Abstractive Summarization of Line Graphs from Popular Media

Charles F. Greenbacker **Peng Wu**
Sandra Carberry **Kathleen F. McCoy** **Stephanie Elzer***
Department of Computer and Information Sciences
University of Delaware, Newark, Delaware, USA
[charlieg|pwu|carberry|mccoy]@cis.udel.edu
*Department of Computer Science
Millersville University, Millersville, Pennsylvania, USA
elzer@cs.millersville.edu

Abstract

Information graphics (bar charts, line graphs, etc.) in popular media generally have a discourse goal that contributes to achieving the communicative intent of a multimodal document. This paper presents our work on abstractive summarization of line graphs. Our methodology involves hypothesizing the intended message of a line graph and using it as the core of a summary of the graphic. This core is then augmented with salient propositions that elaborate on the intended message.

1 Introduction

Summarization research has focused primarily on summarizing textual documents, and until recently, other kinds of communicative vehicles have been largely ignored. As noted by Clark (1996), language is more than just words — it is any signal that is intended to convey a message. Information graphics (non-pictorial graphics such as bar charts, line graphs, etc.) in popular media such as *Newsweek*, *Businessweek*, or newspapers, generally have a communicative goal or intended message. For example, the graphic in Figure 1 is intended to convey a changing trend in sea levels — relatively flat from 1900 to 1930 and then rising from 1930 to 2003. Thus, using Clark's view of language, information graphics are a means of communication.

Research has shown that the content of information graphics in popular media is usually not repeated in the text of the accompanying article (Carberry et al., 2006). The captions of such graphics are also often uninformative or convey little of the graphic's high-level message (Elzer et al., 2005). This contrasts with scientific documents in which graphics are often used to visualize data, with explicit references to the graphic being used to explain their content (e.g., "*As shown in Fig. A...*"). Information graphics in popular media contribute to the overall communicative goal of a multimodal document and should not be ignored.

Our work is concerned with the summarization of information graphics from popular media. Such summaries have several major applications: 1) they can be integrated with the summary of a multimodal document's text, thereby producing a richer summary of the overall document's content; 2) they can be stored in a digital library along with the graphic itself and used to retrieve appropriate graphics in response to user queries; and 3) for individuals with sight impairments, they can be used along with a screen reader to convey not only the text of a document, but also the content of the document's graphics. In this paper we present our work on summarizing line graphs. This builds on our previous efforts into summarizing bar charts (Demir et al., 2008; Elzer et al., 2011); however, line graphs have different messages and communicative signals than bar charts and their continuous nature requires different processing. In addition, a very different set of visual features must be taken into account in deciding the importance of including a proposition in a summary.

2 Methodology

Most summarization research has focused on extractive techniques by which segments of text are extracted and put together to form the summary.

Proceedings of the Workshop on Automatic Summarization for Different Genres, Media, and Languages, pages 41–48,
Portland, Oregon, June 23, 2011.

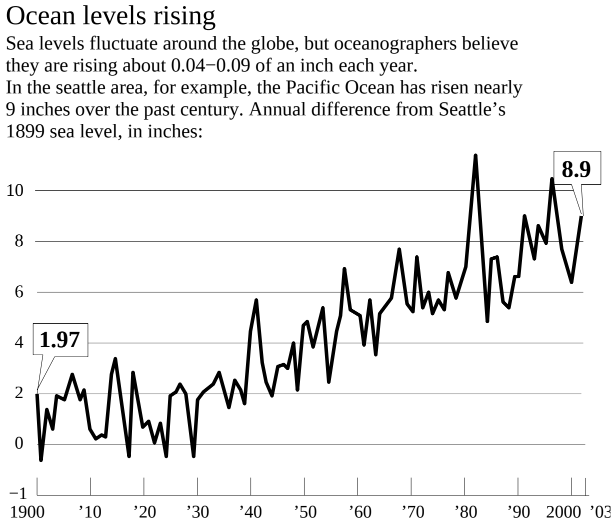

Figure 1: From "Worry flows from Arctic ice to tropical waters" in USA Today, May 31, 2006.

However, the *Holy Grail* of summarization work is abstractive summarization in which the document's content is understood and the important concepts are integrated into a coherent summary. For information graphics, extractive summarization might mean treating the text in the graphic (e.g., the caption) as if it were document text. One could imagine perhaps expanding this view to include selecting particular data points or segments and constructing sentences that convey them. Abstractive summarization, on the other hand, requires that the high-level content of the graphic be identified and conveyed in the summary. The goal of our work is abstractive summarization. The main issues are identifying the knowledge conveyed by a graphic, selecting the concepts that should be conveyed in a summary, and integrating them into coherent natural language sentences.

As noted in the Introduction, information graphics in popular media generally have a high-level message that they are intended to convey. This message constitutes the primary communicative or discourse goal (Grosz and Sidner, 1986) of the graphic and captures its main contribution to the overall discourse goal of the entire document. However, the graphic also includes salient features that are important components of the graphic's content. For example, the graphic in Figure 1 is very jagged with sharp fluctuations, indicating that short-term changes have been inconsistent. Since the graphic's intended message represents its primary discourse goal, we contend that this message should form the core or focus of the graphic's summary. The salient features should be used to augment the summary of the graph and elaborate on its intended message. Thus, our methodology consists of the following steps: 1) hypothesize the graphic's primary discourse or communicative goal (i.e., its intended message), 2) identify additional propositions that are salient in the graphic, and 3) construct a natural language summary that integrates the intended message and the additional salient propositions into a coherent text.

Section 3 presents our methodology for hypothesizing a line graph's intended message or discourse goal. It starts with an XML representation of the graphic that specifies the x-y coordinates of the sampled pixels along the data series in the line graph, the axes with tick marks and labels, the caption, etc.; constructing the XML representation is the responsibility of a Visual Extraction Module similar to the one for bar charts described by Chester and Elzer (2005). Section 4 presents our work on identifying the additional propositions that elaborate on the intended message and should be included in the summary. Section 5 discusses future work on realizing the propositions in a natural language summary, and Section 6 reviews related work in multimodal and abstractive summarization.

3 Identifying a Line Graph's Message

Research has shown that human subjects have a strong tendency to use line graphs to portray trend relationships, as well as a strong tendency to describe line graphs in terms of trends (Zacks and Tversky, 1999). We analyzed a corpus of simple line graphs collected from various popular media including *USA Today*, *Businessweek*, and *The (Wilmington) News Journal*, and identified a set of 10 high-level message categories that capture the kinds of messages that are conveyed by a simple line graph. Table 1 defines four of them. The complete list can be found in (Wu et al., 2010b). Each of these messages requires recognizing the visual trend(s) in the depicted data. We use a support vector machine (SVM) to first segment the line graph into a sequence of visually-distinguishable trends; this sequence is then input into a Bayesian network that reasons with evidence from the graphic

Intention Category	Description
RT: Rising-trend	There is a rising trend from <param_1> to <param_2>.
CT: Change-trend	There is a <direction_2> trend from <param_2> to <param_3> that is significantly different from the <direction_1> trend from <param_1> to <param_2>.
CTR: Change-trend-return	There is a <direction_1> trend from <param_3> to <param_4> that is different from the <direction_2> trend between <param_2> and <param_3> and reflects a return to the kind of <direction_1> trend from <param_1> to <param_2>.
BJ: Big-jump	There was a very significant sudden jump in value between <param_1> and <param_2> which may or may not be sustained.

Table 1: Four categories of High Level Messages for Line Graphs

in order to recognize the graphic's intended message. The next two subsections outline these steps. (Our corpus of line graphs can be found at www.cis.udel.edu/∼carberry/Graphs/viewallgraphs.php)

3.1 Segmenting a Line Graph

A line graph can consist of many short, jagged line segments, although a viewer of the graphic abstracts from it a sequence of visually-distinguishable trends. For example, the line graph in Figure 1 consists of two trends: a relatively stable trend from 1900 to 1930 and a longer, increasing trend from 1930 to 2003. Our Graph Segmentation Module (GSM) takes a top-down approach (Keogh et al., 2001) to generalize the line graph into sequences of rising, falling, and stable segments, where a segment is a series of connected data points. The GSM starts with the entire line graph as a single segment and uses a learned model to recursively decide whether each segment should be split into two subsegments; if the decision is to split, the division is made at the point being the greatest distance from a straight line between the two end points of the original segment. This process is repeated on each subsegment until no further splits are identified. The GSM returns a sequence of straight lines representing a linear regression of the points in each subsegment, where each straight line is presumed to capture a visually-distinguishable trend in the original graphic.

We used Sequential Minimal Optimization (Platt, 1999) in training an SVM to make segment splitting decisions. We chose to use an SVM because it works well with high-dimensional data and a relatively small training set, and lessens the chance of overfitting by using the maximum margin separating hyperplane which minimizes the worst-case generalization errors (Tan et al., 2005). 18 attributes, falling into two categories, were used in building the data model (Wu et al., 2010a). The first category captures statistical tests computed from the sampled data points in the XML representation of the graphic; these tests estimate how different the segment is from a linear regression (i.e., a straight line). The second category of attributes captures global features of the graphic. For example, one such attribute relates the segment size to the size of the entire graphic, based on the hypothesis that segments comprising more of the total graph may be stronger candidates for splitting than segments that comprise only a small portion of the graph.

Our Graph Segmentation Module was trained on a set of 649 instances that required a split/no-split decision. Using leave-one-out cross validation, in which one instance is used for testing and the other 648 instances are used for training, our model achieved an overall accuracy rate of 88.29%.

3.2 A Bayesian Recognition System

Once the line graph has been converted into a sequence of visually-distinguishable trends, a Bayesian network is built that captures the possible intended messages for the graphic and the evidence for or against each message. We adopted a Bayesian network because it weighs different pieces of evidence and assigns a probability to each candidate intended message. The next subsections briefly outline the Bayesian network and its evaluation; details can be found in (Wu et al., 2010b).

Structure of the Bayesian Network Figure 2 shows a portion of the Bayesian network constructed for Figure 1. The top-level node in our Bayesian network represents all of the high-level message cat-

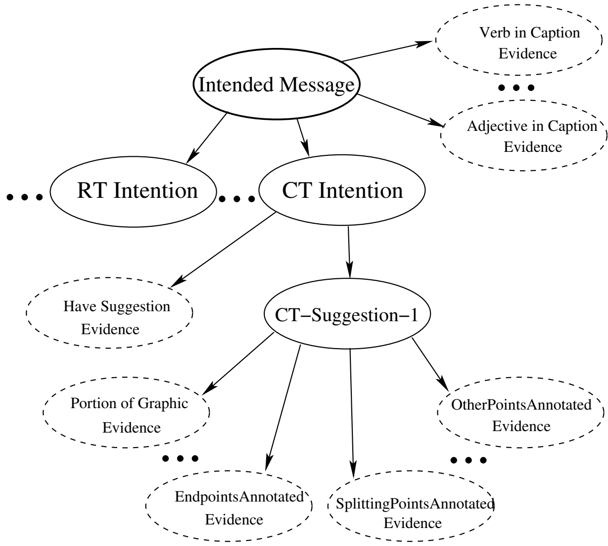

Figure 2: A portion of the Bayesian network

egories. Each of these possible non-parameterized message categories is repeated as a child of the top-level node; this is purely for ease of representation. Up to this point, the Bayesian network is a static structure with conditional probability tables capturing the a priori probability of each category of intended message. When given a line graph to analyze, an extension of this network is built dynamically according to the particulars of the graph itself. Candidate (concrete) intended messages, having actual instantiated parameters, appear beneath the high-level message category nodes. These candidates are introduced by a Suggestion Generation Module; it dynamically constructs all possible intended messages with concrete parameters using the visually-distinguishable trends (rising, falling, or stable) identified by the Graph Segmentation Module. For example, for each visually-distinguishable trend, a Rising, Falling, or Stable trend message is suggested; similary, for each sequence of two visually-distinguishable trends, a Change-trend message is suggested. For the graphic in Figure 1, six candidate messages will be generated, including RT(1930, 2003), CT(1900, stable, 1930, rise, 2003) and BJ(1930, 2003) (see Table 1).

Entering Evidence into the Bayesian Network Just as listeners use evidence to identify the intended meaning of a speaker's utterance, so also must a viewer use evidence to recognize a graphic's intended message. The evidence for or against each of the candidate intended messages must be entered into the Bayesian network. We identified three kinds of evidence that are used in line graphs: attention-getting devices explicitly added by the graphic designer (e.g., the annotation of a point with its value), aspects of a graphic that are perceptually-salient (e.g., the slope of a segment), and clues that suggest the general message category (e.g., a verb [or noun derived from a verb such as *rebound*] in the caption which might indicate a Change-trend message). The first two kinds of evidence are attached to the Bayesian network as children of each candidate message node, such as the child nodes of "CT-Suggestion-1" in Figure 2. The third kind of evidence is attached to the top level node as child nodes named "Verb in Caption Evidence" and "Adjective in Caption Evidence" in Figure 2.

Bayesian Network Inference We evaluated the performance of our system for recognizing a line graph's intended message on a corpus of 215 line graphs using leave-one-out cross validation in which one graph is held out as a test graph and the conditional probability tables for the Bayesian network are computed from the other 214 graphs. Our system recognized the correct intended message with the correct parameters for 157 line graphs, resulting in a 73.36% overall accuracy rate.

4 Identifying Elaborative Propositions

Once the intended message has been determined, the next step is to identify additional important informational propositions[1] conveyed by the line graph which should be included in the summary. To accomplish this, we collected data to determine what kinds of propositions in what situations were deemed most important by human subjects, and developed rules designed to make similar assessments based on the graphic's intended message and visual features present in the graphic.

4.1 Collecting Data from Human Subjects

Participants in our study were given 23 different line graphs. With each graph, the subjects were provided

[1]We define a "proposition" as a logical representation describing a relationship between one or more concepts, while a "sentence" is a surface form realizing one or more propositions.

Figure 3: From "This Cable Outfit Is Getting Tuned In" in *Businessweek* magazine, Oct 4, 1999.

with an initial sentence describing the overall intended message of the graphic. The subjects were asked to add additional sentences so that the completed summary captured the most important information conveyed by the graphic. The graphs were presented to the subjects in different orders, and the subjects completed as many graphs as they wanted during the one hour study session. The set covered the eight most prevalent of our intended message categories and a variety of visual features. Roughly half of the graphs were real-world examples from the corpus used to train the Bayesian network in Section 3.2, (e.g., Figure 3), with the others created specifically to fill a gap in the coverage of intended messages and visual features.

We collected a total of 998 summaries written by 69 human subjects for the 23 different line graphs. The number of summaries we received per graph ranged from 37 to 50. Most of the summaries were between one and four sentences long, in addition to the initial sentence (capturing the graphic's intended message) that was provided for each graph. A representative sample summary collected for the line graph shown in Figure 3 is as follows, with the initial sentence provided to the study participants in italics:

> *This line graph shows a big jump in Blonder Tongue Laboratories stock price in August '99.* The graph has many peaks and valleys between March 26th 1999 to August '99 but maintains an average stock price of around 6 dollars. However, in August '99 the stock price jumps sharply to around 10 dollars before dropping quickly to around 9 dollars by September 21st.

4.2 Extracting & Weighting Propositions

The data collected during the study was analyzed by a human annotator who manually coded the propositions that appeared in each individual summary in order to determine, for each graphic, which propositions were used and how often. For example, the set of propositions coded in the sample summary from Section 4.1 were:

- $volatile(26Mar99, Aug99)$
- $average\_val(26Mar99, Aug99, \$6)$
- $jump\_1(Aug99, \$10)$
- $steep(jump\_1)$
- $decrease\_1(Aug99, \$10, 21Sep99, \$9)$
- $steep(decrease\_1)$

From this information, we formulated a set of rules governing the use of each proposition according to the intended message category and various visual features. Our intuition was that by finding and exploiting a correlation between the intended message category and/or certain visual features and the propositions appearing most often in the human-written summaries, our system could use these indicators to determine which propositions are most salient in new graphs. Our rules assign a weight to each proposition in the situation captured by the rule; these weights are based on the relative frequency of the proposition being used in summaries reflecting similar situations in our corpus study. The rules are organized into three types:

1. Message Category-only (M):
   IF $M = m$ **THEN** select P with weight w_1
2. Visual Feature-only (V):
   IF $V = v$ **THEN** select P with weight w_2
3. Message Category + Visual Feature:
   IF $M = m$ and $V = v$
   THEN select P with weight w_2

We constructed type 1 (Message Category-only) rules when a plurality of human-written summaries in our corpus for all line graphs belonging to a given message category contain the proposition. A weight was assigned according to the frequency with which the proposition was included. This weighting, shown in Equation 1, is based on the proportion of summaries for each line graph in the corpus having intended message *m* and containing proposition *P*.

$$w_1 = \prod_{i=1}^{n} \frac{P_i}{S_i} \tag{1}$$

In this equation, n is the number of line graphs in this intended message category, S_i is the total number of summaries for a particular line graph with this intended message category, and P_i is the number of these summaries that contain the proposition.

Intuitively, a proposition appearing in all summaries for all graphs in a given message category will have a weight of 1.0, while a proposition which never appears will have a weight of zero. However, a proposition appearing in all summaries for half of the graphs in a category, and rarely for the other half of the graphs in that category, will have a much lower weight than one which appears in half of the summaries for all the graphs in that category, even though the overall frequencies could be equal for both. In this case, the message category is an insufficient signal, and it is likely that the former proposition is more highly correlated to some particular visual feature than to the message category.

Weights for type 2 and type 3 rules (Visual Feature-only and Message Category + Visual Feature) are slightly more complicated in that they involve a measure of degree for the associated visual feature rather than simply its presence. The definition of this measure varies depending on the nature of the visual feature (e.g., steepness of a trend line, volatility), but all such measures range from zero to one. Additionally, since the impact of a visual feature is a matter of degree, the weighting cannot rely on a simple proportion of summaries containing the proposition as in type 1 rules. Instead, it is necessary to find the covariance between the magnitude of the visual feature ($|v|$) and how frequently the corresponding proposition is used ($\frac{P}{S}$) in the corpus summaries for the n graphs having this visual feature, as shown in Equation 2.

$$Cov(|v|, \frac{P}{S}) = \left[\left(\frac{\sum_{i=1}^{n} |v_i|}{n} \frac{\sum_{i=1}^{n} \frac{P_i}{S_i}}{n} \right) - \frac{\sum_{i=1}^{n} |v_i| \frac{P_i}{S_i}}{n} \right] \tag{2}$$

Then for a particular graphic whose magnitude for this feature is $|\overline{v}|$, we compute the weight w_2 for the proposition P as shown in Equation 3.

$$w_2 = |\overline{v}| * Cov(|v|, \frac{P}{S}) \tag{3}$$

This way, the stronger a certain visual feature is in a given line graph, the higher the weight for the associated proposition.

Type 3 rules (Message Category + Visual Feature) differ only from type 2 rules in that they are restricted to a particular intended message category, rather than any line graph having the visual feature in question. For example, a proposition comparing the slope of two trends may be appropriate for a graph in the Change-trend message category, but does not make sense for a line graph with only a single trend (e.g., Rising-trend).

Once all propositions have been extracted and ranked, these weights are passed along to a graph-based content selection framework (Demir et al., 2010) that iteratively selects for inclusion in the initial summary those propositions which provide the best coverage of the highest-ranked information.

4.3 Sample Rule Application

Figures 1 and 4 consist of two different line graphs with the same intended message category: Change-trend. Figure 1 shows a stable trend in annual sea level difference from 1900 to 1930, followed by a rising trend through 2003, while Figure 4 shows a rising trend in Durango sales from 1997 to 1999, followed by a falling trend through 2006. Propositions associated with type 1 rules will have the same weights for both graphs, but propositions related to visual features may have different weights. For example, the graph in Figure 1 is far more volatile than the graph in Figure 4. Thus, the type 2 rule associated with volatility will have a very high weight for the graph in Figure 1 and will almost certainly be included in the initial summary of that line graph (e.g.,

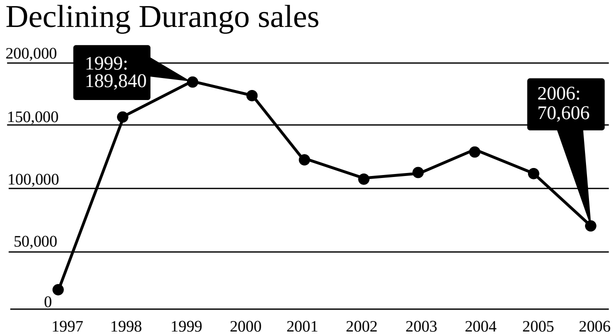

Figure 4: From "Chrysler: Plant had $800 million impact" in The (Wilmington) News Journal, Feb 15, 2007.

"*The values vary a lot...*", "*The trend is unstable...*"), possibly displacing a type 1 proposition that would still appear in the summary for the graph in Figure 4.

5 Future Work

Once the propositions that should be included in the summary have been selected, they must be coherently organized and realized as natural language sentences. We anticipate using the FUF/SURGE surface realizer (Elhadad and Robin, 1996); our collected corpus of line graph summaries provides a large set of real-world expressions to draw from when crafting the surface realization forms our system will produce for the final-output summaries. Our summarization methodology must also be evaluated. In particular, we must evaluate the rules for identifying the additional informational propositions that are used to elaborate the overall intended message, and the quality of the summaries both in terms of content and coherence.

6 Related Work

Image summarization has focused on constructing a smaller image that contains the important content of a larger image (Shi et al., 2009), selecting a set of representative images that summarize a collection of images (Baratis et al., 2008), or constructing a new diagram that summarizes one or more diagrams (Futrelle, 1999). However, all of these efforts produce an image as the end product, not a textual summary of the content of the image(s).

Ferres et al. (2007) developed a system for conveying graphs to blind users, but it generates the same basic information for each instance of a graph type (e.g., line graphs) regardless of the individual graph's specific characteristics. Efforts toward summarizing multimodal documents containing graphics have included naïve approaches relying on captions and direct references to the image in the text (Bhatia et al., 2009), while content-based image analysis and NLP techniques are being combined for multimodal document indexing and retrieval in the medical domain (Névéol et al., 2009).

Jing and McKeown (1999) approached abstractive summarization as a text-to-text generation task, modifying sentences from the original document via editing and rewriting. There have been some attempts to do abstractive summarization from semantic models, but most of it has focused on text documents (Rau et al., 1989; Reimer and Hahn, 1988), though Alexandersson (2003) used abstraction and semantic modeling for speech-to-speech translation and multilingual summary generation.

7 Discussion

Information graphics play an important communicative role in popular media and cannot be ignored. We have presented our methodology for constructing a summary of a line graph. Our method is abstractive, in that we identify the important high-level knowledge conveyed by a graphic and capture it in propositions to be realized in novel, coherent natural language sentences. The resulting summary can be integrated with a summary of the document's text to produce a rich summary of the entire multimodal document. In addition, the graphic's summary can be used along with a screen reader to provide sight-impaired users with full access to the knowledge conveyed by multimodal documents.

Acknowledgments

This work was supported in part by the National Institute on Disability and Rehabilitation Research under Grant No. H133G080047.

References

Jan Alexandersson. 2003. *Hybrid Discourse Modeling and Summarization for a Speech-to-Speech Translation System*. Ph.D. thesis, Saarland University.

Evdoxios Baratis, Euripides Petrakis, and Evangelos Milios. 2008. Automatic web site summarization by image content: A case study with logo and trademark images. *IEEE Transactions on Knowledge and Data Engineering*, 20(9):1195–1204.

Sumit Bhatia, Shibamouli Lahiri, and Prasenjit Mitra. 2009. Generating synopses for document-element search. In *Proceeding of the 18th ACM Conference on Information and Knowledge Management*, CIKM '09, pages 2003–2006, Hong Kong, November. ACM.

Sandra Carberry, Stephanie Elzer, and Seniz Demir. 2006. Information graphics: an untapped resource for digital libraries. In *Proc. of the 29th Annual Int'l ACM SIGIR Conf. on Research & Development in Information Retrieval*, SIGIR '06, pages 581–588, Seattle, August. ACM.

Daniel Chester and Stephanie Elzer. 2005. Getting computers to see information graphics so users do not have to. In *Proceedings of the 15th International Symposium on Methodologies for Intelligent Systems (LNAI 3488)*, ISMIS 2005, pages 660–668, Saratoga Springs, NY, June. Springer-Verlag.

Herbert Clark. 1996. *Using Language*. Cambridge University Press.

Seniz Demir, Sandra Carberry, and Kathleen F. McCoy. 2008. Generating textual summaries of bar charts. In *Proceedings of the 5th International Natural Language Generation Conference*, INLG 2008, pages 7–15, Salt Fork, Ohio, June. ACL.

Seniz Demir, Sandra Carberry, and Kathleen F. McCoy. 2010. A discourse-aware graph-based content-selection framework. In *Proceedings of the 6th International Natural Language Generation Conference*, INLG 2010, pages 17–26, Trim, Ireland, July. ACL.

Michael Elhadad and Jacques Robin. 1996. An overview of SURGE: a re-usable comprehensive syntactic realization component. In *Proceedings of the 8th International Natural Language Generation Workshop (Posters & Demos)*, Sussex, UK, June. ACL.

Stephanie Elzer, Sandra Carberry, Daniel Chester, Seniz Demir, Nancy Green, Ingrid Zukerman, and Keith Trnka. 2005. 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*, pages 223–230, Ann Arbor, June. ACL.

Stephanie Elzer, Sandra Carberry, and Ingrid Zukerman. 2011. The automated understanding of simple bar charts. *Artificial Intelligence*, 175:526–555, February.

Leo Ferres, Petro Verkhogliad, Gitte Lindgaard, Louis Boucher, Antoine Chretien, and Martin Lachance. 2007. Improving accessibility to statistical graphs: the iGraph-Lite system. In *Proc. of the 9th Int'l ACM SIGACCESS Conf. on Computers & Accessibility*, ASSETS '07, pages 67–74, Tempe, October. ACM.

Robert P. Futrelle. 1999. Summarization of diagrams in documents. In I. Mani and M. Maybury, editors, *Advances in Automatic Text Summarization*. MIT Press.

Barbara Grosz and Candace Sidner. 1986. Attention, Intentions, and the Structure of Discourse. *Computational Linguistics*, 12(3):175–204.

Hongyan Jing and Kathleen R. McKeown. 1999. The decomposition of human-written summary sentences. In *Proc. of the 22nd Annual Int'l ACM SIGIR Conf. on Research & Development in Information Retrieval*, SIGIR '99, pages 129–136, Berkeley, August. ACM.

Eamonn J. Keogh, Selina Chu, David Hart, and Michael J. Pazzani. 2001. An online algorithm for segmenting time series. In *Proceedings of the 2001 IEEE International Conference on Data Mining*, ICDM '01, pages 289–296, Washington, DC. IEEE.

Aurélie Névéol, Thomas M. Deserno, Stéfan J. Darmoni, Mark Oliver Güld, and Alan R. Aronson. 2009. Natural language processing versus content-based image analysis for medical document retrieval. *Journal of the American Society for Information Science and Technology*, 60(1):123–134.

John C. Platt. 1999. Fast training of support vector machines using sequential minimal optimization. In B. Schölkopf, C. J. C. Burges, and A. J. Smola, editors, *Advances in kernel methods: support vector learning*, pages 185–208. MIT Press, Cambridge, MA, USA.

Lisa F. Rau, Paul S. Jacobs, and Uri Zernik. 1989. Information extraction and text summarization using linguistic knowledge acquisition. *Information Processing & Management*, 25(4):419 – 428.

Ulrich Reimer and Udo Hahn. 1988. Text condensation as knowledge base abstraction. In *Proceedings of the 4th Conference on Artificial Intelligence Applications*, CAIA '88, pages 338–344, San Diego, March. IEEE.

Liang Shi, Jinqiao Wang, Lei Xu, Hanqing Lu, and Changsheng Xu. 2009. Context saliency based image summarization. In *Proceedings of the 2009 IEEE international conference on Multimedia and Expo*, ICME '09, pages 270–273, New York. IEEE.

Pang-Ning Tan, Michael Steinbach, and Vipin Kumar. 2005. *Introduction to Data Mining*. Addison Wesley.

Peng Wu, Sandra Carberry, and Stephanie Elzer. 2010a. Segmenting line graphs into trends. In *Proceedings of the 2010 International Conference on Artificial Intelligence*, ICAI '10, pages 697–703, Las Vegas, July.

Peng Wu, Sandra Carberry, Stephanie Elzer, and Daniel Chester. 2010b. Recognizing the intended message of line graphs. In *Proc. of the 6th Int'l Conf. on Diagrammatic Representation & Inference*, Diagrams '10, pages 220–234, Portland. Springer-Verlag.

Jeff Zacks and Barbara Tversky. 1999. Bars and lines: A study of graphic communication. *Memory & Cognition*, 27:1073–1079.