Modeling a Graph Viewer's Effort in Recognizing Messages Conveyed by Grouped Bar Charts

Richard Burns[1], Sandra Carberry[2], and Stephanie Elzer Schwartz[3]

[1] Dept. of Computer Science, West Chester University, West Chester, PA 19383 USA
rburns@wcupa.edu
[2] Dept. of Computer Science, University of Delaware, Newark, DE 19716 USA
carberry@cis.udel.edu
[3] Dept. of Computer Science, Millersville University, Millserville, PA 17551 USA
stephanie.schwartz@millersville.edu

Abstract. Information graphics (bar charts, line graphs, etc.) in popular media generally have a high-level message that they are intended to convey. These messages are seldom repeated in the document's text yet contribute to understanding the overall document. The relative perceptual effort required to recognize a particular message is a communicative signal that serves as a clue about whether that message is the one intended by the graph designer. This paper presents a model of relative effort by a viewer for recognizing different messages from grouped bar charts. The model is implemented within the ACT-R cognitive framework and has been validated by human subjects experiments. We also present a statistical analysis of the contribution of effort in recognizing the intended message of a grouped bar chart.

1 Introduction

Information graphics are non-pictorial visual devices, such as simple bar charts, line graphs, pie charts, and grouped bar charts. They are incorporated into a multimodal document in order to achieve one or more communicative goals [12,11]. In the case of scientific documents, the communicative goal might be to present data or to help the reader visualize information. However, when information graphics appear in popular media such as periodicals (*USA Today*, *Wall Street Journal*) and magazines (*The Economist*, *Time*), they generally have a high-level message that they are intended to convey. For example, consider the graphics in Figures 1 and 2 which ostensibly convey that *"Women are more likely than men to delay medical treatment"* and that *"food prices are lower in Iraq than in the United States"*. Although the caption in Figure 1 explicitly states the graphic's message, the caption in Figure 2 does not help recognize the message of that graphic. A study by Carberry et al. [5] found that a graphic's message is often not contained in the graphic's caption or in the article accompanying the graphic. Yet the graphic's message is integral to understanding the full content of a multimodal document.

We are developing systems for recognizing the intended message of an information graphic in popular media. Our work has several applications. The first

S. Carberry et al. (Eds.): UMAP 2013, LNCS 7899, pp. 114–126, 2013.

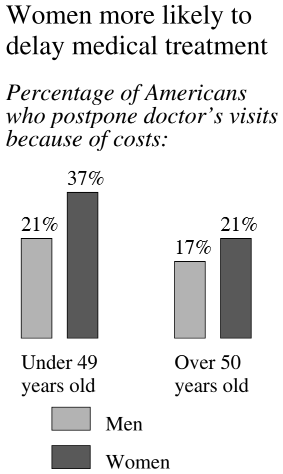

Fig. 1. "Snapshot" graphic from *USA Today*, June 16, 2003

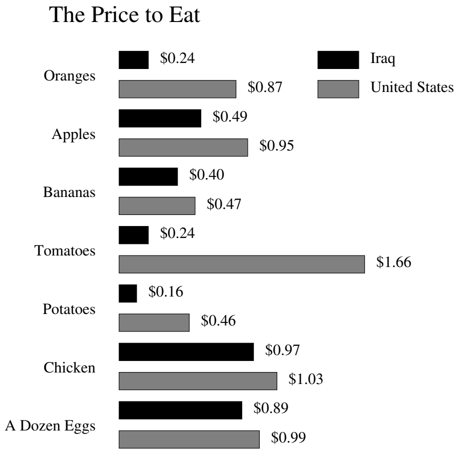

Fig. 2. Graphic from *USA Today*, "Markets' prices shelve thrill of new selections", March 10, 2005

is a system that provides alternative access to information graphics for sight-impaired individuals by conveying their high-level content to the user via speech [8]. The second is the retrieval of information graphics from a digital library where the graphic's message is used to capture its high-level content [5]. The third application is the summarization of multimodal documents that takes into account their information graphics rather than ignoring them or merely considering only their captions [21].

Previous work has focused on message recognition for simple bar charts [9] and line graphs [22]. Grouped bar charts are another type of information graphic. Although grouped bar charts and simple bar charts both display bars that depict quantifiable relationships among the values of entities, grouped bar charts also contain a *grouping* dimension. For example, Figures 1, 2, and 5 respectively contain two groups of two bars each, seven groups of two bars each, and three groups of four bars each. Consequently, grouped bar charts convey a much wider variety of messages than simple bar charts or line graphs, and thus the recognition of their messages is much more complex.

The overall objective of our research is a system for recognizing the high-level messages conveyed by a grouped bar chart [4]. An important component of the system is a model that estimates the relative effort that a graph viewer would have to expend in order to recognize a particular message for a given grouped bar chart. Consider the graphic in Figure 4 which depicts the same data as is displayed in Figure 3. Although the graphic in Figure 3 facilitates an easy comparison of Internet usage between the United States and China for each year from 2002 to 2008, such a comparison is more difficult in Figure 4 due to the different organization. Thus while the primary message conveyed by the graphic

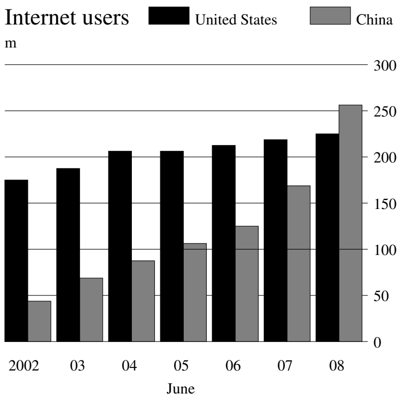

Fig. 3. Graphic from *The Economist Daily Chart*, July 31, 2008

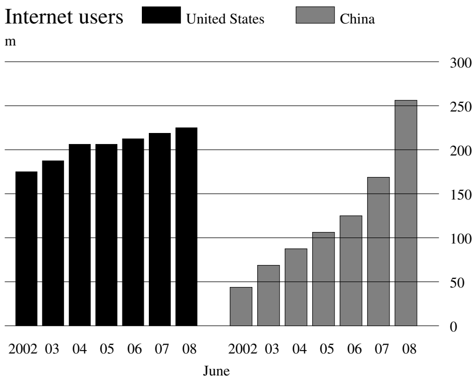

Fig. 4. A different organization of Figure 3

in Figure 4 is the rising trend in Internet usage in both countries[1], the primary message of the graphic in Figure 3 is that the gap in Internet usage between the two countries has decreased and in fact reversed (with China now having higher usage than the United States). This correlates with an observation by Larkin and Simon [13] that graphics that are informationally equivalent (that is, they convey the same data) are not necessarily computationally equivalent (it can be more difficult to extract certain information from one graphic than from the other). The AutoBrief project [11] hypothesized that graph designers construct graphics that enable the easy performance of tasks that are needed to recognize the graphic's communicative goal. Thus we view the relative effort needed to recognize a particular message from a graphic as evidence of whether that was the message intended by the graph designer — that is, the more effort required to recognize a particular message relative to other messages, the less likely that was the message that the graph designer intended to convey.

This paper presents our model of the relative effort that is required for a viewer to recognize messages from grouped bar charts and its effect on our overall system. Our model is implemented in the ACT-R cognitive framework [2] and is based on research in the area of graph comprehension as well as our own motivating eye-tracking experiments. Validation experiments quantitatively and qualitatively support our model.

Section 2 of this paper discusses related work. Section 3 of the paper describes the types of messages that grouped bar chart information graphics convey in popular media. Section 4 then presents our model that estimates the relative effort required for a user to recognize a particular message given a graphic. It presents the cognitive research underlying the model, describes its implementation, and presents the results of an experiment validating the model. Section 5

[1] And perhaps that it is rising faster in China.

very briefly describes how the model, along with other communicative signals, is incorporated into a Bayesian message recognition system.

2 Related Work

Elzer [9] and Wu [22] have implemented intended message recognition systems for simple bar charts and line graphs, respectively. Their systems are similar to our grouped bar chart system in that they also use Bayesian networks to probabilistically capture the relationships between high-level intended messages and communicative signals. However, grouped bar charts are more complex than simple bar charts and they convey a much richer set of messages. The system for simple bar charts modeled relative effort [10], but it followed the GOMS paradigm [6] in which perceptual tasks were decomposed into primitive tasks whose effort estimates were summed. However, grouped bar charts require more complex reasoning that also takes into account peripheral vision, high-level visual patterns, the re-encoding of graph objects, and other aspects of perceptual processing that were not considered in the effort model for simple bar charts. Consequently, our effort model for grouped bar charts is implemented within the ACT-R cognitive modeling framework [2] which facilitates such complex reasoning.

3 Messages Conveyed in Grouped Bar Charts

We collected 330 grouped bar charts from a variety of popular media sources and assembled them into a corpus that is available online.[2] In analyzing the corpus, we identified 25 different *message categories* that capture the kinds of messages conveyed by grouped bar charts [4]. Parameters in the message categories become instantiated to fully capture the intended *messages.* Each graph in the collected corpus was examined by a team of annotators who identified the graphic's high-level *primary* message and *secondary* message[3], based on our generalization of message categories, terminology, and parameters.[4] Consensus-based annotation [3] was performed to resolve cases of disagreement to enable the inclusion of difficult examples where the message was not obvious and there was not complete agreement amongst the annotators [14]. The final consensus for the intended messages in the corpus is also published online.[5] In this section, we briefly present some of the most commonly occurring message categories.[6]

Trend Messages. Trend messages convey a general trend that is either rising, falling, or steady. The trend exists over a set of ordinal data points. Trend messages can be within-groups in which case each group of bars comprises a graph

[2] Accessible at: `http://www.cis.udel.edu/~burns/corpus`
[3] A second intended message that is not as apparent.
[4] The full terminology is presented in [4].
[5] Available at: `http://www.cis.udel.edu/~burns/corpus/view-consensus.php`
[6] Space limitations preclude the description of all identified message categories.

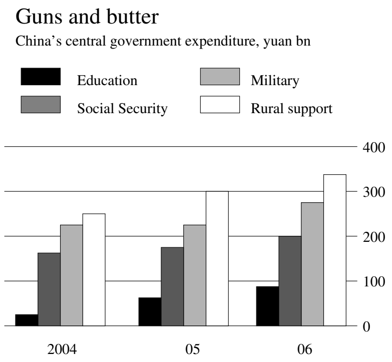

Fig. 5. Graphic from *The Economist*, "Planning the new socialist countryside", March 9, 2006

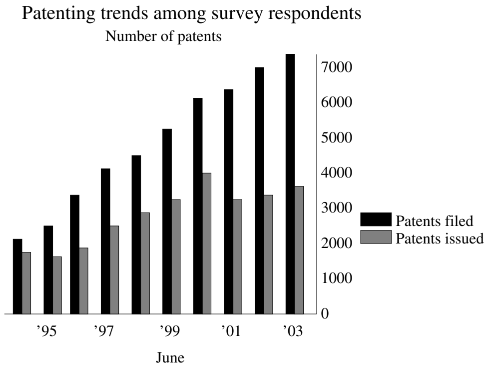

Fig. 6. Graphic from *Technology Review*, "A Mixed Bag of U.S. Institutions", July 2005

group entity or across-groups in which case the i^{th} bar from each group forms a graph series entity. For example, the grouped bar chart in Figure 5 ostensibly conveys the primary message that *"China increased spending on education, social security, military, and rural support from 2004 to 2006"*; it is an *across-groups* message since the i^{th} bar from each of the three groups comprises the i^{th} trend. We generalize this and similar messages into the *Rising-Trend* message category.

Relationship Messages. Relationship messages capture the consistency of relative values for a set of graphed entities, or the inconsistency of one set of relative values with respect to the other sets. For example, the graphic in Figure 7 ostensibly conveys the message that *"the increased funding to the area of Life Sciences is in contrast to the steady or decreased funding to the other areas"*. This messages contrasts *Life Sciences* with the other entities, and the comparison with respect to research funding is within-groups. We identify it as an *Entity-Relationship-Contrast* message category. Messages that convey the *identical* relative ordering of values of a set of graphed entities (that is, there is no contrasting entity) are generalized into the *Same-Relationship-All* message category. The *Opposite-Relationship* message category captures messages that convey two entities with a different relative ordering of bar values.

Gap Messages. The gap message category captures high-level messages involving either one *gap*, or a trend in the size of multiple *gaps*, where a *gap* is the approximate absolute difference between two values within the same entity. Gap messages can refer to gaps within-groups or to gaps across-groups. There are several interesting types of gap messages that occur in grouped bar charts.

Figure 6 displays a graphic whose message falls into the *Gap-Increasing* message category, where the graph is intended to convey that the trending of the gaps (gaps within-groups) is increasing. Ostensibly, the graph conveys that the *"gap between the number of patents filed and issued increased over the period from 1994 to 2003"*.

The *Gap-Crossover* message category captures messages conveying that the trending of one entity surpasses the trending of another entity. For example, the grouped bar chart in Figure 3 conveys that *"the gap between the number of Internet users in the US and China has steadily decreased until now China has more Internet users than the US"*.

Comparison Messages. Some grouped bar chart messages compare either the gap of a single entity to the gaps of the other entities, or the entire single entity itself to all of the other entities. These message categories are called *Gap-Comparison* and *Entity-Comparison*, respectively.

Consider the grouped bar chart in Figure 8. Its *primary* intended message is *"that the percentage of pirated software is greater in China than in the World"*. However, to a much lesser degree, the graphic *secondarily* conveys *"that the decrease in piracy between 1994 and 2002 is less in China than in the World"*.The former message captures the comparison of the *size* of piracy in China with the other entity (the World) and is an *Entity-Comparison* message, whereas the latter represents a comparison of the *gap* between piracy in 1994 and 2002 for China with the other entity (the World) and is a *Gap-Comparison* message.

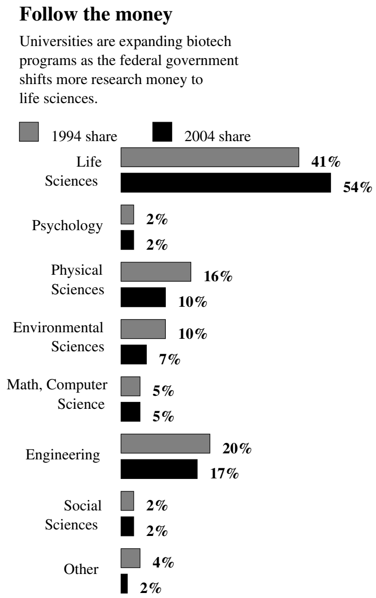

Fig. 7. Graphic from *USA Today*, "Universities gird for battle for biosciences supremacy", June 24, 2005

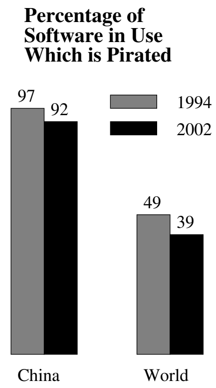

Fig. 8. Graphic from *NewsWeek*, "Microsoft Cozies Up to China", June 28, 2004

4 Effort

From a given graphic, one can extract a number of different messages, such as a trend within groups, a relationship across groups, a comparison of gaps among group entities, etc. Green et al. [11] hypothesizes that the design of a graph should facilitate as much as possible the tasks that the graph viewer will need to perform in order to understand the graphic's intended message. Thus, because our motivation is an overall intention recognition system that can hypothesize the messages that are most likely to be the ones intended by the graph designer, the ability to model which messages in a graphic are relatively easy to recognize in comparison with which messages are more difficult to recognize, may be a useful factor for our overall system. This is the motivation for our *relative effort model* that we now present.

4.1 Cognitive Underpinnings

Our implemented cognitive model was motivated by research by cognitive psychologists as well as motivational eye-tracking experiments that we performed. The following briefly summarizes these principles and observations.

Pinker [15] states that high-level visual patterns, such as straight lines and "U" shapes, are easily recognized by the human visual system. Shah et al. [17] notes how graph comprehension utilizes *bottom-up* visual pattern recognition for the perception of trends (fixating on adjacent bars to determine whether the direction is generally increasing, decreasing or steady). In our motivational experiments, we also observed that subjects were able to quickly identify relationships conveyed by adjacent bars whose values represented a straight line, more so than bar patterns which did not capture a familiar perceptual pattern.

Peripheral vision is the ability to visually process objects that are not in direct line-of-sight. For guided search tasks, Anderson [1] showed how multiple objects can be processed in parallel through the use of peripheral vision. In our motivational eye-tracking experiments, we also observed how subjects processed entities in a graph by using peripheral vision — that is, all of the entities in a graph were processed without fixating on each individual entity. For example, we frequently observed instances in which subjects could correctly identify trends in grouped bar charts without looking at every bar.

We define an *exception* as one or more bars that do not conform to an overall trend. From our motivational experiments, we found that exceptions do impact the overall effort required to recognize a trend. When presented with noisier graphs with a greater amount of "trend exceptions", subjects frequently reattended to areas around the exception location, and overall, took longer to perform high-level recognition tasks on the graphic.

We observed that the presence of *visual clutter* and violations of the *proximity compatibility principle* (as defined by Wickens and Carswell [20], that perceptual proximity of elements is advised if the elements are part of the same task and is otherwise discouraged) cause an increase in processing time for subjects.

Simkin and Hastie [18] describe *superimposition* as an elementary spatial reasoning graph process where the graph viewer spatially moves objects in the graph around to ease comparison with other graph objects. Trickett and Trafton [19] additionally hypothesize that superimposition is used for the mental averaging within a group for performing the task of comparing the heights of groups in grouped bar charts.

4.2 ACT-R Model of Effort

We implemented our model of effort in the ACT-R cognitive modeling framework [2] with the EMMA add-on [16]. Models implemented in ACT-R are intended to reflect the ACT-R theory of human cognition. Model accuracy is usually demonstrated by comparing the model on some task to that of a human performing the same task. However, it is important to emphasize that our goal is not to construct a cognitive model that simulates how humans comprehend graphs, but rather to create a model that estimates the *relative* difficulty for a user to recognize one message vs. another in the same graphic.

ACT-R models how the cognitive system *uses* visual attention, but it is unable to automatically recognize that the data points representing the tops of bars can be encoded to form a visual pattern, unless that relationship is explicitly declared. Therefore to implement pattern recognition, a small preprocessing script was also implemented. High-level patterns that the script identifies are then declared in model to simulate *top-down* encoding (recognizing the direction of a trend with only a few fixations because the bars generally form a straight line, or quickly recognizing that several groups of bars each form a common visual pattern such as a "U" shape and thus convey similar relationships).

We implemented 12 different cognitive submodels in our overall model of relative effort for grouped bar charts. Some submodels estimate the relative effort for multiple messages categories. For example, the same cognitive submodel can process both *Rising-Trends* and *Falling-Trends* that exist *across-groups.* The following presents the significant parts of the submodels that estimate the relative difficulty for a user to recognize the message categories presented in Section 3.[7]

Trends (Within-Groups) Model. This model estimates the relative perceptual effort required for the recognition of trends *within-groups.* The model attends to and encodes each group until all groups have been processed. The total processing time for the model is dependent on the cost of encoding each group as well as the number of groups in the graph. The increase in cost as a result of additional groups was significant in the motivational eye-tracking experiments.

High-level visual patterns may exist in a group and are identified through the preprocessing. Exceptions are also possible. As expected, because of high-level visual pattern recognition ability and peripheral vision, the motivational eye-tracking experiments also showed that the number of bars per group did not significantly affect recognition time. The model ultimately encodes each group top-down or bottom-up into a trend representation when a trend exists.

Relationship (Within-Groups) Model. The *Relationship* model is very similar to the *Trend* model. Each group is encoded until all of the group entities are processed. Entities with contrasting relationships are re-encoded. Any visual patterns are identified in preprocessing.

Gap Trend (Gaps Within-Groups) Model. The design of this model follows observations from the motivational eye-tracking data that high-level visual patterns are utilized in the recognition of the *Gap-Increasing (gaps within-groups)* messages and that additional fixations tend to occur around the *"crossover*

[7] Space precludes us from describing all of the submodels in our system.

point" in a *Gap-Crossover (gaps within-groups)* message.[8] The model repeatedly alternatives between attending to each series, to simulate the encoding of gaps between adjacent groups. In addition, high-level visual patterns across groups are processed by the model which allows some bar entities to be encoded without an explicit attention and ultimately speeds up the overall processing time. Crossover points are identified by the model as instances of visual clutter that induce additional attentions.

Entity-Comparison (Group Entity Instantiation) Model. Unlike the previous models, this model also requires an instantiated parameter: a specific group entity to compare with the other groups. Thus, the model's estimate of effort is dependent on the instantiation.

The *Entity-Comparison* message category sometimes captures a message of *rank*, such as *"the instantiated entity is the 2nd tallest group"*. The model processes the graphic by beginning with an instantiated entity and repeatedly finding the next tallest entity in the grouped bar chart until no more exist. Thus, the model attends to a subset of entities in the graphic and compares the instantiated entity with all of the entities in that subset. Because it is easier to recognize the rank of an entity in a grouped bar chart if the entities are sorted by bar height, preprocessing in the model determines if the entities in the grouped bar chart are sorted by bar height. If they are, the model will recognize the rank of the instantiated entity more quickly by following in a straight path all of the entities that are taller than it.

Gap-Comparison (Gaps Within-Groups Instantiation) Model. This model requires an instantiation of a gap that exists within a group entity. Intuitively, the recognition of the size of a gap is dependent on its similarity to the size of the gaps in the other group entities; thus, it is important which entity is instantiated. The model first encodes the gap of the instantiated entity. Then the model encodes all of the other gaps in the grouped bar chart while re-attending to any whose gap size is approximately similar to that of the instantiated entity.

4.3 Validation Experiment

Design. We validated our model by comparing the relative effort estimates for a given set of grouped bar charts against the relative effort required by human subjects performing the same tasks on the same set of graphics.[9] 46 human subjects participated in the experiment, each performing graph tasks on 48 grouped bar charts. Each subject was initially presented with learning and practice slides that explained the types of tasks that they would be asked to perform. Then the appropriate task was prompted to the subjects before each graph in the actual experiment. As an example, a prompt for the *Trend* graph task was: *"In the following graph, is each country's revenue generally increasing? are all revenues*

[8] In Figure 3, the "crossover point" is between the *07* and *08* groups.
[9] Graphs and subjects were different than in the motivational experiments.

increasing except for one country? except for 2 or more countries? or do all of the revenues first increase and then decrease?"

Quantitative Results. For each grouped bar chart, the average mean timing for a subject to perform the prompt task was calculated.[10] These means were ranked to produce an ordered set. The times estimated by the model for the same set of grouped bar charts and graph tasks were also ranked.

The Spearman rank-order correlation measures the relation between two sets of rank-ordered data. Values approaching 1.0 indicate a strong correlation between the ranking of two sets. The overall Spearman correlation for the ranking of all 48 tasks and graphics is $\rho = .725$, $p < .001$. This strong correlation suggests that the models capture the relative effort for recognizing different messages from a graphic and thus should serve as a useful piece of evidence in our overall intention recognition system (Section 5).

Qualitative Results. Additionally, the subject data from the validation experiment was qualitatively consistent with many of the intuitions that were incorporated into the design of the models. For example:

– subjects recognized trends within-groups with less effort when there were fewer groups and more visual patterns
– additional bars per group increased the effort for within-group relationship comparisons
– relationship comparisons within-groups were generally less effortful than across-groups
– additional groups in a graph increased the effort required for recognizing the group with the largest gap unless that largest gap was extremely salient
– a group was more easily identified in entity comparisons when its bar entities were each taller than the bars comprising the other groups
– subjects recognized gap trends where the gap was within groups much easier than when the graph was designed with the gaps across groups

5 Role of Effort in Message Recognition

Our overall system that automatically recognizes the intended message of a graphic is implemented as a Bayesian network graphical model [4]. Given a grouped bar chart, a computer vision system [7] first processes the graphic and extracts its features: the positioning of bars, their bar heights, etc. These features are passed to the effort models and Bayesian network.

Various pieces of *communicative evidence* are input into the overall system so that the Bayesian network can hypothesize the most likely intended message of a graph. One major piece of evidence is the relative effort required to recognize a message. For each possible message that might be recognized from a graphic, effort is discretized into three categories: *Easy*, *Medium*, and *Hard*, capturing how relatively easy or difficult it would be for a viewer to recognize that message from

[10] Incorrect responses by subjects were omitted.

Table 1. Impact of Evidence in the Bayesian System

Included Evidence in Overall System	Accuracy	McNemar's Test
Baseline: None	98 / 330 (29.7%)	$\chi^2 = 15.803$, $p < .0001$
Effort Only	153 / 330 (46.4%)	

the graphic. Many other types of communicative evidence are also incorporated into the system, such as if a group entity is much taller than the others. Using leave-one-out cross-validation, the overall system accuracy for recognizing the primary intended message of a grouped bar chart is 65.6%, which far exceeds a baseline accuracy of 18.8% that results from selecting the most commonly occurring possible message.

It is interesting to consider the impact of effort on our Bayesian recognition system. As noted earlier, it is common for grouped bar charts to have both a primary and a secondary message. Our annotators also annotated our corpus for secondary messages and found that 177 of 330 grouped bar charts had *strong* secondary messages. These secondary messages were only identified by the annotators when they were quickly apparent and recognizable with minimal effort.

We hypothesize that effort is an important factor for the recognition of messages. Communicative signals other than effort (coloring, salience by height or position, salience by mention in a caption, etc.) contribute to the recognition of a graphic's primary message, and the absence of one kind of communicative signal can be compensated for by the presence of other communicative signals. On the other hand, these other communicative signals may detract from the recognition of a secondary message that relies mostly on being readily apparent. Thus to see if our effort model has a positive impact on our recognition of messages, we ran an experiment that considers both a graph's primary and secondary message (if any). We first ran our Bayesian system without any evidence nodes to establish a baseline, and then ran the system once again with only effort as evidence. When no evidence was considered (only the a priori probabilities of messages are present), the system's baseline for correctly predicting either the primary or secondary message of a graphic within the top two messages that it hypothesizes is 29.7%. When we add only effort evidence into the system, this performance improves to 46.4% — demonstrating that the learned probabilistic relationships between intended messages and relative effort is a beneficial evidence source for the overall system. These results are shown in Table 1 along with a statistical significance measurement as calculated by McNemar's test, which is used on nominal, matched-pair data to show the statistical significance of change.

6 Conclusion

Prior work has modeled the relative difficulty for a user to recognize primary messages in simple bar charts. However, grouped bar charts convey a far richer set of messages—including secondary messages—that require a richer model of relative effort. This paper has presented our model of relative effort for grouped bar charts, including the cognitive underpinnings of the model and its validation. It also briefly explored the benefit of a model of relative effort as an evidence source in our overall intention recognition system that aims to automatically identify both primary and secondary messages in grouped bar charts.

References

1. Anderson, J.R., Lebiere, C.: The Atomic Components of Thought. Lawrence Erlbaum Associates, Mahwah (1998)
2. Anderson, J.R., Matessa, M., Lebiere, C.: Act-r: A theory of higher level cognition and its relation to visual attenion. Human-Computer Interaction 12, 439–462 (1997)
3. Ang, J., Dhillon, R., Krupski, A., Shriberg, E., Stolcke, A.: Prosody-based automatic detection of annoyance and frustration in human-computer dialog. In: Proceedings of the International Conference on Spoken Language Processing, pp. 2037–2040 (2002)
4. Burns, R.: Automated Intention Recognition of Grouped Bar Charts in Multimodal Documents. Ph.D. thesis, University of Delaware (2012)
5. Carberry, S., Elzer, S., Demir, S.: Information graphics: An untapped resource of digital libraries. In: Proceedings of 9th International ACM SigIR Conference on Research and Development on Information Retrieval, pp. 581–588. ACM, New York (2006)
6. Card, S., Moran, T.P., Newell, A.: The Psychology of Human-Computer Interaction. Lawrence Erlbaum Associates, Hillsdale (1983)
7. Chester, D., Elzer, S.: Getting computers to see information graphics so users do not have to. In: Hacid, M.-S., Murray, N.V., Raś, Z.W., Tsumoto, S. (eds.) ISMIS 2005. LNCS (LNAI), vol. 3488, pp. 660–668. Springer, Heidelberg (2005)
8. Demir, S., Oliver, D., Schwartz, E., Elzer, S., Carberry, S., McCoy, K.F.: Interactive sight into information graphics. In: Proceedings of the 2010 International Cross Disciplinary Conference on Web Accessibility (W4A), pp. 16:1–16:10. ACM, New York (2010)
9. Elzer, S., Carberry, S., Zukerman, I.: The automated understanding of simple bar charts. Artificial Intelligence 175(2), 526–555 (2011)
10. Elzer, S., Green, N., Carberry, S., Hoffman, J.: A model of perceptual task effort for bar charts and its role in recognizing intention. International Journal on User Modeling and User-Adapted Interaction 16, 1–30 (2006)
11. Green, N.L., Carenini, G., Kerpedjiev, S., Mattis, J., Moore, J.D., Roth, S.F.: 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 (2004)
12. Iverson, G., Gergen, M.: Statistics: The Conceptual Approach. Springer, New York (1997)
13. Larkin, J.H., Simon, H.A.: Why a diagram is (sometimes) worth a thousand words. Cognitive Science 11, 65–99 (1987)
14. Litman, D.J., Forbes-Riley, K.: Predicting student emotions in computer-human tutoring dialogues. In: Proceedings of the 42nd Annual Meeting on Association for Computational Linguistics, pp. 352–359. Association for Computational Linguistics, Morristown (2004)

15. Pinker, S.: A theory of graph comprehension. No. 15 in Occasional Papers, Center for Cognitive Science, Massachusetts Institute of Technology, Cambridge, MA (1981)
16. Salvucci, D.D.: An integrated model of eye movements and visual encoding. Cognitive Systems Research 1, 201–220 (2001)
17. Shah, P., Mayer, R.E., Hegarty, M.: Graphs as aids to knowledge construction: Signaling techniques for guiding the process of graph comprehension. Educational Psychology 91, 690–702 (1999)
18. Simkin, D., Hastie, R.: An information-processing analysis of graph perception. American Statistical Association 82, 454–465 (1987)
19. Trickett, S.B., Trafton, J.G.: Toward a comprehensive model of graph comprehension: Making the case for spatial cognition. In: Barker-Plummer, D., Cox, R., Swoboda, N. (eds.) Diagrams 2006. LNCS (LNAI), vol. 4045, pp. 286–300. Springer, Heidelberg (2006)
20. Wickens, C.D., Carswell, C.M.: The proximity compatibility principle: Its psychological foundation and relevance to display design. Human Factors 37, 473–494 (1995)
21. Wu, P., Carberry, S.: Toward extractive summarization of multimodal documents. In: Proceedings of the Workshop on Text Summarization at the Canadian Conference on Artificial Intelligence, pp. 53–61 (2011)
22. Wu, P., Carberry, S., Elzer, S., Chester, D.: Recognizing the intended message of line graphs. In: Goel, A.K., Jamnik, M., Narayanan, N.H. (eds.) Diagrams 2010. LNCS (LNAI), vol. 6170, pp. 220–234. Springer, Heidelberg (2010)