

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Exploiting Cognitive Psychology Research for Recognizing Intention in Information Graphics

Permalink

<https://escholarship.org/uc/item/14t8d8bj>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 25(25)

ISSN

1069-7977

Authors

Elzer, Stephanie
Green, Nancy
Carberry, Sandra

Publication Date

2003

Peer reviewed

Exploiting Cognitive Psychology Research for Recognizing Intention in Information Graphics

Stephanie Elzer

Dept of Computer Science
University of Delaware
Newark, DE 19716 USA
(elzer@cis.udel.edu)

Nancy Green¹

Dept of Mathematical Sciences
Univ. of N. Carolina at Greensboro
Greensboro, NC 27402, USA
(nlgreen@uncg.edu)

Sandra Carberry

Dept of Computer Science
University of Delaware
Newark, DE 19716 USA
(carberry@cis.udel.edu)

Abstract

This paper outlines our approach to a novel application of plan inference, recognizing the intended message of an information graphic, focusing on how results from research by cognitive psychologists have been incorporated into the design of our system. 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.

Introduction

Information graphics (line graphs, bar charts, etc.) are pervasive in popular media such as newspaper and magazine articles. Our analysis of a corpus of information graphics from such documents indicates that information graphics generally have a communicative goal (for example, to convince a viewer that a particular mutual fund has an upward trend and that the fund's performance is better than that of the S&P-500 and so is a good buy) and that information graphics often carry information content that is not available from the text alone. Unfortunately, individuals with impaired eyesight have limited access to information graphics and thus cannot fully utilize this information resource. Although some projects have attempted to reproduce the image in an alternative medium, such as via soundscapes [Meijer, 1992], these approaches are ineffective with complex graphics such as multiple line graphs and they require the user to develop a "mental map" of the information graphic, which puts congenitally blind users at a disadvantage since they do not have the personal knowledge to assist them in the interpretation of the image [Kennel, 1996]. The overall goal of our project is an interactive natural language system that will provide the user with the same knowledge about an information graphic, to any desired degree of detail, that would be obtained by viewing it. The envisioned interactive system would infer the intended message of the information graphic, provide an initial summary including the intended message along with notable features of the graphic, and then respond to follow-up questions from the user.

Cognitive psychologists have investigated how humans perceive and process information, and artificial intelligence researchers have implemented systems that exhibit intelligent behavior. Although cognitive modelling is central to projects such as SOAR [SOAR, 2003], ACT-R [ACT-R, 2003], and EPIC [EPIC, 2003], too often AI researchers do not incorporate the results of cognitive psychology research into artificial intelligence systems. This paper outlines our approach to recognizing the intended message of an information graphic, focusing on how results from research by cognitive psychologists have been incorporated into the design of our system.

Information Graphics as Language

As noted by Clark [1996], 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. 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 [Grice, 1969, Clark, 1996]. 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, the salience of entities in the graphic, and mutual beliefs.

Beginning with the seminal work of Wilensky [1981] who recognized the importance of inferring characters' goals in order to understand a story and Perrault and Allen [1980] 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 discourse and dialogue. Given domain knowledge in the form of operators that decompose goals into a sequence of subgoals, along with evidence in the form of an observed action (such as a character's action in a story or a speech act), a plan inference system chains backwards on the plan operators to deduce one or more high-level goals that might have led the agent to perform the observed action as part of an overall plan for achieving his goal(s).

When designing an information graphic, the designer

¹The work of the second author was supported by the National Science Foundation under Grant No. 0132821.

has one or more high-level communicative goals. Consequently, he constructs an information graphic that he believes will enable the viewer to perform certain perceptual and cognitive tasks which, along with other knowledge, will enable the viewer to recognize the message that the designer intends the graphic to convey [Kerpedjiev and Roth, 2000]. 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.

In our research, we extend plan inference techniques (that have been used successfully on natural language discourse) to inferring intention from information graphics. Our plan operators capture knowledge about how the graphic designer's goal of conveying a message can be achieved via the viewer performing certain perceptual and cognitive tasks, as well as knowledge about how perceptual and cognitive tasks decompose into sets of simpler tasks. Using these plan operators, we can chain from evidence provided by the information graphic to eventually reach a high-level goal that captures the message underlying the graphic in the same way that plan inference systems chain from a speech act to the probable goals of the speaker. However, extending plan inference techniques to the recognition of intentions from information graphics is not a straightforward task [Elzer et al., 2003]. Several questions must be addressed:

1. *What should constitute the evidence from the graphic that should be used to start the plan inference process?*
2. *How can evidence be used to guide the search through the space of possible plans that could be produced via chaining?*

In addressing each of these questions, we have made recourse to results from cognitive psychology research.

Starting Point for Plan Inference

Given a set of data, the graphic designer has many alternative ways of designing a graphic. As Larkin and Simon [1987] note, information graphics that are informationally equivalent (all of the information in one graphic can also be inferred from the other) are not necessarily computationally equivalent (enabling the same inferences to be drawn quickly and easily). Peebles and Cheng [in press] take this one step further by observing that even in graphics that are informationally equivalent, the design of the graphic can affect viewers' performance of graph reading tasks. Much of this can be attributed to the fact that design choices made while constructing an information graphic will facilitate some perceptual tasks more than others. Following the AutoBrief work on generating graphics to achieve 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 [Kerpedjiev and Roth, 2000].

We contend that the designer made these choices in order to make "important" tasks as easy as possible.

This can be done through a variety of techniques such as choice of graphic type (for example, bar chart versus pie chart) and the organization and presentation of data. If, for instance, the graphic designer wants the viewer to find the exact value represented by the top of a bar in a bar chart, this task could be made easier by annotating the bar with its exact value. If the graphic designer wants the viewer to compare the relative values of two bars, this task could be facilitated by putting the bars immediately beside each other and highlighting the bars to draw attention to them.

In order to apply plan inference techniques to recognizing the intended message of an information graphic, we must identify the evidence in the graphic that should be used to start the plan inference process. Our methodology is to apply the results of research from cognitive psychology to construct rules that estimate the effort required for different perceptual tasks within a given information graphic, and thereby identify the perceptual tasks that the graphic designer has best enabled in the graphic. Our working hypothesis is that the *easiest* tasks are good candidates for tasks that the viewer was intended to perform, since the designer went to the effort of making them easy to accomplish. We can then use this set of the easiest perceptual tasks along with any unusually salient tasks (discussed later in this paper) as a starting point for our plan inference process. By reasoning about the more complex tasks in which these perceptual tasks play a role, we can hypothesize the message that the graphic designer intended the viewer to extract from the graphic. The component of our system that is responsible for estimating effort is called APTE (Analysis of Perceptual Task Effort).

Analysis of Perceptual Task Effort

The goal of APTE is to determine whether a task is easy or hard with respect to other perceptual tasks that could be performed from an information graphic. In order to estimate the relative effort involved in performing a task, we adopt a GOMS-like approach [Card et al., 1983], decomposing each task into a set of expected component tasks. Following other cognitive psychology research, we take the principle measure of the effort involved in performing a task as the amount of time that it takes to perform the task, and our effort estimates are based on time estimates for the component tasks. Wherever possible, we utilize the estimates applied by Lohse [1993] in his UCIE system, a cognitive model of information graphic perception that was intended to simulate and predict human performance on graphic comprehension tasks. In doing this, we are not attempting to develop a predictive model of our own – our aim is to be able to identify the tasks that the designer would expect to have best facilitated by his design choices in order to apply that information to the plan inference process.

Structure of Rules

APTE contains a set of rules that estimate how well a task is enabled in an information graphic. Each rule

Rule-1: Estimate effort for task Perceive-dependent-value(<viewer>, <g>, <att>, <e>, <v>)
 Graphic-type: bar-chart
 Gloss: Compute effort for finding the exact value <v> for attribute <att> represented by top <e> of a bar in graph <g>
 B1-1: IF the top of bar is annotated with a value,
 THEN effort=150 + 300
 B1-2: IF the top <e> of bar aligns with a labelled tick mark on the dependent axis,
 THEN effort=scan + 150 + 300

Figure 1: A rule for estimating effort for the primitive perceptual task *Perceive-value*

Rule-2: Estimate effort for task Perceive-info-to-interpolate(<viewer>, <g>, <axis>, <e>, <l₁>, <l₂>, <f>)
 Graphic-type: bar-chart
 Gloss: Compute effort for finding the information needed for interpolation, including the labels <l₁> and <l₂> on either side of entity <e> on axis <axis> in graph <g>, and the fraction <f> that is the distance between <l₁> and entity <e> on <axis> relative to the distance between <l₁> and <l₂>
 B2-1: IF <axis> is labelled with values THEN effort=scan + 150 + ((230 + 150 + 300) x 2)

Figure 2: A rule for estimating effort for the primitive perceptual task *Perceive-info-to-interpolate*




Figure 3: Information Graphic Example

captures a perceptual task that can be performed on a particular type of information graphic (line graph, bar chart, and so forth), along with the conditions (design choices) that affect the difficulty of performing that task. The conditions for the tasks are ordered so that the conditions producing the lowest estimates of effort appear first. Often several conditions within a single rule will be satisfied – this might occur, for example, in the rule shown in Figure 1 which estimates the effort of determining the exact value represented by the top of a bar in a bar chart. Condition-computation pair B1-1 estimates the effort involved when the bar is annotated with the value; this condition is illustrated by the bars that are annotated with their values in the bar chart in Figure 3. The second condition-computation pair, B1-2, is applicable when the top of the bar aligns with a labelled tick mark on the dependent axis (in a vertical bar chart, the dependent axis is the y-axis); this condition is illustrated by all bars except the second bar in Figure 3. If the top of a bar both falls on a tick mark and has its value annotated at the top of the bar (as in the fourth bar in the bar chart in Figure 3), 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. When multiple conditions are applicable, the first condition that is satisfied will be applied to calculate the effort estimate, thereby estimating the least expected effort required to perform the task.

Applying Estimates of Component Tasks

Researchers have examined many different perceptual tasks, although often studying individual perceptual tasks in isolation. As mentioned earlier, we have followed Lohse's approach [1993] in breaking down our tasks into component tasks. We then utilize existing time estimates (primarily the estimates applied in Lohse's UCIE system [1993]) for those component tasks wherever possible. For some perceptual tasks, this has been a sufficient foundation for our rules. For example, we developed effort estimates for the rule shown in Figure 1 in this manner. In the case of condition-computation pair B1-1, finding the exact value for a bar where the bar is annotated with the value, the effort is estimated as 150 units for discriminating the label (based on work by Lohse [1993]) and 300 units for recognizing a 6-letter word [John and Newell, 1990]. In the case of B1-2, finding the exact value for a bar where the top of the bar is aligned with a tick mark on the axis, the effort estimate includes scanning over to the dependent axis in order to read the value (measured in terms of distance in order to estimate the degrees of visual arc scanned [Kosslyn, 1989]) in addition to the effort of discriminating and recognizing the label.

Notice that the rule shown in Figure 1 does not cover the situation where the viewer must interpolate between two labelled values on the dependent axis. Performing

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_{12 2. Interpolate(<viewer>, <l₁₂ }

Figure 4: Plan operator that employs both perceptual and cognitive subgoals

interpolation to find the exact value represented by the top of a bar involves 1) the perceptual task of finding the intersecting location on the axis and recognizing the two surrounding labels and 2) the mental or cognitive task of interpolating to find the appropriate value. In our system, this more complex goal of finding the value represented by a bar in a bar chart via interpolation is captured by the plan operator shown in Figure 4, whose body consists of the perceptual task *perceive-info-to-interpolate* and the cognitive task *interpolate*. Associated with the *perceive-info-to-interpolate* task is an APTE rule (Figure 2) for estimating the effort to perform this perceptual task; similarly, there is a cognitive rule (not discussed in this paper) for estimating the effort associated with the cognitive task *interpolate*. We developed the effort estimate for the *perceive-info-to-interpolate* task by applying estimates for the component tasks — the effort of the scan to the dependent axis (based on [Kosslyn, 1989]), the effort of discriminating the intersection location on the axis (150 units based on [Lohse, 1993]), plus the effort of the saccade to each label (230 units [Russo, 1978]) along with the effort involved in discriminating and recognizing the labels. The cumulative effort associated with the *Find-value* goal will be the sum of the effort associated with each subgoal in the body of the operator (Figure 4).

Exploiting Cognitive Psychology Principles

For more complex tasks that have not been explicitly studied by cognitive psychologists, we have applied existing principles and laws in the development of our rules for estimating perceptual effort. An example of this is the class of comparison tasks (for example, finding the maximum or minimum value represented by the tops of the bars in a bar chart or comparing the tops of two bars to determine the relative difference in value), where the proximity compatibility principle espoused by Wickens and Carswell [1995] plays a major role. This principle is based on two types of proximity; *perceptual proximity* refers to how perceptually similar two elements of a display are (in terms of spatial closeness, color, shape, etc.) while *processing proximity* refers to how closely linked the two elements are in terms of completing a particular task. If the elements must be used together (integrated) in order to complete a task, they have close processing proximity. The proximity compatibility principle states that if there is close processing proximity between two elements, then close perceptual proximity is advised. If

two elements are intended to be processed independently, then distant perceptual proximity is advised. Violating the principle will increase the effort required for a viewer to process the information contained in the display.

We assume that the graphic designer attempted to follow the proximity compatibility principle in designing the information graphic so as to facilitate intended tasks and make them easier to perform than if the principle were violated. This assumption is reflected in the rule shown in Figure 5, where the effort required to perform the integrated task of determining the relative difference between two bars is different based on the spatial proximity of the two bars. If the bars are adjacent, the effort of doing the comparison will be lower than if the bars are not adjacent. We also apply this principle when defining the effort of performing the same perceptual tasks on different types of information graphics. For example, the elements (points) in a line graph have a higher perceptual proximity than the bars in a bar chart (this example of perceptual proximity applies the Gestalt law of good continuation [Pomerantz and Kubovy, 1986]). This means that it will be easier to perform integrated tasks with the points on a line in a line graph than it will be to perform the same task with the bars in a bar chart.

Weber's Law [Cleveland, 1985] has also played a critical role in our rules. Many of the tasks for which we have had to develop effort estimates, including the comparison tasks described above, involve discriminating between two or more graphical elements; these tasks require the viewer to make comparative judgments of length, area, and angle. In order to define the conditions affecting the complexity of these judgments, we have applied Weber's Law [Cleveland, 1985]. One of the implications of Weber's Law is that a fixed percentage increase in line length or area is required to enable discrimination between two entities with a particular probability (and the probability of discrimination is affected not by object size, but by the percentage increase). Weber's Law has influenced the thresholds used in rules for estimating effort such as *Rule-3* in Figure 5 where thresholds in the percentage difference in the height of the bars influence the effort required to perceptually discriminate the relative difference between the values represented by the bars.

In some cases, the optimal combination of component tasks (also representing the optimal scan path) does not take into account the escalating complexity captured by

Rule-3: Estimate effort for task Perceive-relative-diff(<viewer>, <g>, <e1>, <e2>, <b1>, <b2>, <r>, <d>)

Graphic-type: bar-chart

Gloss: Compute effort for finding the relative difference <r> in value (greater than/less than/equal to) and degree <d> of difference (high/low/none) represented by the tops <e1> and <e2> of two bars <b1> and <b2> in graph <g>

B3-1: IF bar <b1> and bar <b2> are adjacent and the height difference is >20%
THEN effort=92 + 230

B3-2: IF bar <b1> and bar <b2> are adjacent and the height difference is >5%
THEN effort=92 + 460

B3-3: IF bar <b1> and bar <b2> are not adjacent and the height difference is >20%
THEN effort=92 + 460

B3-4: IF bar <b1> and bar <b2> are not adjacent and the height difference is >5%
THEN effort=92 + 690

Figure 5: A rule for estimating effort for the primitive perceptual task *Perceive-relative-diff*

the conditions of the rule. For example, the optimal scan path would be the same for all conditions of *Rule-3* (Figure 5) even though the difficulty of making the required perceptual judgment can vary greatly. Therefore, when estimating the effort in such cases, we estimate the expected number of saccades that will be required by the average viewer in order to perform the necessary perceptual judgment. Thus the effort estimates shown in Figure 5 show the estimate of 92 units to perform a perceptual judgment [Welford, 1973] along with a multiple of 230 units where 230 represents the estimate for a single saccade [Russo, 1978].

Output

APTE takes as input an XML representation of the information graphic provided by the vision component of our system. APTE rules are applied to produce an effort estimate for each applicable rule (some tasks will not be able to be performed perceptually on a given graphic). The lowest effort instantiation is chosen for each task that can be performed. So for tasks like finding the exact value represented by the top of the bar where the task could be performed on any bar in the bar chart, the bar producing the lowest possible effort estimate will be chosen. If the bars are not annotated with values, this would be the bar with the shortest scan to the dependent axis. This is consistent with the idea that the graphic designer will make the important tasks easy to perform. The set of lowest effort tasks form part of the evidence used as input to the plan inference process, which can then chain to higher level goals whose operators contain one or more of these tasks as subgoals.

Additional Impact

We have also exploited the ideas put forth by cognitive psychologists in several other areas of our system. For example, designers of information graphics employ salience techniques to highlight particular items in the display; this can be done by coloring specific bars in a bar chart or by drawing an arrow to an element of a graphic. To account for salient entities in a graphic, we

made recourse to the work of Kosslyn [1994] who described the nerve cells in the visual system as 'difference detectors,' responding first to features of a display that are brighter, darker or otherwise different from their surroundings. Kosslyn's work suggests that when an element has been made salient (such as the red bar(s) in a bar chart), the viewer's eye is naturally drawn to that element before any information about the bar (for example, its label) is even known. To capture this natural perceptual behavior, we include any perceptual tasks that can be performed using the salient items of a graph and the effort estimates for those tasks generated by APTE in the set of tasks used to begin our plan inference process.

When performing plan inference, chaining among the operators produces a search space that is quite large; methods must be developed to guide the search. Several criteria come into play in evaluating a partial plan, where a partial plan consists of the tasks in the operators along a candidate path. One of the measures is the effort involved in performing a partial plan, which is estimated as the sum of the effort assigned to the component tasks. The proximity compatibility principle [Wickens and Carswell, 1995] also plays a vital role in evaluating partial plans. Since this principle posits that similarly encoded elements should be processed together, partial plans that use the similarly encoded evidence in an integrated fashion are rated higher than those that do not. For example, given an information graphic that contains two red bars, a plan that involves comparing the two red bars will be rated higher than a plan that entails finding the value of just one of the red bars. This reflects the fact that the first plan embodies the proximity compatibility principle by integrating the similarly encoded elements into a single task.

Conclusion

This paper presented a novel application of cognitive psychology research to the problem of recognizing the intended message underlying an information graphic. In future work, we will consider the impact on graph interpretation of the designer's beliefs about the knowl-

edge and skills of the intended audience, since individual differences have been shown to impact the graph comprehension process [Shah, 2002]. 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; the system will provide an initial summary containing the intended message of the graphic along with other important salient characteristics, and will respond to follow-up questions about the graphic's content.

References

- [ACT-R, 2003] ACT-R (2003). The ACT-R home page. <http://act-r.psy.cmu.edu/>, Retrieved on January 30th.
- [Card et al., 1983] Card, S. K., Moran, T. P., and Newell, A. (1983). *The Psychology of Human-Computer Interaction*. Lawrence Erlbaum Associates, Inc., Hillsdale, NJ.
- [Clark, 1996] Clark, H. (1996). *Using Language*. Cambridge University Press.
- [Cleveland, 1985] Cleveland, W. S. (1985). *The Elements of Graphing Data*. Chapman and Hall, New York.
- [Elzer et al., 2003] Elzer, S., Green, N., Carberry, S., and McCoy, K. (2003). Extending plan inference techniques to recognize intentions in information graphics. In *Proceedings of the Ninth International Conference on User Modeling*. To appear.
- [EPIC, 2003] EPIC (2003). The brain, cognition, and action laboratory: EPIC. <http://www.umich.edu/bcalab/epic.html>, retrieved on January 30th, 2003.
- [Grice, 1969] Grice, H. P. (1969). Utterer's meaning and intentions. *Philosophical Review*, pages 147–177.
- [John and Newell, 1990] John, B. E. and Newell, A. (1990). Toward an engineering model of stimulus-response compatibility. In Gilmore, R. W. and Reeve, T. G., editors, *Stimulus-response compatibility: An integrated approach*, pages 107–115. North-Holland, New York.
- [Kennel, 1996] Kennel, A. R. (1996). Audiograf: A diagram-reader for the blind. In *Second Annual ACM Conference on Assistive Technologies*, pages 51–56.
- [Kerpedjiev and Roth, 2000] Kerpedjiev, S. and Roth, S. F. (2000). Mapping communicative goals into conceptual tasks to generate graphics in discourse. In *Intelligent User Interfaces*, pages 157–164.
- [Kosslyn, 1994] Kosslyn, S. (1994). *Elements of Graph Design*. W. H. Freeman and Company.
- [Kosslyn, 1989] Kosslyn, S. M. (1989). Understanding charts and graphs. *Applied Cognitive Psychology*, 3:185–226.
- [Larkin and Simon, 1987] Larkin, J. H. and Simon, H. A. (1987). Why a diagram is (sometimes) worth a thousand words. *Cognitive Science*, 11:65–99.
- [Lohse, 1993] Lohse, G. L. (1993). A cognitive model for understanding graphical perception. *Human-Computer Interaction*, 8:353–388.
- [Meijer, 1992] Meijer, P. B. (1992). An experimental system for auditory image representations. *IEEE Transactions on Biomedical Engineering*, 39(2):112–121.
- [Peebles and Cheng, res] Peebles, D. and Cheng, P. C.-H. (in press). Modeling the effect of task and graphical representation on response latency in a graph reading task. *Human Factors*.
- [Perrault and Allen, 1980] Perrault, C. R. and Allen, J. F. (1980). A plan-based analysis of indirect speech acts. *American Journal of Computational Linguistics*, 6(3-4):167–182.
- [Pomerantz and Kubovy, 1986] Pomerantz, J. R. and Kubovy, M. (1986). Theoretical approaches to perceptual organization. In Boff, K. R., Kaufman, L., and Thomas, J. P., editors, *Handbook of Perception and Human Performance*, pages 36.1–36.46. Wiley, New York.
- [Russo, 1978] Russo, J. E. (1978). Adaptation of cognitive processes to eye movement systems. In Senders, J. W., Fisher, D. F., and Monty, R. A., editors, *Eye movements and higher psychological functions*. Lawrence Erlbaum Associates, Inc., Hillsdale, NJ.
- [Shah, 2002] Shah, P. (2002). Graph comprehension: The role of format, content, and individual differences. In Anderson, M., Meyer, M. B., and Olivier, P., editors, *Diagrammatic Representation and Reasoning*. Springer Verlag.
- [SOAR, 2003] SOAR (2003). The Soar home page. <http://ai.eecs.umich.edu/soar/main.html>, Retrieved January 30th, 2003.
- [Welford, 1973] Welford, A. T. (1973). Attention, strategy, and reaction time. In Kornblum, S., editor, *Attention and Performance IV*, pages 37–54. Academic, New York.
- [Wickens and Carswell, 1995] Wickens, C. D. and Carswell, C. M. (1995). The proximity compatibility principle: Its psychological foundation and relevance to display design. *Human Factors*, 37(3):473–494.
- [Wilensky, 1981] Wilensky, R. (1981). Meta-planning: Representing and using knowledge about planning in problem solving and natural language understanding. *Cognitive Science*, 5:197–233.