Visual and Spatial Factors in a Bayesian Reasoning Framework for the Recognition of Intended Messages in Grouped Bar Charts

Richard Burns and **Sandra Carberry**
Department of Computer Science
University of Delaware

Stephanie Elzer
Department of Computer Science
Millersville University

Abstract

The overall goal of our research is the automatic recognition of the intended message of a grouped bar chart. This paper presents our preliminary work on a system that utilizes the communicative signals in a grouped bar chart as evidence in a Bayesian network that hypothesizes the primary message conveyed by the graphic. The paper discusses the kinds of communicative signals present in grouped bar charts and an ACT-R model for computationalizing one important communicative signal, the relative effort involved in performing the perceptual tasks necessary for the recognition. It also describes our Bayesian network and its implementation on a subset of the kinds of messages that can be conveyed by grouped bar charts.

Introduction

Grouped bar charts are a type of information graphic often utilized by graphic designers as a tool for visually displaying quantifiable relationships of values that hold over a set of dependent entities in two dimensions. Usually their presence in popular media is to communicate to the graph viewer a high-level contextual message which involves the graphed elements. That is, most information graphics have an intention to communicate. However, the high-level content of an information graphic is often not repeated in the accompanying text or caption of the graphic (Carberry, Elzer, and Demir 2006). Thus, it is necessary to integrate the high-level message conveyed by an information graphic with the article's text in order to completely understand a multimodal document. This paper presents preliminary work on a system for automatically recognizing the high-level intentions from grouped bar chart information graphics. The general framework is based on a methodology already shown to be successful in the domain of simple bar charts (Elzer et al. 2005b).

The system utilizes the graphical features in grouped bar charts as communicative signals and uses them as evidence in a Bayesian network to hypothesize the most likely intended message for a grouped bar chart. The communicative signals used include the salient coloring of bars, as well as their ordering and positioning. Data-dependent evidence can also be viewed as a communicative signal, such as whether entities are salient temporally (for example, a bar entity which represents the current year) or by location.

Additional evidence can be found by estimating the relative perceptual effort of recognizing each possible high-level message. This follows from the AutoBrief project where it was posited that graphic designers will attempt to design a graph in such a way as to facilitate the necessary perceptual tasks for the viewer to recognize the intended communication with as little effort as possible (Green et al. 2004). Thus the relative effort required to recognize a message serves as a signal as to whether that message is what the graph is intended to convey.

At least three applications can greatly benefit from this research. The first is a system which provides sight-impaired individuals with alternative access to information graphics in multimodal articles. Many of these individuals are able to interact with a computer by first utilizing a screen reader, such as JAWS. But information graphics cause great difficulty. For information graphics in popular media, where the reader's primary focus is the understanding of the content of the article, conveying the high-level knowledge captured by an information graphic might be better than merely reciting the raw data points of the graphic. The second is to provide techniques for indexing grouped bar charts in digital libraries. Such information graphics are a rich knowledge resource that is currently ignored in information retrieval. The third is a summarization system for multimodal documents, where the system can use the message inferred for a grouped bar chart to provide a more complete summary of the document's content.

We will first introduce the types of grouped bar chart messages that our system will automatically recognize. Then we introduce the communicative signals utilized to probabilistically recognize these messages. Following that, we will discuss the modeling of our "effort" communicative signal which estimates the relative perceptual effort required for a graph viewer to recognize a message given a graphic. Finally, we will describe our Bayesian structure and its implementation for a subset of the kinds of messages that can be conveyed by grouped bar charts.

Grouped Bar Charts

Grouped bar chart information graphics found in popular media are usually intended to convey some high-level message. For example, Figure 1 is taken from the *Wall Street Journal* and it ostensibly has a message which is to convey the general decrease in profits for car manufacturers Ford and GM from 1998 to 2005.

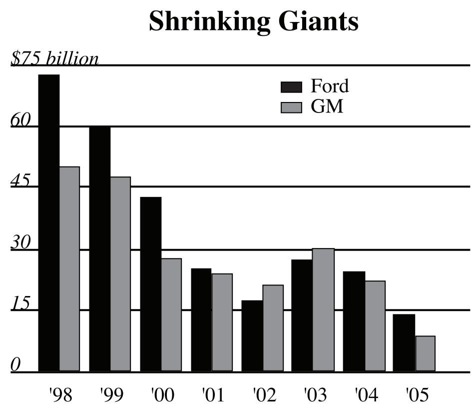

Figure 1: Graphic from *Wall Street Journal*, "Auto Industry, at a Crossroads, Finds Itself Stalled by History", January 2, 2006.

Currently we have collected 150 grouped bar charts from popular media as well as the graphic's accompanying article if the graphic is part of a multimodal document. The grouped bar charts are from a diverse range of sources, including *USA Today*, *Time*, *New Your Times*, *BusinessWeek*, *Philadelphia Inquirer*, *Forbes*, and *The Economist*.

Through a corpus analysis, it was apparent that a vast majority of the grouped bar charts have intended messages that are generalizable into common high-level *message categories*. The following message categories capture the kinds of high-level messages conveyed in the corpus that are currently fully implemented in our system.

1. **Rising-Trends-All**
2. **Falling-Trends-All**

   Many graphics have a high-level message which is generalizable to multiple rising (or falling) trends, where the trends are *generally* rising (falling) for a set of data points over *ordinal* entities. Here, "generally" is intended to include cases where specific data point pairs may exist which are not strictly increasing (decreasing).

   For example, the graphic in Figure 1 can be captured by the message category *Falling-Trends-All*. In addition, notice that the graphic's data does not strictly fall but instead contains a brief increase at the years 2002 and 2003, an "exception" to the general decreasing trend.
3. **Same-Relationship-All**

   The Same-Relationship-All message category captures messages where the relative values of entities in a set is consistent over the sets of entities, given entities that are either ordinal or not ordinal. For example, a *Same-Relationship-All* exists in Figure 2 and the graphic ostensibly conveys the message that food prices are greater in the United States than in Iraq for the displayed grocery groups.

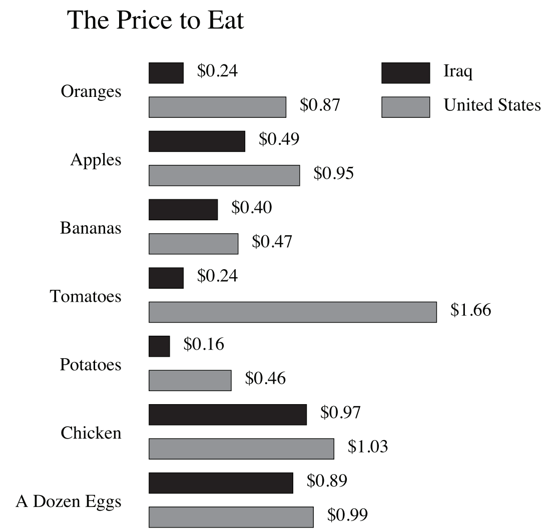

Figure 2: Graphic from *USA Today*, "Markets' prices shelve thrill of new selections", March 10, 2005.

4. **Entity-Relationship-Contrast**

   The Entity-Relationship-Contrast message category captures the high-level intention to convey that the relative ordering of the values for a set of entities differs from the relative ordering of values for the other sets of entities. For example in Figure 3, the increase in instructional time spent on reading is contrasted with the decrease in time devoted to the other subjects.

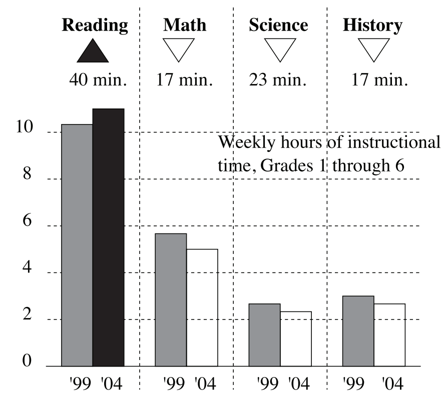

Figure 3: Graphic from *Time*, "How to Fix No Child Left Behind", June 4, 2007.

5. **Contrast-Trend**

   This message category is similar to Entity-Relationship-Contrast, except that the contrast is in the context of a trend instead of the relationships among entities in a set.
6. **Gap-Increasing**

   The Gap-Increasing message category captures the intention to convey that a trend *generally* exists over the absolute size of the "gaps", and that this trend is increasing.

For example in Figure 4, the "gap" between the percentage of rural households compared with the percentage of urban households which have Internet access is generally increasing as the income level of a household increases.

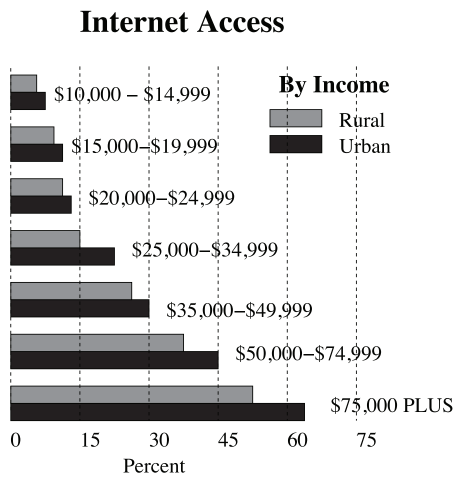

Figure 4: Graphic from *Business Week*, "A Small Town Reveals America's Digital Divide", October 4, 1999.

7. **Gap-Decreasing**
   This message category is identical to Gap-Increasing except that the gap-trend is falling instead of rising.

The graphics in our corpus have been annotated with their intended message by human annotators using consensus-based annotation. Those graphics which were annotated with a message that falls into one of the above discrete message categories were used to train and test our preliminary Bayesian system which will be described later in this paper.

Communicative Signals

Viewing graphics as an extended form of language, graphic designers attempt to assist the graph viewer in recognizing the intended message of a graphic by deliberately incorporating communicative signals in the graphic. We identified many communicative signals present in grouped bar charts through a corpus analysis. They include: highlighted entities, the positioning of entities, the perceptual effort required to recognize a specific message, and other implicit signals such as bar height, and verbs and nouns in the caption of the graphic.

Figure 3 shows a graphic from *Time*, where coloring creates salience. The '04 bar in the first group is colored differently from the '04 bars in the other groups, thereby drawing attention to the increased instruction on reading, in contrast with the decrease in instruction time for the other subjects.

We identified instances where sets of bars were salient because of design choices made by the graphic designer. For example, the position of a group or position of a bar within each group can make a set of bars salient. In Figure 5, the group "Life Sciences" is made salient by its position as the first group in the graphic. "Life Sciences" is listed first and is not part of a natural (such as alphabetical) ordering of the groups.

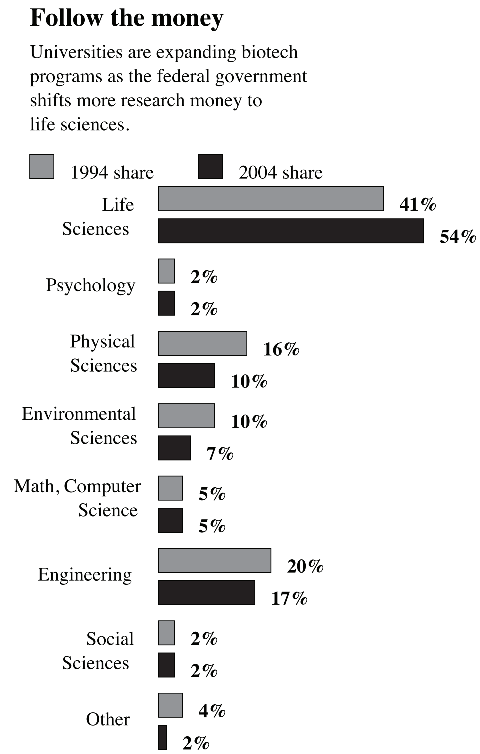

Figure 5: Graphic from *USA Today*, "Universities grid for battle for biosciences supremacy", June 24, 2005.

Kosslyn argues that graphic designers should follow the principle to "*put the most important independent variable on the x-axis, and treat the others as parameters*" (Kosslyn 1994). He claims that the variable on the independent axis will be perceived as more important because of the human nature to visually group elements that are close together. This recommendation, Kosslyn claims, is based on psychological principles of relevance and limited processing capacity. Thus the independent axis entity provides a communicative signal about how the entities in the graphic are intended to be compared.

Captions on graphics in popular media are often very general and uninformative. Even when the caption attempts to convey some of a graphic's message, the caption might be ill-formed or require analogical reasoning and domain knowledge to understand it; thus a general-purpose natural language system would have difficulty processing and understanding captions (Elzer et al. 2005a). Nonetheless, captions do contain communicative signals which can be identified via simple shallow processing (Elzer et al. 2005a). Verbs in a caption, such as the verb *shrink* which is the root form of the adjective *shrinking* that appears in the caption on Figure 1, suggest a general category of message such as *Falling-Trends-All*. Nouns in the caption of a grouped bar chart serve to make an entity salient. For example, the overall caption for the graphic in Figure 6 is "Boys Don't

Cry: Men and Depression". As a result, the entity "men" is salient, and this suggests that the intended message might be a comparison involving men.

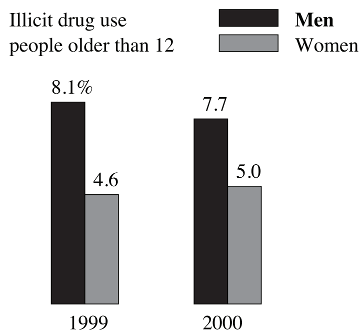

Figure 6: Graphic from *NewsWeek*, "Stop Pretending Nothing's Wrong", June 16, 2003.

The AutoBrief project posited that graphic designers attempt to design a graph in such a way as to facilitate the necessary perceptual tasks for the viewer to recognize the intended communication with as little effort as possible (Green et al. 2004). Our "effort" communicative signal captures this: an estimate of the relative perceptual effort required for a graph viewer to recognize some message given a grouped bar chart. From a recognition viewpoint, this relative effort is a communicative signal about which tasks the graphic designer intended the viewer to perform. If recognizing a message requires significantly more perceptual effort than some other message, then it is less likely that the graphic designer intended that the first message be inferred as the graphic's primary message.

For example in Figure 7, although both graphics contain the same data, individually they convey two different messages. The high-level message conveyed by the top graphic is ostensibly that male salaries are greater than female salaries in all of the subject areas, while the message conveyed by the bottom graphic is ostensibly a message of rank: that engineering and the physical sciences have the greatest salaries for both men and women. While these messages can be inferred from either graphic, the design of the graphic affects how easy it is to actually perform that recognition. Since a viewer expects the graph designer to facilitate the necessary perceptual tasks by making them as easy as possible, perceptual task effort is also a communicative signal. This example also correlates with Larkin and Simon who observe that informationally equivalent graphics are not necessarily computationally equivalent (Larkin and Simon 1987), and Peebles and Cheng who note that seemingly minor design changes can greatly affect performance on graph reading tasks (Peebles and Cheng 2003).

Other communicative signals are also present in grouped bar charts, one being large bar values that are implicit to the data. For example, the group "Life Sciences" in Figure 5 is salient by its relative bar height.

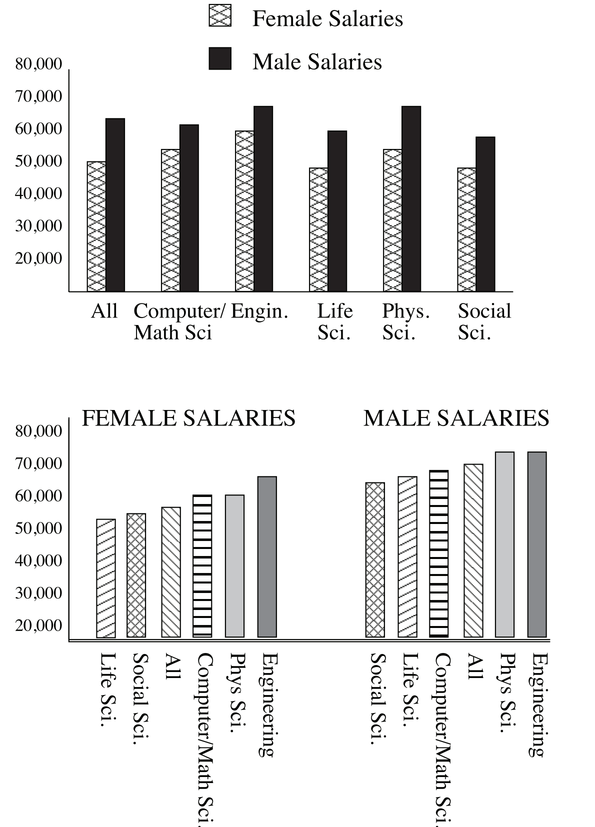

Figure 7: Two computationally inequivalent bar charts from the same data.

Model of Relative Effort

As discussed above, one communicative signal is the relative perceptual effort for a graph viewer to recognize some message. To build a model of relative task effort, we drew on research by cognitive psychologists and performed pilot eye-tracking experiments with human subjects to identify the factors in grouped bar charts which affect the required recognition effort.

High-level Visual Patterns Pinker identified various high-level visual patterns such as linear lines or quadratic curves which are easily identifiable for most viewers (Pinker 1990). Shah hypothesized that graph viewers use bottom-up encoding in graph comprehension and noted that the grouping of data points in graph design will influence the perceived pattern recognition of trends (Shah, Mayer, and Hegarty 1999). This relates to our observations from our pilot studies 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.

We also observed that high-level visual patterns can be easily perceived by the human-visual system, and noticed fewer fixations on high-level visual pattern objects, as well as a shorter task completion time on graphics which contained these patterns. The effort required to extract messages from grouped bar charts is significantly affected when these graphics contain high-level visual patterns. As an example, it is very easy to recognize that the bar heights in Figure 8 form an increasing pattern. Other patterns, such as "U" shapes also appear to be perceived quickly without a fixation on each bar in the pattern.

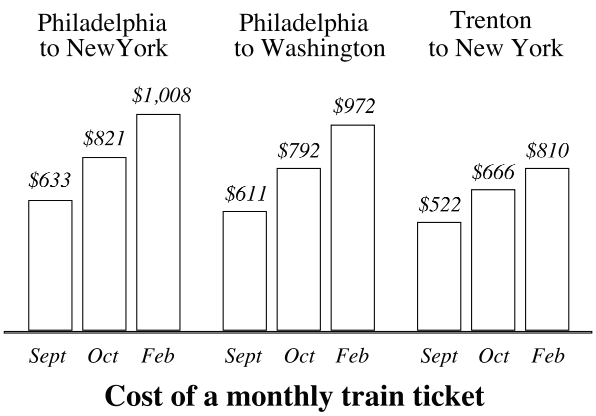

Figure 8: High-level visual patterns allow the trends within each group to be easily perceived. Graphic from *Philadelphia Inquirer*, "Amtrak revives plan for fare increases", September 28, 2005.

Peripheral Vision The concept of peripheral vision is closely related to the idea of high-level visual patterns as well as the ability for multiple objects to be processed in parallel in a guided search (Anderson and Lebiere 1998). Unfortunately, since eye-trackers can only record gazes, it is difficult to observe the exact specifics of how peripheral vision is utilized by a subject. However, we have observed phenomena in our experiments which can be explained by the use of peripheral vision. For instance, subjects did not always fixate on the bars in the first and last groups of a graph. We theorize that subjects are still attending to these groups though, and we also hypothesize that peripheral vision explains why most subjects did not fixate on every bar when recognizing an easily identifiable trend.

Exceptions The presence of exceptions appears to impact the effort required to recognize a trend. In our pilot eye-tracking experiments, when subjects were given a graphic (such as Figure 1) with an overall trend that contained exceptions, subjects tended to still identify the trend, but took longer to do so compared to a similar graph without exceptions. We observed an increased number of fixations around the exception area, causing the increase in processing time.

Clutter Wickens and Carswell showed that performance in comprehension tasks degrades when visual clutter increases (Wickens and Carswell 1995). Visual clutter is the close spatial proximity of two perceptually or semantically contrasting elements which should not be compared. The encoding time for an element increases because of the close proximity of the other element's "noise". We observed increased time for recognition tasks when visual clutter is present in grouped bar charts.

Spatial Reasoning Ability *Superimposition* is described by Simkin as an elementary graph process which allows the graph viewer to spatially move graph objects to create overlap and ease comparison with other graph objects (Simkin and Hastie 1987). Trickett and Trafton hypothesize that superimposition is used in the mental averaging of bar heights within a group when performing the task of comparing the height of two groups (Trickett and Trafton 2006). Based on the nature of the gap, we hypothesize that it is possible for graph viewers to use superimposition during the recognition of a *Gap-Increasing* or *Gap-Decreasing* message, by spatially moving the absolute height difference of two bars, the "gap", onto the gap of the next group. In this way, through a series of superimpositions, it is possible to recognize that the gap is increasing or decreasing across the graph.

Model Based on the psychological literature and the data from our pilot eye-tracking studies, we built a model to estimate this relative effort using the ACT-R programmable cognitive framework(Anderson, Matessa, and 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 for grouped bar charts. Our model is augmented with EMMA (Salvucci 2001), an add-on to ACT-R. EMMA separates fixations from attentions, and accounts for the ability to attend to objects without ever fixating on them. EMMA was incorporated to handle our observation of peripheral vision being used by subjects who could still attend to bar groups without ever fixating on them.

A more natural use of ACT-R would be to cognitively model how humans comprehend graphs, but it is important to emphasize that the purpose of our model is *not* a cognitive model of graph comprehension. Rather, its goal is to estimate the relative difficulty of a task on a given graphic by incorporating the known factors that make recognition tasks on one graphic more difficult than on another graphic.

Currently, a model of relative task effort has been designed and validated for the message categories described earlier: *Rising-Trends-All*, *Falling-Trends-All*, *Same-Relationship-All*, *Entity-Relationship-Contrast*, *Contrast-Trend*, *Gap-Increasing*, and *Gap-Decreasing* over a set of grouped bar charts with a variety of features such as a varied number of bars per group, number of groups, the presence of high-level patters, exceptions, etc. (Burns, Elzer, and Carberry 2009).

Bayesian Network

Our system utilizes a Bayesian network to automatically hypothesize a grouped bar chart's most likely intended message. The Bayesian network is given an XML representation of a grouped bar chart which specifies the height of the bars in each group, the bar and group labels, coloring of bars, any annotated values, the graph's caption, etc. Communicative signals, such as whether a group of bars is colored differently from the other groups, is extracted from the XML representation and passed to the Bayesian network as evidence. In addition, the ACT-R effort model is used to estimate the relative perceptual effort required for each potential message.

As discussed earlier, human annotators identified the intended message of each graphic, along with the communicative signals present in the graphic. These annotated graphs were used to compute the conditional probability tables that capture the causal dependencies between the messages and the communicative signals. Thus the network is able to represent that *given some intended high-level message for a graphic, which communicative signals are likely to be present or absent in that graphic?* By applying Bayes' rule, our network can thus reason in the opposite direction: *given the communicative signals in a graphic, what is the likelihood of some high-level message being the actual intended message for the graphic?*

Network Structure The intended message of the graphic is captured in a top-level node called *IntendedMessage*. The possible *general* message categories that can be inferred for a graphic are children of this top-level node. (Each message category has a node dedicated entirely to its own message rather than as a state in the top-level node, to keep the conditional probability tables smaller.)

The next level below the *general* message category nodes contains *specific* instantiations for each node — that is, the message category along with instantiations of its parameters. For example, a message in the *Entity-Relationship-Contrast* category must specify which set of entities are being contrasted with the other entities in the graphic.

As noted earlier, verbs in a caption can suggest a general category of message. For example, a verb from a graphic's caption whose stem is "increase" would give supporting evidence for the Bayesian network to possibly hypothesize the *Rising-Trends-All* and *Gap-Increasing* message categories for the graphic. Thus the presence or absence of such verb evidence is attached to the top-level intended message node. A more detailed description of how verbs are classified into discrete classes is not provided due to space limitations.

Other communicative signals, such as whether a group is colored differently from other groups or the relative perceptual effort required to recognize a specific message, provide evidence for/against particular instantiated messages. Thus the presence or absence of these communicative signals is captured in evidence nodes attached to each instantiated message node in the Bayesian network.

The top-level of the Bayesian network and the next level down, moving from general message categories to specific message instantiations, is shown in Figure 9 and Figure 10, respectively.

Training and Validation From our grouped bar chart corpus, 41 out of the 150 graphics were annotated with *Rising-Trends-All*, *Falling-Trends-All*, *Contrast-Trend*, *Same-Relationship-All*, *Entity-Relationship-Contrast*, *Gap-Increasing*, or *Gap-Decreasing* messages. The network was tested using leave-one-out-cross-validation on this subset, and successfully recognized the graphic's annotated intention for 30 of the 41 graphics—a very promising success rate for our initial system. The graphics that it failed on appear to heavily utilize domain and world knowledge in their messages.

Figure 11 is one such graphic that the Bayesian network failed on. The network hypothesized that the graphic's most likely intended message was *Same-Relationship-All, across groups*: that U.S. advertisement spending was greatest in Network TV, then Cable TV, and then U.S. Internet for each of the years 2003, 2004, and 2005. However, the graphic was annotated as a *Contrast-Trend(Group 1)*: that there is a relationship between the decrease in advertisement spending in Network TV with the increase in spending to Cable TV and the U.S. Internet. The caption of the graphic, "New Media?", helps signal a comparison of the newer forms of communication (Cable TV and Internet) against the more established Network TV entity. However, recognizing *"New Media"* as such a communicative signal requires domain knowledge about new and old forms of media, something that has not yet been incorporated into our system — and is a very difficult problem.

Related Work

Elzer built a simple bar chart intention recognition system which achieved an accuracy rate of 79.1% on a corpus of 110 popular media graphics (Elzer et al. 2005b). It heavily utilized several psychological insights and references including Lohse's framework (Lohse 1993) which is based on the GOMS paradigm. It also used a Bayesian network structure to model observed communicative signals. However, our work is significantly different from the simple bar chart system by: (1) considering the more complex multi-dimensional messages that can be conveyed in grouped bar charts, (2) identifying the unique communicative signals present in grouped bar charts, and (3) estimating task effort by using the ACT-R modeling framework instead of the GOMS paradigm which appears to be insufficient for the high-level messages in complex graphics.

Kerpedjiev et. al. proposes a methodology for automatically generating graphics that realize desired intentions (Kerpedjiev et al. 1998). The system uses hierarchical planning to transform high-level communication goals into subgoals which are then translated into conceptual graph tasks. These tasks describe operations that users can perform on the graph to extract information. The system utilizes a task model which selects specific graphical subgoals given an intention. The goal is a generated graphic whose features allow the graphic viewer to recognize the desired communication with as little effort as possible. The Postgraphe system generates graphics based on the input of a communicative intention and a data set (Fasciano and Lapalme 2000). PostGraphe uses a schema-based planning mechanism to distribute the communicative goals between generated graphics and text. The resulting output is a multimodal document. However, these efforts are in the area of automatically generating graphics that capture a desired communication goal. This is the opposite direction of our work, where our goal is automatic intention recognition given a graphic.

Mittal in the SAGE system, implements a process which automatically generates captions which can be used to explain data in novel or creative graphics (Mittal et al. 1998).

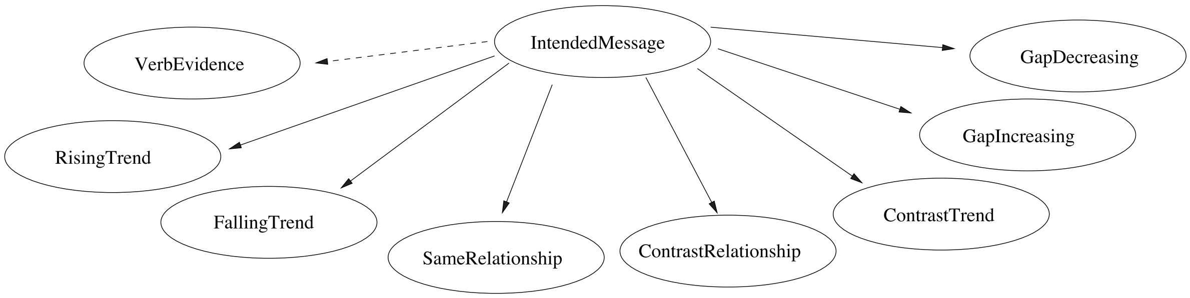

Figure 9: Top level of the Bayesian network.

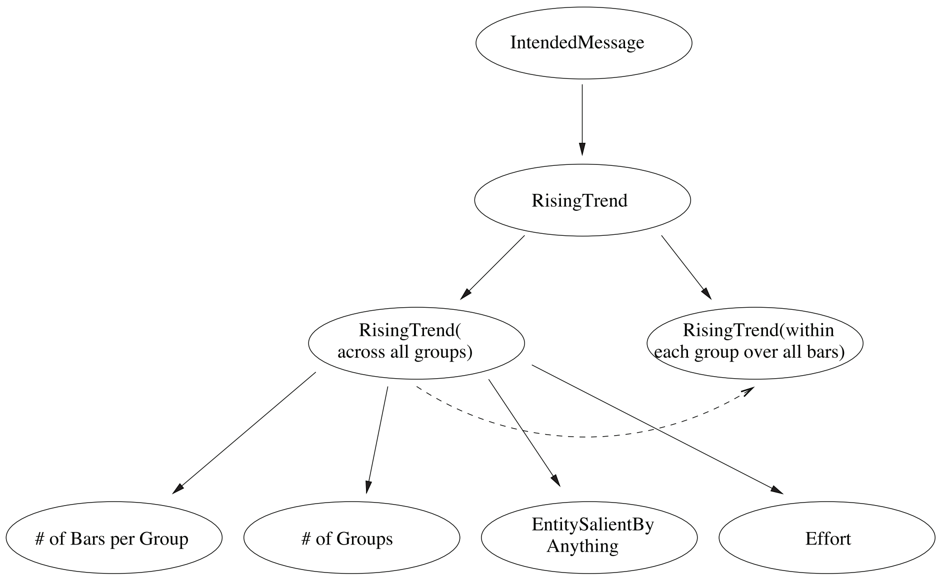

Figure 10: Descendants of the *RisingTrend* high-level message category node.

The system takes as input a set of data and a communicative goal to be communicated, and outputs a 2-dimensional graphic with a corresponding caption which helps explain the relations of data entities in the graphic. Although the concept of generating good captions bears some similarity to identifying the intended message of a graphic, Mittal is given the data points that will be displayed and the communicative goal of the graphic. In our work, the communicative goal must be inferred by reasoning about the communicative signals entered into the graphic by the graph designer.

Conclusion and Future Work

This paper has presented our preliminary system for automatically hypothesizing the intended message of a grouped bar chart by probabilistically reasoning about the communicative signals in the graphic. The paper discusses the communicative signals that are present in grouped bar charts, including the relative perceptual effort required to recognize a message and our ACT-R model for estimating this relative task effort. The paper describes the structure of our Bayesian network for hypothesizing the intended message of a grouped bar chart and our preliminary results on graphics whose annotated messages are members of a subset of the possible message categories. Our preliminary system is only implemented for this subset of messages.

Our current research is focused on extending our ACT-R model of relative task effort and our Bayesian network to the other message types. In addition, we are collecting additional grouped bar charts in order to provide a larger training set which we hope will improve our recognition accuracy. Nevertheless, our preliminary results indicate that our methodology is very promising.

Acknowledgements

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

References

Anderson, J. R., and Lebiere, C. 1998. *The Atomic Components of Thought*. Mahwah, NJ: Lawrence Erlbaum As-

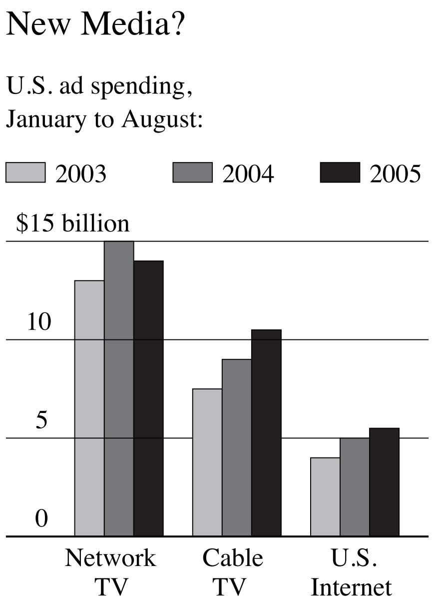

Figure 11: The Bayesian network failed on this graphic because of domain knowledge.

sociates.

Anderson, J. R.; Matessa, M.; and Lebiere, C. 1997. Act-r: A theory of higher level cognition and its relation to visual attenion. *Human-Computer Interaction* 12:439–462.

Burns, R.; Elzer, S.; and Carberry, S. 2009. Modeling relative task effort for grouped bar charts. In Taatgen, N., and van Rijn, H., eds., *Proceedings of the 31st Annual Conference of the Cognitive Science Society*, 2292–2297. Austin, TX: Cognitive Science Society.

Carberry, S.; Elzer, S.; and 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*, 581–588. New York, NY: ACM.

Elzer, S.; Carberry, S.; Chester, D.; Demir, S.; Green, N.; Zukerman, I.; and Trnka, K. 2005a. Exploring and exploiting the limited utility of captions in recognizing intention in information graphics. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics*, 223–230.

Elzer, S.; Carberry, S.; Zukerman, I.; Chester, D.; Green, N.; and Demir, S. 2005b. A probabilistic framework for recognizing intention in information graphics. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 223–230. Morristown, NJ: Association for Computational Linguistics.

Fasciano, M., and Lapalme, G. 2000. Intentions in the coordinated generation of graphics and text from tabular data. *Knowledge and Information Systems* 2(3):310–339.

Green, N. L.; Carenini, G.; Kerpedjiev, S.; Mattis, J.; Moore, J. D.; and Roth, S. F. 2004. Autobrief: an experimental system for the automatic generation of briefings in integrated text and information graphics. *International Journal of Human-Computer Studies* 61(1):32–70.

Kerpedjiev, S.; Green, N.; Moore, J.; and Roth, S. 1998. Saying it in graphics: from intentions to visualizations. In *In Proceedings of the Symposium on Information Visualization (InfoVis '98). IEEE Computer Society Technical Committee on Computer Graphics*, 97–101. IEEE.

Kosslyn, S. M. 1994. *Elements of Graph Design*. W. H. Freeman and Company.

Larkin, J. H., and 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.

Mittal, V. O.; Carenini, G.; Moore, J. D.; and Roth, S. 1998. Describing complex charts in natural language: A caption generation system. *Computational Linguistics* 24(3):431–467.

Peebles, D., and 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*, 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.; Mayer, R. E.; and 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., and Hastie, R. 1987. An information-processing analysis of graph perception. *American Statistical Association* 82:454–465.

Trickett, S. B., and 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*, 286–300. Berlin Heidelberg New York: Springer-Verlag.

Wickens, C. D., and Carswell, C. M. 1995. The proximity compatibility principle: Its psychological foundation and relevance to display design. *Human Factors* 37:473–494.