diff. [compact] [p-density] > [sparse] [p-density]
33 [small] [main] effect [p-density]
34 compare left and right patterns...
35 [0.5] diff in distance between points. [right] bigger
36 [moderate] diff & [same] point order so [moderate] [interaction]
37 for [sparse] [p-density] [low] [n-level] = [high] [n-level]
38 for [compact] [p-density] [high] [n-level] > [low] [n-level]

References

Ali, N., & Peebles, D. (in press). The effect of Gestalt laws of perceptual organisation on the comprehension of three-variable bar and line graphs. *Human Factors*.

Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York, NY: Oxford University 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.

Casner, S. M. (1991). A task-analytic approach to the automated design of graphic presentations. *ACM Transactions on Graphics*, *10*, 111–151.

Cleveland, W. S., & McGill, R. (1984). Graphical perception: Theory, experimentation, and application to the development of graphical methods. *Journal of the American Statistical Association*, *79*, 531–554.

Gillan, D. J. (1994). A componential model of human interaction with graphs: 1. linear regression modelling. *Human Factors*, *36*(3), 419–440.

Gunzelmann, G. (Ed.). (2011). *Modeling spatial cognition [special issue]* (Vol. 3) (No. 4). Topics in Cognitive Science.

Kosslyn, S. M. (1989). Understanding charts and graphs. *Applied Cognitive Psychology*, *3*, 185–226.

Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). SOAR: An architecture for general intelligence. *Artificial Intelligence*, *33*, 1–64.

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

McClelland, J. L. (2009). The place of modeling in cognitive science. *Topics in Cognitive Science*, *1*(1), 11–38.

Meyer, D. E., & Kieras, D. E. (1997). A computational theory of executive cognitive processes and multiple-task performance 1. basic mechanisms. *Psychological Review*, *104*, 3–65.

Peebles, D., & Ali, N. (2009). Differences in comprehensibility between three-variable bar and line graphs. In *Proceedings of the thirty-first annual conference of the cognitive science society* (pp. 2938–2943). Mahwah, NJ: Lawrence Erlbaum Associates.

Peebles, D., & Ali, N. (in preparation). *Comparing novice and expert comprehension of three-variable bar and line graphs.*

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*, 28–45.

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.

Shah, P., & Freedman, E. G. (2011). Bar and line graph comprehension: An interaction between top-down and bottom-up processes. *Topics in Cognitive Science*, *3*(3), 560–578.

Tabachneck-Schijf, H. J. M., Leonardo, A. M., & Simon, H. A. (1997). CaMeRa: A computational model of multiple representations. *Cognitive Science*, *21*, 305–350.

Wertheimer, M. (1938). Laws of organization in perceptual forms. In W. D. Ellis (Ed.), *A source book of Gestalt psychology.* London: Routledge & Kegan Paul.