Extending Plan Inference Techniques to Recognize Intentions in Information Graphics*

Stephanie Elzer[1], Nancy Green[2], Sandra Carberry[1], and Kathleen McCoy[1]

[1] Department of Computer Science, University of Delaware
Newark, DE 19716 USA
[2] Computer Science Division – Department of Mathematical Sciences
University of North Carolina at Greensboro, Greensboro, NC 27402, USA

Abstract. Plan inference techniques have been used extensively to understand natural language dialogue. But as noted by Clark[5], language and communication are more than just utterances. This paper presents the problems that we have had to address and the solutions that we have devised in designing a system to recognize intentions from *information graphics*. Our work is part of a larger project to develop an interactive natural language system that provides an alternative means for individuals with sight-impairments to access the content of information graphics.

1 Introduction

The amount of information available electronically has increased dramatically over the past decade. Unfortunately, many knowledge sources are provided in a single format and thus are not accessible to everyone. For example, individuals with impaired eyesight have limited access to graphical displays, thus preventing them from fully utilizing available information resources. Although research has investigated alternative modes of presentation of graphical information for people who have visual impairments, their focus is on rendering graphical elements in an alternative medium and they have serious limitations. For example, it would be extremely difficult for a user to compare two related lines on a line graph via a soundscape[13]. The underlying hypothesis of our work is that alternative access to what the graphic looks like is not enough — the user should be provided with the message and knowledge that one would gain from viewing the graphic in order to enable effective and efficient use of this information resource.

Our overall goal is to develop an interactive natural language system that infers the intended message underlying an *information graphic* (a non-pictorial graphic such as a bar chart or a line graph),[3] provides an initial summary that includes the intended message along with notable features of the graphic, and then responds to follow-up questions from the user. Recognizing the intended message of an information graphic also has other applications. For example, as multimodal communication becomes more prevalent, we envision users engaging in interactive communication via text and graphics; an artificial agent will need to be able to recognize the intentions that the user wants to convey via his information graphics in order to respond appropriately.

* The work of the second author was supported by the National Science Foundation under Grant No. 0132821.

[3] *Information graphics* are distinguished from depictions of concrete entities and scientific visualizations of spatial data[4].

Plan recognition, a subarea of user modeling, plays a central role in our work since identifying the intended message of the graphic designer is critical. Although one might suggest relying on captions to provide the message of a graphic, Corio found in a large corpus study[6] that captions are often missing or very general. Once the intentions underlying the graphic have been inferred, they can be used 1) to construct a summary that includes the message intended by the person who constructed the graphic (when the system is serving as an alternative communication system for an individual with sight impairments), or 2) to respond appropriately to the user (when the system is an artificial agent engaged in a multimodal interaction with a user). In addition, the system should determine whether the intended message is warranted by the displayed data. Mittal[14] identified a number of strategies that are frequently employed to construct deceptive graphics, such as truncating the vertical axis of a bar chart (starting the axis at a value larger than 0), thereby magnifying differences in the heights of the bars when their actual values are proportionally quite close. A summary should call attention to such discrepancies, whereas an artificial agent should note the user's intention to be deceptive and react accordingly.

This paper focuses on the novel application of a theory of plan-based intention recognition, and presents our solutions to issues that we have had to address in extending plan inference, which has previously been applied to utterances that are part of a dialogue, to recognize the intended message underlying an information graphic.

2 Recognizing Intention from Information Graphics

Language research has posited that a speaker or writer executes a speech act whose intended meaning he expects the listener to be able to deduce, and that the listener identifies the intended meaning by reasoning about the observed signals and the mutual beliefs of author and interpreter[7, 5]. But as noted by Clark[5], language is more than just words. It is any "signal" (or lack of signal when one is expected), where a signal is a deliberate action that is intended to convey a message. Although some information graphics are only intended to display data values,[4] the overwhelming majority of the graphics that we have examined (taken from newspaper, magazine, and web articles) appear to have some underlying goal, such as getting the reader to believe that a particular mutual fund has fared much better than the S&P-500 and thus that the reader should purchase the mutual fund.

Applying Clark's view of language to information graphics, it is reasonable to presume that the author of an information graphic similarly expects the viewer to deduce from the graphic the message that he intended to convey by reasoning about the graphic itself and the salience of entities in the graphic. Beginning with the seminal work of Allen[15] who developed a system for deducing the intended meaning of an indirect speech act, researchers have applied plan inference techniques to a variety of problems associated with understanding utterances, particularly utterances that are part of a dialogue. But extending plan inference techniques to the recognition of intentions from information graphics is not a straightforward task and requires that a number of issues be addressed.

In the case of information graphics, the designer has one or more high-level goals which cause him to construct a graphic that he believes will lead the viewer to perform certain perceptual and cognitive tasks[9] which, along with other knowledge, will cause

[4] [19] used pattern recognition techniques to summarize interesting features of automatically generated graphs of time-series data from a gas turbine engine.

the viewer to recognize the message that the designer intends the graphic to convey. By *perceptual tasks* we mean tasks that can be performed by simply viewing the graphic, such as finding the top of a bar in a bar chart; by *cognitive tasks* we mean tasks that are done via mental computations, such as computing the difference between two numbers. Section 2.1 and Section 2.2 discuss the kinds of knowledge that must be explicitly available for plan inference from information graphics. In particular, they present our approach that encodes knowledge about perceptual and cognitive tasks in plan operators and encodes knowledge about the effort required for different perceptual tasks in rules associated with primitive subgoals. These sections also describe how the operators and rules are used in plan inference.

The graphic designer has many alternative ways of designing a graphic, and the design choices facilitate some perceptual tasks more than others. Following the AutoBrief work[9] on generating graphics that fulfill communicative goals, we hypothesize that the designer chooses a design that best facilitates the tasks that are most important to conveying his intended message, subject to the constraints imposed by competing tasks. Section 2.2 presents our approach to capturing knowledge about the effort required for different perceptual tasks, as well as our approach to identifying tasks that the graphic designer intended to be salient for the viewer, and how all of this information is used as a starting point for plan inference. Section 2.3 presents our approach to guiding the search through the space of candidate plans.

2.1 Plan Operators for Information Graphics

In their work on multimedia generation, the AutoBrief group proposed that speech act theory can be extended to the generation of graphical presentations[9]. During the first phase of graphics generation in AutoBrief, media-independent communicative goals are mapped to perceptual and cognitive tasks that the graphics should support. For example, if the goal is for the viewer to believe that Company A had the highest profits of a set of companies, then it would be desirable to design a graphic that facilitates the tasks of comparing the profits of all the companies, locating the maximum profit, and identifying the company associated with the maximum. In the second phase of graphics generation, a specification of the tasks that the graphic should support, along with a description of the data, is input to an automatic graphic designer that uses constraint satisfaction, along with knowledge about the effectiveness of different design techniques for supporting different kinds of tasks, to design the graphic.

AutoBrief used algorithms to map communicative goals to partially ordered sequences of tasks and to reason about how to realize the graphic. For plan recognition we need to explicitly encode, in such a way that the plan inference system has access to it, detailed knowledge about how communicative goals decompose into perceptual and cognitive tasks and how different perceptual tasks by the viewer can be enabled by particular realizations of the graphic. Then we can reason backwards from the observed graphic to hypothesize what goals might have motivated its design. Our operators decompose knowledge goals (such as getting the viewer to believe that a mutual fund has risen substantially in value over the past decade) into tasks that the viewer must be able to perform using the information graphic. Such tasks may be further decomposed into a sequence of simpler

Goal:	Find-value(<viewer>, <g>, <e>, <ds>, <att>, <v>)
Gloss:	Given graphical element <e> in graphic <g>, <viewer> can find the value <v> in dataset <ds> of attribute <att> for <e>
Data-req:	Dependent-variable(<att>, <ds>)
Body:	1. Perceive-dependent-value(<viewer>, <g>, <att>, <e>, <v>)

Fig. 1. Operator for achieving a goal perceptually

tasks, which eventually decompose into perceptual or cognitive primitives.[5] Associated with each perceptual primitive are rules that consider the various ways that the perceptual task could be enabled and estimate the viewer effort that would be required for each choice[6](see Section 2.2).

Our plan operators for achieving goals via information graphics consist of:

– **Goal:** the goal that the operator achieves
– **Data-req:** requirements which the data must satisfy in order for the operator to be applicable in a graphic planning paradigm
– **Display-const:** features that constrain how the information graphic is eventually constructed if this operator is part of the final plan
– **Body:** lower-level subgoals that must be accomplished in order to achieve the overall goal of the operator

Plan inference reasons backwards from an XML representation of an observed graphic that is provided by a computer vision module. The display constraints are used to eliminate operators from consideration (i.e., if the graphic does not capture the operator's constraints on the display, then the operator could not have been part of a plan that produced the graphic). The data requirements are used to instantiate parameters in the operator (i.e., the data must have had certain characteristics for the operator to have been included in the graphic designer's plan, and these often limit how the operator's arguments can be instantiated). Goals can often be accomplished in several different ways. Figures 1 and 2 present two operators that can be used to achieve the goal of enabling the viewer to find the value of an attribute for a graphical element <e> (for example, the y-value of a point on a line graph or the y-value for the top of a bar on a vertical bar chart). The body of the operator in Figure 1 consists of a primitive perceptual task *Perceive-dependent-value*, in which the viewer simply perceives the value. Since *Perceive-dependent-value* is a primitive task, there is no operator that decomposes it further; if *Perceive-dependent-value* is a task in the final plan for the graphic, then the graphic would need to be realized so that the attribute value could be directly perceived, such as by annotating the element in the graphic with its value as is done for the bars in Figure 4. As mentioned earlier, this knowledge about how a graphic might be realized to enable a primitive perceptual task is captured in the rules discussed in Section 2.2 that compute the effort required to perform the task given

[5] We are treating subgoals as *primitives* when they are simple tasks that cannot be decomposed further with our operators. This is not to be confused with a psychological primitive.

[6] Of course, a planner would need to ensure that all of the tasks comprising the plan could be achieved in a single graphic. In addition, it would need to decide to what degree each task would be enabled if it was not possible to design a graphic that would enable them all to the fullest extent.

Goal:	Find-value(<viewer>, <g>, <e>, <ds>, <att>, <v>)
Gloss:	Given graphical element <e> in graphic <g>, <viewer> can find the value <v> in dataset <ds> of attribute <att> for <e>
Data-req:	Natural-quantitative-ordering(<att>)
Display-const:	Ordered-values-on-axis(<g>, <axis>, <att>)
Body:	1. Perceive-info-to-interpolate(<viewer>, <g>, <axis>, <e>, $<l_1>$, $<l_2>$, <f>) 2. Interpolate(<viewer>, $<l_1>$, $<l_2>$, <f>, <v>)

Fig. 2. Operator that employs both perceptual and cognitive subgoals

different realizations. On the other hand, the operator in Figure 2 specifies how the same goal can be achieved, admittedly with more effort, using a combination of perceptual and cognitive tasks. The first subgoal, *Perceive-info-to-interpolate*, is a primitive perceptual task in which the viewer perceives the labels $<l_1>$ and $<l_2>$ immediately below and above the location on <axis> corresponding to graphical element <e> of graph <g> and the fraction <f> of the distance that this location lies between $<l_1>$ and $<l_2>$. The second subgoal, *Interpolate*, is a primitive cognitive task in which the viewer computes (via interpolation) the value <v> of attribute <att> for graphical element <e> based on $<l_1>$, $<l_2>$ and <f>. The operator in Figure 2 places constraints on the graphical display if this operator is used to construct the plan, namely that the values of the desired attribute be displayed on <axis> in ascending or descending order. The reason for this is that if the surrounding values produced by achieving the first subgoal are not from an ordered set of labels on the axis, then interpolation to get the value of <v> is not possible.

2.2 Starting Point for Plan Inference

In plan recognition, we must reason about the graphical choices that resulted in the graphic. We contend that the designer made these choices in order to make "important" tasks as easy or as salient as possible. The graphic designer can make a task easy for the viewer to perform by the choice of graphic type (for example, bar chart versus pie chart) and the organization and presentation of data. The graphic designer might also intend a task to be particularly salient, or of notable significance, to the viewer. We must have mechanisms for identifying the easiest and most salient tasks so that we can use these tasks as a starting point for the plan inference process.

Estimating Viewer Effort In order to estimate which primitive tasks are easy/hard (presumably the important ones are among the easiest tasks), our system uses rules that estimate the viewer effort involved in performing different primitive tasks. Our estimates of viewer effort are based on research by cognitive psychologists such as Lohse[12] who developed a cognitive model of information graphic perception that was intended to simulate human performance on graphic comprehension tasks. Figure 3 presents a rule for estimating the effort involved in finding the top <e> of a bar in a bar chart given the bar's <label>. Each rule consists of a set of condition-computation pairs, ordered so that the computations producing the lowest estimates of effort appear first. The conditions specify characteristics of the graphic which are necessary for the associated computation to be

Rule1: Estimate effort for task Perceive-element(<viewer>, <g>, <e>, <label>)
Graphic-type: bar-chart
Gloss: Compute effort for finding the top <e> of a bar whose label is <label> in graphic <g>
B1-1: IF labels on the independent axis appear in sorted order,
THEN cost=scan + 150 + 300 + 230
B1-2: If labels on the independent axis do not appear in sorted order,
THEN cost=scan + ((150 + 300)$\times$ number-of-preceding-items) + 150 + 300 + 230

Fig. 3. A rule for estimating effort for a primitive perceptual task

applicable. In Figure 3 the first condition-computation pair, labelled B1-1, is applicable if the labels appear in sorted order; in such cases, the effort is estimated as the cost of scanning along the x-axis until reaching the label (measured in terms of the degrees of visual arc scanned[10]), 150 units for discriminating the label (based on work by Lohse[12]), 300 units for recognizing a 6-letter word[8], and 230 units for making the saccade up to the top of the bar[17]. The second condition-computation pair, labelled B1-2, is applicable if the labels do not appear in sorted order; in this case, the costs for discrimination and recognition are charged against each label up to and including the one being sought. Often several conditions within a single rule will be satisfied; this might occur for example if the top of a bar in a vertical bar chart both falls on a tick mark and has its value annotated at the top of the bar; the easiest way to get the value represented by the top of the bar would be to read the annotated value although it could also be obtained by scanning across to the tick mark on the dependent axis. The computation associated with the first satisfied condition in a rule is used, thereby estimating the least effort required to perform the task.

The vision component gives the plan inference module an XML schema representing the information graphic. The effort estimates for the primitive perceptual tasks are then generated, and the least costly ones are selected. Operators containing these primitive tasks as subgoals are then used to begin the plan inference process. As bottom-up chaining suggests higher-level operators for consideration, the effort expended by the viewer to achieve the operator's goal is estimated on the basis of the effort expended to achieve the subgoals in the operator's body. This can cause a downward expansion of the operator into primitive tasks whose effort is computed.

Identifying Salient Tasks A plan inference system should exploit all available evidence in recognizing intention. For information graphics, this entails not only reasoning about the set of perceptual tasks that are best enabled by the graphic, but also identifying any tasks that the graphic designer intended to be salient for the viewer. We have identified three sources of salience: captions, a model of mutual beliefs about entities of interest to members of the viewing audience, and highlighted entities in the information graphic.

Noun phrases in captions are potentially salient to the intended message of the graphic. For example, consider the graphic in Figure 4 and suppose that it had the caption "United States Tops China in Gold Production". The presence of the nouns *United States* and *China* in the caption suggests that they are relevant to the designer's intended message. Consequently, UNITED-STATES and CHINA would be used to instantiate <label> in two instances of the perceptual task Perceive-element(U, G, <e>, <label>), where U and G are constants designating the particular viewer and graphic. This produces the

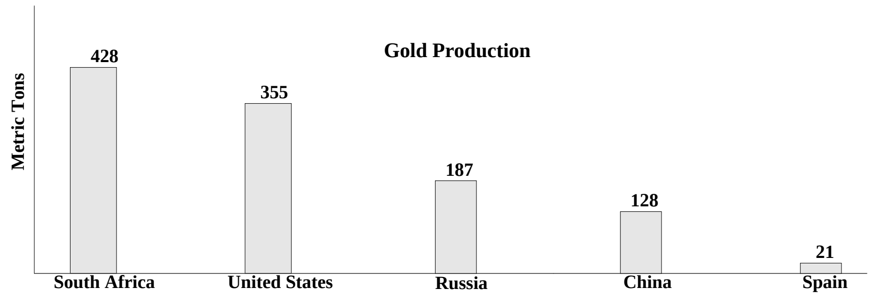

Fig. 4. A sample information graphic

Goal: Believe(<viewer>, Comparison(<g>, $<att_y>$, <ds>, $<v_1>$, $<v_2>$, $<att_x>$, $< rel >$))

Gloss: Viewer to believe from graphic <g> that the value of attribute $<att_y>$ for the element $<e_1>$of dataset <ds> whose $<att_x>$ value is $<v_1>$ is <rel> (greater-than, less-than, equal-to) the value of attribute $<att_y>$ for the element $<e_2>$ whose $<att_x>$ value is $<v_2>$

Data-req: Natural-quantitative-ordering($<att_y>$)
Independent-variable($<att_x>$, <ds>)
Related($<att_y>$, <ds>, $<att_x>$, $<v_1>$, $<v_2>$, <rel>)

Body: 1. Perceive-element(<viewer>, <g>, $<e_1>$, $<v_1>$)
2. Perceive-element(<viewer>, <g>, $<e_2>$, $<v_2>$)
3. Determine-relationship(<viewer>, <g>, $<att_y>$, $<e_1>$, $<e_2>$, <rel>)

Fig. 5. Operator for determining the relationship between two graphical elements

---

salient primitive perceptual tasks Perceive-element(U, G, <e>, UNITED-STATES) and Perceive-element(U, G, <e>, CHINA) which are instantiations of subgoals in the operator shown in Figure 5. The results of perceiving these elements would be used to instantiate the arguments $<e_1>$ and $<e_2>$ in the operator in Figure 5 with TOP-BAR-2 and TOP-BAR-4 from the graphic in Figure 4, and the data requirements in the operator, along with an analysis of the graphic, would lead to the instantiation of other arguments, thereby suggesting that part of the graphic designer's intentions might be to get the viewer to believe that gold production in the United States exceeds that in China. Further chaining would lead to positing higher-level goals, and other aspects of the graphic would also suggest additional inferred plans.

A model of the intended recipient of the information graphic also plays a role in the plan recognition process. In designing the information graphic, the graphic designer takes into account mutual beliefs about entities that will be particularly salient to his audience. For example, if an information graphic appears in a document targeted at residents of Pittsburgh, then both the designer and the viewer will mutually believe that entities such as Pittsburgh, its football and baseball teams, etc. will be particularly salient to the viewer. The viewer model captures these beliefs, and our approach is to treat them in a manner similar to how we handle noun phrases in captions. Verb phrases in captions also provide evidence, but they suggest particular operators of interest rather than instantiations of operators, and thus we associate verbs with particular operators in the plan library.

One might wonder why we do not deal almost exclusively with captions to infer the intentions of the information graphic. Corio[6] performed a large corpus study of information graphics and noted that captions often do not give any indication of what the information graphic conveys. Our examination of a collection of graphics supports his findings. Thus we must be able to infer the message underlying a graphic when captions are missing or of little use.

Graphic designers also use techniques to highlight particular aspects of the graphic, thus making them more salient to the viewer. Such techniques include the use of color or shading for individual elements of a graphic, annotations such as an asterisk, an arrow pointing to a particular location on the graphic, or a pie chart with a single piece "exploded." Mittal[14] discusses a variety of such design techniques in the context of distorting the message inferred from the graphic. Our working hypothesis is that if the graphic designer goes to the effort of employing such attention-getting devices, then the highlighted items are almost certainly part of the intended message. Thus we treat highlighted entities in the information graphic as suggesting instantiations of primitive perceptual tasks that produce particularly salient tasks. Suppose for example that there was no caption on the information graphic depicted in Figure 4, but that the bars for United States and for China were highlighted by shading them darker than the other bars. This suggests that these bars are particularly relevant to the intended message of the information graphic. Consequently, we use the attributes of the bars to instantiate primitive perceptual tasks and produce tasks that are hypothesized to be salient.

2.3 Guiding the Search

A plan inference system must select, from among many plausible goals, the best hypothesis about the agent's intentions. Moreover, in any reasonably sized system, it will be necessary to guide the search through the space of possible plans and goals so that only a small proportion of the plan space is examined. This has generally been done either via Bayesian approaches that estimate the probability of different hypotheses[1] or heuristics that suggest which hypotheses should be considered first[2, 3, 11, 15]. Since we do not have available the probabilities necessary to construct a Bayesian system, we have chosen to use heuristics to guide the search. The question arises as to what features should be taken into account in heuristics that are used for plan inference from information graphics.

The evidence available for evaluating possible hypotheses includes:

- the effort expended by the viewer in carrying out the perceptual and cognitive tasks required for the message to be recognized. As higher-level goals are inferred, the effort required by the viewer to achieve those goals is computed from the effort estimates for its constituent subgoals. Since the graphic designer is assumed to be felicitous and trying to effectively convey his intended message, he is expected to construct a graphic that enables the requisite tasks so that they can be performed easily. Thus the greater the amount of effort required, the less likely it is that a candidate plan represents the designer's intentions.
- the extent to which elements of the information graphic (especially salient elements) and the perceptual tasks of least effort play a role in the inferred plan. If the graphic designer went to the effort of including particular graphical elements in the graphical

display and made particular perceptual tasks easiest or salient, then it is reasonable to believe that he made these design choices in order to facilitate the viewer recognizing his intended message. Thus the percentage of such elements and tasks that play a role in a candidate plan should influence how favorably we view it as a hypothesis about the designer's intentions.

– the extent to which parameters in an inferred plan are instantiated and the basis for the particular instantiations, as discussed below.

Arguments in operators may be instantiated for a variety of reasons. Perhaps the graphic permits only one instantiation (for example, if *Recognize-maximum* is an operator that is produced during chaining and the XML representation of the graphic indicates that there is a single graphical element whose value is greater than the others in the graph). Or the instantiation may be due to highlighting in the graphic, features extracted from the caption, or the model of mutual beliefs about entities of interest. Or the parameters may be instantiated with the values that produce the lowest effort estimations.

The basis for instantiating an argument in a primitive perceptual task, and therefore in a hypothesized plan, impacts confidence about whether the plan really represents the designer's intentions. In the case of plan inference from information graphics, if only one instantiation is possible or if an instantiation is suggested by highlighting or a caption or entities that are particularly salient to the targeted audience, that partial plan should be evaluated more favorably since the designer of the graphic has provided explicit reasons for the viewer to use these instantiations in recognizing his intentions.

Moreover, the proximity compatibility principle[18] dictates that the ratings of partial plans be increased further if the plans contain operators with arguments instantiated from multiple elements of the graphic that are similarly highlighted. The proximity compatibility principle is based on perceptual proximity (how perceptually similar two elements of a display are) and processing proximity (how closely linked the two elements are in terms of completing a task). According to the proximity compatibility principle, if two elements of a graphic are to be used in the same task, then the elements should be realized so that they have close perceptual proximity. [18] showed that violating this principle increased the cost of performing tasks that used multiple elements of a graphic. For example, the points in a line graph have a higher perceptual proximity than the bars in a bar chart. (This example applies the Gestalt law of good continuation[16].) The higher perceptual proximity of the points on a line graph means that it is easier to perform integrated tasks, such as recognizing a trend, in such a graph than with the bars in a bar chart. Along the same lines, if multiple items are similarly highlighted, as in our example where the bars for the United States and China were shaded the same but darker than the other bars, then the proximity compatibility principle suggests that they were intended to be part of an integrated reasoning task and thus part of the same intention.

3 Summary

Clark[5] has argued that communication includes more than natural language utterances, and that alternative modes of communication, such as hand signals, facial gestures, and drawings, bear many commonalities with natural language communication. This paper has presented a novel use of plan inference — namely, to infer the intended message underlying an information graphic — and it has presented our solutions to issues that we have had to address in extending plan inference from natural language utterances to information graphics. Our work is part of a larger project to develop an interactive natural language system that provides an alternative means for individuals with sight-impairments to access the content of information graphics.

References

1. David Albrecht, Ingrid Zukerman, and Ann Nicholson. Bayesian models for keyhole plan recognition in an adventure game. *User Modeling and User-Adapted Interaction*, pages 5–47, 1998.
2. Liliana Ardisonno and Dario Sestero. Using dynamic user models in the recognition of the plans of the user. *User Modeling and User-Adapted Interaction*, 5(2):157–190, 1996.
3. Sandra Carberry. *Plan Recognition in Natural Language Dialogue*. ACL-MIT Press Series on Natural Language Processing. MIT Press, Cambridge, Massachusetts, 1990.
4. S.K. Card, J. Mackinlay, and B. Shneiderman. *Readings in Information Visualization: Using Vision to Think*. Morgan-Kaufmann, 1999. Chapter 1.
5. Herbert Clark. *Using Language*. Cambridge University Press, 1996.
6. Marc Corio and Guy Lapalme. Generation of texts for information graphics. In *Proceedings of the 7th European Workshop on Natural Language Generation EWNLG'99*, pages 49–58, 1999.
7. H. Paul Grice. Utterer's Meaning and Intentions. *Philosophical Review*, 68:147–177, 1969.
8. B. E. John and A. Newell. Toward an engineering model of stimulus response compatibility. In R. W. Gilmore and T. G. Reeve, editors, *Stimulus-response compatibility: An integrated approach*, pages 107–115. North-Holland, New York, 1990.
9. Stephan Kerpedjiev and Steven Roth. Mapping communicative goals into conceptual tasks to generate graphics in discourse. In *Proceedings of the International Conference on Intelligent User Interfaces*, pages 60–67, 2000.
10. Stephen M. Kosslyn. Understanding charts and graphs. *Applied Cognitive Psychology*, 3:185–226, 1989.
11. Diane Litman and James Allen. A Plan Recognition Model for Subdialogues in Conversation. *Cognitive Science*, 11:163–200, 1987.
12. Gerald Lee Lohse. A cognitive model for understanding graphical perception. *Human-Computer Interaction*, 8:353–388, 1993.
13. Peter B. L. Meijer. An experimental system for auditory image representations. *IEEE Transactions on Biomedical Engineering*, 39(2):291–300, February 1992.
14. Vibhu Mittal. Visual prompts and graphical design: A framework for exploring the design space of 2-D charts and graphs. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*, pages 57–63, 1997.
15. Raymond Perrault and James Allen. A Plan-Based Analysis of Indirect Speech Acts. *American Journal of Computational Linguistics*, 6(3-4):167–182, 1980.
16. J. R. Pomerantz and M. Kubovy. Theoretical approaches to perceptual organization. In K. R. Boff, L. Kaufman, and J. P. Thomas, editors, *Handbook of Perception and Human Performance*, pages 36.1–36.46. Wiley, New York, 1986.
17. J. E. Russo. Adaptation of cognitive processes to eye movement systems. In J. W. Senders, D. F. Fisher, and R. A. Monty, editors, *Eye movements and higher psychological functions*, pages 89–109. Lawrence Erlbaum Associates, Inc., Hillsdale, NJ, 1978.
18. Christopher D. Wickens and C. Melody Carswell. The proximity compatibility principle: Its psychological foundation and relevance to display design. *Human Factors*, 37(3):473–494, 1995.
19. Jin Yu, Jim Hunter, Ehud Reiter, and Somayajulu Sripada. Recognising visual patterns to communicate gas turbine time-series data. In A. Macintosh, R. Ellis, and F. Coenen, editors, *Proceedings of ES2002*, pages 105–118, 2002.