

# Modeling Relative Task Effort for Grouped Bar Charts

**Richard Burns (burns@cis.udel.edu)**

Department of Computer and Information Sciences, University of Delaware  
Newark, DE 19716 USA

**Stephanie Elzer (elzer@cs.millersville.edu)**

Department of Computer Science, Millersville University  
Millersville, PA 17551 USA

**Sandra Carberry (carberry@cis.udel.edu)**

Department of Computer and Information Sciences, University of Delaware  
Newark, DE 19716 USA

## Abstract

The overall goal of our research is a system which can recognize the intended message of a grouped bar chart by reasoning about the communicative signals contained in the graphic. One such communicative signal is the relative effort required to perform different perceptual tasks on the graphic. This paper presents our methodology for estimating relative task effort. Based on graph comprehension research and our motivational eye tracking experiments, we hypothesize a set of factors that should be taken into account in a model of task effort. We present our model, implemented in the ACT-R framework, and discuss the results of a final set of eye tracking experiments that validate our model as a predictor of relative task effort.

**Keywords:** Graph Comprehension; Grouped Bar Charts; Cognitive Modeling; Bayesian Reasoning; Eye Tracking.

## Introduction

Information graphics, ranging from simple bar charts and line graphs to grouped bar charts and multiple line graphs, play a prominent role in today's information age. When such graphics appear in popular media, they generally have a communicative goal or message that they are intended to convey. For example, the message conveyed by the top graphic in Figure 1 is ostensibly that female salaries lag behind male salaries in all of the disciplines listed. Our research has shown that when information graphics appear in popular media, their intended message is very often not repeated in the article's text and cannot be gleaned from the graphic's caption (Carberry, Elzer, & Demir, 2006). Thus, recognizing the graphic's message is fundamental to the two applications that we are pursuing: 1) constructing a rich summary of a multimodal document and 2) providing sight-impaired individuals with alternative access to information graphics in popular media via a brief summary of the graphic.

Our previous research produced a Bayesian network that uses the communicative signals present in a simple bar chart as evidence in hypothesizing the graphic's message (Elzer et al., 2005). We are now extending our methodology to grouped bar charts, which are considerably more complex than simple bar charts and present additional challenges. One of the most important communicative signals is the relative effort of different perceptual tasks.<sup>1</sup> For example, consider

<sup>1</sup>Other communicative signals include whether a group is highlighted, the position of a group in the chart, etc.




Figure 1: Two grouped bar charts from the same data.

the two graphics in Figure 1, both of which display the same data. In the top graphic, it is easy to perceptually compare the salaries of males and females in each discipline, and the graphic's intended message is ostensibly that female salaries lag behind male salaries in all of these disciplines. While these comparisons can also be made in the bottom graphic, they are much more difficult due to the design of the graphic. Thus, the bottom graphic appears to convey a different message. This correlates with Larkin and Simon (1987) who observe that informationally equivalent graphics are not necessarily computationally equivalent, and with Peebles and Cheng (2003) who note that seemingly minor design changes can greatly affect a graph viewer's performance on graph reading tasks. Thus perceptual task effort affects the message that is conveyed by a grouped bar chart.

This paper presents our model of relative task effort for grouped bar charts. Note that our goal is not a cognitive model of human graph comprehension, but rather a model that takes into account cognitive and perceptual aspects of the graphic to estimate the relative effort required to perform relevant recognition tasks. Our aim is to develop this model with solid cognitive underpinnings and to validate it with sound experimental data. In future research, we will incorporate the results of this model into a Bayesian network and use this evidence alongside other communicative signal evidence to hypothesize the intended message of a grouped bar chart.

### Related Work

There has been much work on graph comprehension. Classical graph work such as Lohse (1993), Meyer (2000), and Simkin and Hastie (1987) involve the modeling of *specific* fact-retrieval tasks on a graph. Shah, Freedman, and Vekiri (2005) are concerned with how quantitative information is comprehended from graphics. Freedman and Shah (2002) use a constraint-satisfaction approach and consider the graphic viewer’s prior knowledge. Our work differs from these and other efforts in that rather than modeling human processing, our emphasis lies in estimating and *ranking* the relative effort of different message recognition tasks. Our work is similar in that it must take into account both graphic complexity and the human-visual architecture.

### Messages and Task Effort

In this paper, we will consider five general categories of messages that are common in grouped bar charts, and we will model the effort required to perform the task of extracting these messages from a graphic. We hypothesize that the more difficult a message is to extract, the less likely it is that the graphic was intended to convey that message.

- **Same-relation:** a set of entities have the same relationship to one another over an ontology. The top graphic in Figure 1 has an intended message that falls into the *same-relation* category, as does Figure 3 whose intended message is that rural households have less internet access than urban households at all income levels.
- **Contrast-relation:** a set of entities have a different relation to one another at one place in an ontology compared with everywhere else in the ontology. Figure 1 would contain such a *contrast-relation* message if female salaries exceeded males salaries for one of the listed disciplines.
- **Same-trend:** several entities have the same trend (increasing, decreasing, or stable) over some ordinal set (such as years, ages, etc.). Figure 2 is an example of a grouped bar chart whose intended message falls into the *same-trend* category. A *same-trend* message can also be extracted from Figure 3.
- **Contrast-trend:** one entity has a different trend from the other entities. This would be the case if the group labeled




Figure 2: Graphic from the *Philadelphia Inquirer*, “Commuters Facing Fare Hikes”, September 28, 2005.




Figure 3: Graphic from *Business Week*, “A Small Town Reveals America’s Digital Divide”, October 4, 1999.

“*Philadelphia to New York*” in Figure 2 displayed decreasing prices from Sept. to Feb.

- **Gap-trend:** The gap between entities is generally changing in the same direction (increasing or decreasing) over an ontology. Such a *gap-trend* message could be extracted from Figure 3 (although it is effortful to do so) since the gap between internet access for rural and urban households is *generally* increasing across income levels.

We have developed a model for estimating the relative effort involved in extracting each of the above message types for a given grouped bar chart, based on the design of the graphic. In our work on simple bar charts, effort was modeled using a GOMS approach (Card, Moran, & Newell, 1983) which decomposed tasks into perceptual subtasks, such as ‘find the top of a bar’. The effort for primitive subtasks was estimated based on cost estimates by cognitive psychologists, and the effort estimates for the subtasks were summed to provide an effort estimate for the composite task. Although this GOMS approach worked well for modeling relative effort in simple bar charts, grouped bar charts are much more complex and can involve high dimensional relationships over multiple sets of entities, as well as the inherent cognitive limitations in extracting those messages. Thus we chose to use the ACT-

R programmable framework (Anderson, Matessa, & Lebiere, 1997) which is an implemented cognitive theory with visual and declarative modules that are relevant to the perception and memory issues in modeling task effort. However, our goal is *NOT* to construct a cognitive model that simulates how humans comprehend graphs; instead we want to identify the factors that make tasks on one graphic more difficult than on another graphic and utilize them in a model that estimates the relative difficulty of a task on a given graphic.

### Factors that Affect Effort

We performed preliminary (motivational) eye tracking experiments with human subjects to gain insight into the factors that affect task effort in grouped bar charts and to motivate the design of our effort model. Subjects were given a grouped bar chart and asked to perform a recognition task, during which fixations and their durations were measured along with the time to complete the task. Our observations from these experiments, along with previous research by cognitive psychologists and graph designers, suggest factors that should be incorporated into our model of task effort.

**High-level Visual Patterns** The presence of high-level visual patterns that can be easily perceived by the human visual system appears to significantly affect the effort required to extract messages from grouped bar charts. When such patterns were present, we observed fewer bar fixations and shorter time for task completion in our preliminary eye tracking experiments.

Pinker (1990) identified various high-level visual patterns such as linear lines or quadratic curves which are easily identifiable for most viewers. Shah, Mayer, and Hegarty (1999) show that graph viewers use bottom-up visual pattern recognition in graph comprehension and note that the grouping of data points in graph design will influence the perceived pattern recognition of trends. Applying this to identifying the presence of messages in grouped bar charts, we note that attentions on successive bars which are positioned in a relatively straight line are not needed to determine that those bars *are* indeed in a straight line. For example, one can very easily recognize at a high-level that the bar heights in Figure 2 form an increasing pattern. Other patterns, such as a “U” shape in a *same-relation* message with 3 or more bars per group, also appear to be perceived quickly without a fixation on each bar in the pattern.

**Peripheral Vision** The concept of peripheral vision is closely related to the idea of high-level visual patterns in the domain of information graphics and the acceptance that multiple objects can be processed in parallel in a guided search (Anderson & Lebiere, 1998). Salvucci (2001) has shown the ability to attend to an object without fixating on it, in the domains of equation solving, reading, and visual search. We have observed in our eye tracking experiments that subjects do not always fixate on the first and last groups of a graph, but theorize that subjects still attend to them. We

also hypothesize that the use of peripheral vision explains why subjects in our eye tracking experiments did not fixate on every bar when identifying a *same-trend* message in a grouped bar chart.

**Exceptions** We will refer to one or a couple of bars that do not conform to an overall trend, such as the bars in the groups '02 and '03 in Figure 4, as an *exception*. When given a graphic with an overall trend but containing some exceptions, subjects tend to identify the trend, but take longer to do so compared to a similar graph without exceptions. This increased processing time is the general result of more fixations occurring on the graph, especially around the exception area.




Figure 4: A graphic with exceptions. Graphic from *Wall Street Journal*, “Auto Industry, at a Crossroads, Finds Itself Stalled by History”, January 2, 2006.

**Clutter** Wickens and Carswell (1995) showed that performance in comprehension tasks degrades when visual clutter increases, where visual clutter is the close spatial proximity of two perceptually or semantically contrasting elements which should not be compared. The encoding time for one of these elements increases because of the close proximity of the other element’s “noise”. We have observed increased time for recognition tasks on grouped bar charts when visual clutter is present.

**Spatial Reasoning Ability** Trickett and Trafton (2006) theorize that spatial cognition is often used in graph comprehension. They hypothesize that one such spatial task is, for example, the mental averaging of bar heights within a group, for performing the task of comparing the height of two groups. They use the work of Simkin and Hastie (1987) in their modeling, whose *superimposition* elementary graph process involves a graph viewer spatially moving graph objects to create overlap and ease comparison with other graph objects. We hypothesize that superimposition is present in the extraction of some *gap-trend* messages.

### Modeling in ACT-R and its Limitations

The ACT-R framework enables us to address many perceptual and memory issues involved in our recognition tasks. For example, the imaginal module allows us to build a problem representation, and to handle frequent comparisons between an

object of current attention and a representation of a previously attended object. However, other aspects of ACT-R are inadequate for our needs. For example, ACT-R is unable to automatically recognize that the data points representing the tops of a series of bars can be encoded bottom-up to form a single line object, unless that relationship is explicitly declared in the model. Therefore, to implement pattern recognition, we wrote a small script which takes as input a graphic's data points and automatically outputs any relevant high-level visual patterns present in the graphic. These high-level patterns are then available for the ACT-R model to use in modeling recognition tasks on the graphic.

We also incorporated EMMA, an ACT-R add-on (Salvucci, 2001) which was designed to model peripheral vision. EMMA adjusts the constant visual encoding time cost of an object into a variable cost, affected by the proximity of the previous attention location. It is able to successfully model that a "skipped fixation" will occur when two attentions are approximately close and are executed successively.

### Modeling Task Effort

This section briefly discusses the five categories of message recognition tasks and the implementation of our ACT-R model of relative task effort.

#### Same-relation and Contrast-relation Tasks

When the relation being examined is among entities represented by adjacent bars in a group (such as the relation between rural and urban internet access in Figure 3 or the relation among train prices in Sept, Oct, and Feb<sup>2</sup> for the three trips depicted in Figure 2), high-level visual patterns may help with recognizing a same or contrast relation. Therefore, when the tops of the bars in a group represent a high-level visual pattern (such as a "V" or a "U" shape), our ACT-R model encodes this pattern and models the same-relation and contrast-relation tasks as a comparison of the pattern with each group. On the other hand, there may not be a common visual pattern; this happens when the tops of the bars do not represent a common pattern (typically occurring with 4 or more bars per group) or when the bars are so similar in height that a more detailed examination is needed to discern which bar is higher. In such cases, bar-by-bar processing of each group is performed to determine whether there is a same or contrast relation message. When there was a contrast-relation, our motivational eye tracking experiments showed re-attentions in the vicinity of the group that is in contrast with the other groups, presumably to double-check the contrasting relation; thus our model also includes an extra check.

In some cases, such as the bottom half of Figure 1, a *same-relation* message (that female salaries lag behind male salaries in each discipline) can be extracted from the graph but the related bars are not adjacent to one another. For such graphics, our ACT-R model must capture the extraction of the

bars being compared (in the case of Figure 1, the female and male salary bars for each discipline must be extracted from different groups) and perform bar-by-bar comparisons to determine whether they have the same relation to one another — resulting in many fixations and much cognitive processing, and therefore producing a high effort estimate.

#### Same-trend and Contrast-trend Tasks

In modeling the same-trend and contrast-trend recognition tasks, we differentiate between *intra*-trends (where each group in the grouped bar chart represents a trend, as in Figure 2) and *inter*-trends (where each series of bars across all groups represents a trend, such as in Figures 3 and 4).

Intuitively, the recognition of an *intra-trend* requires that one realize that a trend exists in each group of a graph. Therefore, the effort for this realization is dependent on both the cost of recognizing each individual group trend and the number of groups. Thus one would expect that as the number of groups in a graph increases, the effort and total time to realize an *intra-trend* should increase as well. This was borne out in our preliminary eye tracking experiments.

Our motivational eye tracking experiments also showed that the number of bars per group did not significantly affect recognition time. We hypothesize that this can be accounted for by the human ability to recognize high-level visual patterns. Because the bars which form the trend over a group can be encoded as a high-level line pattern, the recognition of the trend at each group can occur without an exhaustive fixation on every bar in the group. Thus our ACT-R model scans group-by-group using high-level group attentions, and compares the current group's visual pattern with a pattern in ACT-R's imaginal module representing the pattern of the preceding group attention. Our experiments suggested that subjects re-attend on and around areas which semantically differ from the surrounding areas and their expectations. Thus in the case of a contrast-relation, our ACT-R model captures this extra effort by re-attending to the prior and succeeding groups around any contrast.

*Inter*-trends consist of a series of bars across all groups (for example, the increasing internet access over income levels for both rural and urban residents in Figure 3). In our eye tracking experiments, the fixations by subjects suggested that they were both processing multiple trends in parallel and processing several adjacent bars in a trend with a single fixation. This apparent use of peripheral vision is accounted for in our model in two ways. When the tops of the bars in a series are all approximately near the tops of the bars in another series, our ACT-R model marks one series as a "free" attention, and attending to the non-free series automatically encodes the trends of both that series and the "free" series. In addition, our model does not fixate on each bar in a series, but instead uses EMMA's ability to capture attending to several close adjacent bars with a single fixation.

However, since *inter*-trends are spread across groups, there is opportunity for visual clutter, particularly when one trend contrasts with the other trends in the graph — for example, a

---

<sup>2</sup>However, the intended message in Figure 2 is a *same-trend* and not a *same-relation* message.

decreasing trend when the other trends are increasing. Thus not only does our ACT-R model re-attend to contrasting *inter*-trends, but it also re-attends to the bars in the area where the contrasting trend crosses the other trends since this represents an area of visual clutter.

For both *intra*-trends and *inter*-trends, our ACT-R model re-attends to the bars around any exceptions to the trend. This captures the extra processing observed in our eye tracking experiments and the resulting extra time to perform same-trend and contrast-trend tasks in the presence of exceptions.

### Gap-trend Task

The *gap-trend* messages that we have observed in our corpus of graphs<sup>3</sup> consist of two bars per group. Based on the work of Trickett and Trafton (2006), if the two series of bars across all the groups are both increasing or both decreasing, we hypothesize that superimposition is used in the gap comparison process. Our ACT-R model takes this into account by fixating on a gap and superimposing it on the succeeding gap to recognize whether the *gap-trend* is generally increasing or decreasing over the graph.<sup>4</sup>

If the two series are moving in opposite directions, the high-level visual pattern formed by the bar tops of the two series can generalize to resemble “<” or “>”, and would be easily recognizable.<sup>5</sup> Thus our ACT-R model extracts the high-level visual patterns for the two series and compares them.

As with trend messages, exceptions can occur and result in re-attentions in the area of the exceptions.

## Validation Experiment

### Design

To evaluate our model for estimating relative task effort, we performed an eye tracking validation experiment using a Tobii T60 system with 20 human subjects.

Each subject was initially presented with learning and practice slides which explained the kinds of tasks that they would be asked to perform and the concept of an *exception* to a trend. Each subject was then presented with 52 graphics, and in each case the subject was asked to perform one of the five tasks: determine whether there is 1) a same-trend, 2) a contrast-trend, 3) a same-relation, 4) a contrast-relation, or 5) a gap-trend. 42 of the graphics contained a trend, relation, or gap-trend that the subject was supposed to identify and the other 10 graphics did not. The graphics included both *intra* and *inter* trends and varied in the number of groups, the number of bars per group, and the presence of exceptions and clutter. The order in which the graphics were presented to

<sup>3</sup>Our corpus consists of approximately 150 grouped bar charts extracted from a variety of popular media sources.

<sup>4</sup>ACT-R/S, an ACT-R module that models spatial representations and their size capacity is currently under development. However, our problem is much simpler than their 3-dimensional problem space (Hiatt, Trafton, & Harrison, 2004).

<sup>5</sup>An “X” pattern can also occur, when the relation between the bars changes — for example, when the first bar is taller than the second bar for the beginning groups (with the gap between them decreasing) and then the second bar becomes taller than the first bar.

subjects was randomized. Both the subjects and the graphics used in this validation experiment were different from those used in the preliminary eye tracking experiments.

For each graphic, an untimed instruction slide was first displayed explaining to the subject which task they were being asked to perform on the next graphic. The subject hit the space bar after comprehending this instruction. Then the graphic was displayed on the screen and the subject performed the requested recognition and then hit the space bar again. In addition to gaze points, the elapsed time following the onset display of the graphic and the hitting of the space bar was also recorded. Then a multiple choice prompt was displayed, and the subject was instructed to select the answer corresponding to his/her response to the task.

### Results and Analysis

For each of the 42 graphics that contained a trend, relation, or gap-trend that the subject was supposed to recognize, we computed the mean completion time for the human subjects<sup>6</sup> and the effort estimate produced by our ACT-R model, where the ACT-R effort estimate was in terms of task completion time. Since our goal is to rank tasks on a graphic in terms of their relative effort, we performed a Spearman Rank-Order Correlation, which is used to determine whether two sets of rank-ordered data are related; values approaching 1.0 show a strong correlation. Table 1 shows the results of the correlations between the model’s effort estimates and average task completion time for the human subjects.

Table 1: The Spearman Rank Correlation Coefficient for the task effort produced by our model and the subjects’ average task completion time. (*n* is the number of graphs.)

Task	Spearman correlation (rho)
Same/Contrast-Relation	$\rho = .854, p < .004 (n = 24)$
Same/Contrast-Trend	$\rho = .821, p < .001 (n = 10)$
Gap-Trend	$\rho = .833, p < .02 (n = 8)$
All Tasks	$\rho = .809, p < .001 (n = 42)$

The results in Table 1 show a strong correlation between the relative task effort produced by our model and the ranking of graphics according to the average task completion time by human subjects; this correlation holds both within the categories of tasks and more importantly, across all tasks. For same/contrast-trend tasks which involved an *inter*-trend and for *same-relation* tasks, the main disagreement in the rankings occurred in graphics where the estimated recognition times by the model and the mean completion times by the subjects differed only slightly. For same/contrast-trend tasks that involved an *intra*-trend, differences in the rankings appear to reflect disagreement about how much extra effort is needed when either a contrasting group or an increase in the number of groups occurs. Our model estimates less time on

<sup>6</sup>There were only a few “incorrect” responses overall, and these were not included in the mean time calculation.

graphics with many groups and no contrasting group than it does for graphics with fewer groups that include a contrasting group, whereas subject completion times reflect the opposite. Overall, the strong correlations displayed in Table 1 validate our model of task effort.

### Future Work

In the next phase of our project, our model will be used to provide evidence about relative task effort for use with other communicative signals in a Bayesian network to hypothesize the intended message of a grouped bar chart. For example, we processed both the top and bottom graphic of Figure 1 in our model of task effort. As hoped, the model estimated that the top graphic required significantly less effort to recognize a *same-relation* message than the bottom graphic (approximately only a third of the time required by the bottom graphic). This result suggests that our ACT-R model will produce relative effort estimates that will serve as useful evidence in our Bayesian network.

### Conclusion

We have developed a model of relative task effort on grouped bar charts, implemented within the ACT-R framework. Our model takes into account a number of factors that appear to impact the requisite effort, including the recognition of high-level visual patterns, the use of peripheral vision, the presence of exceptions and clutter, and the potential for spatial reasoning. Our model was validated by a final set of eye tracking experiments in which a strong correlation was shown between the effort estimates produced by our model and the average completion times of human subjects. Future work on this project involves implementing a Bayesian network which will use the communicative signals present in a grouped bar chart, including relative task effort, to probabilistically reason about the graphic's intended message.

### Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. IIS-0534948.

### References

- Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Anderson, J. R., Matessa, M., & Lebiere, C. (1997). Act-r: A theory of higher level cognition and its relation to visual attention. *Human-Computer Interaction*, 12, 439-462.
- Carberry, S., Elzer, S., & Demir, S. (2006). Information graphics: An untapped resource of digital libraries. In *Proceedings of 9th international acm sigir conference on research and development on information retrieval* (p. 581-588). New York, NY: ACM.
- Card, S., Moran, T. P., & Newell, A. (1983). *The psychology of human-computer interaction*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Elzer, S., Carberry, S., Zukerman, I., Chester, D., Green, N., & Demir, S. (2005). A probabilistic framework for recognizing intention in information graphics. In *Proceedings of the international joint conference on artificial intelligence* (p. 223-230). Morristown, NJ: Association for Computational Linguistics.
- Freedman, E. G., & Shah, P. (2002). Toward a model of knowledge-based graph comprehension. In *Second international conference on diagrammatic representation and inference* (p. 18-30). London, UK: Springer-Verlag.
- Hiatt, L. M., Trafton, J. G., & Harrison, A. (2004). A cognitive model for spatial perspective taking. In *Proceedings of the sixth international conference on cognitive modeling* (p. 354-355). Pittsburgh, PA, USA: Carnegie Mellon University/University of Pittsburgh.
- Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth a thousand words. *Cognitive Science*, 11, 65-99.
- Lohse, G. L. (1993). A cognitive model for understanding graphical perception. *Human-Computer Interaction*, 8, 353-388.
- Meyer, J. (2000). Performance with tables and graphs: effects of training and a visual search model. *Ergonomics*, 43(11), 1840-1865.
- 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 *Artificial intelligence and the future of testing* (p. 73-126). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Salvucci, D. D. (2001). An integrated model of eye movements and visual encoding. *Cognitive Systems Research*, 1, 201-220.
- Shah, P., Freedman, E. G., & Vekiri, I. (2005). The comprehension of quantitative information in graphical displays. In P. Shah & A. Miyake (Eds.), *The cambridge handbook of visuospatial thinking* (p. 426-476). New York, NY: Cambridge University Press.
- Shah, P., Mayer, R. E., & Hegarty, M. (1999). Graphs as aids to knowledge construction: Signaling techniques for guiding the process of graph comprehension. *Educational Psychology*, 91, 690-702.
- Simkin, D., & Hastie, R. (1987). An information-processing analysis of graph perception. *American Statistical Association*, 82, 454-465.
- Trickett, S. B., & Trafton, J. G. (2006). Toward a comprehensive model of graph comprehension: Making the case for spatial cognition. In *Proceedings of the fifth international conference on the theory and application of diagrams* (p. 286-300). Berlin Heidelberg New York: Springer-Verlag.
- Wickens, C. D., & Carswell, C. M. (1995). The proximity compatibility principle: Its psychological foundation and relevance to display design. *Human Factors*, 37, 473-494.