From: AAAI Technical Report WS-98-10. Compilation copyright © 1998, AAAI (www.aaai.org). All rights reserved.

The Impact of Perception on Agent Architectures

Micheal Hewett
Department of Computer Sciences
University of Texas at Austin
Austin, Texas 78712-1188
hewett@cs.utexas.edu

Abstract

The advanced state of agent software and computing hardware makes it possible to construct complex agents and robots with multiple streams of input such as vision, speech, gestures and data. Such agents, like people (who also have access to multiple input streams), need to effectively manage the input in order to process important information within useful time bounds. This paper discusses processes and architectural components that are used to manage input data.

In addition to reduced processing load, input management may also enable symbol grounding. However, some effects are not beneficial. For example, the agent will lack a full accounting of all input data, which means that standard explanation techniques will not function correctly. We propose several techniques for overcoming the disadvantages of input management.

1 The Evolution of Highly-Perceptive Agents

The first agents, in the form of knowledge-based systems, were invented over twenty years ago. The computing resources of that time were not capable of supporting high-bandwidth interaction so they relied on simple user interfaces, interacting in a very limited way with one person at a time.

Today we can build agents that are much more perceptually complex, capable of receiving input from many high-bandwidth sources at the same time. Various autonomous robots utilize vision, sonar and infrared sensors while communicating with other agents via wireless networks. The CMU Navlab (Pomerleau 1995) processes visual input and monitors vehicle controls in real time while driving a van at highway speeds. MIT's Intelligent Room (Coen 1998) processes visual, gestural and speech inputs from multiple people while displaying requested information on several displays in the room. ISI's helicopter pilot agents (Hill *et al.* 1997) fly helicopters while scouting for enemy tanks in a full scale battlefield simulation.

While these are impressive gains over systems developed a decade ago, they still have many limitations. The Intelligent Room uses twelve cameras for vision input, yet the test room required extensive modifications to its lighting and decor in order to prevent shadows from interfering with the visual processes. The helicopter pilot agents at ISI once participated in a simulation in which they came over a hill and encountered a field containing ninety tanks. The agents concentrated on processing visual information about the tanks. Meanwhile the helicopters drifted out of position because the helicopter gauges and controls were ignored. As a result, one of the helicopters crashed into a hillside while the other two helicopters crashed into each other. The need to correctly focus attention is just as prevalent in agent perception as it is in human perception. A chillingly similar scenario led to the crash of Eastern Flight 401 on December 29, 1972 near Miami, resulting in over one hundred deaths.

This paper discusses perception management for highly perceptive agents. There has been significant research into the problem of allocating *cognitive* resources to ongoing tasks; see, for example, (Hayes-Roth 1995). The importance of allocating *perceptual* resources has become more important recently, as the continuing advances in computer technology have made feasible advanced input modes such as vision and speech.

In order to process large quantities of input without producing problems such as those described above, we must design the agent's architecture to facilitate complex perception. If one were designing the first intelligent agent today, the specifications for it would almost certainly include the following:

- *network access to the World Wide Web as a source of knowledge*
- *access to other intelligent systems on a network*
- *a variety of methods for asynchronous interaction with people and other systems, including voice, text, gestural and visual input and output*
- *the ability to control external non-intelligent devices such as switches and motors*
- *to selectively incorporate data and results into the knowledge base (to prevent unbounded knowledge base expansion due to the increased bandwidth)*

This paper focuses on agents that have some or all of the above perceptual capabilities. The next section discusses attention mechanisms that help focus and filter perceptual activity. The following section discusses some ramifications of perceptual management, including the inability of standard explanation mechanisms to function in a highly perceptive environment. There are also some beneficial ramifications. One is that perception provides a way to ground symbols–which some say is a prerequisite for truly intelligent systems.

2 Attention-based Perception

In general, an input manager should be able to:

- *regulate the amount of information received over each input channel,*
- *prioritize the input according to dynamically changing preferences,*
- *provide a way for input to activate memory contents,*
- *handle unrecognized input,*
- *serialize parallel inputs when necessary,*

Since people are even more perceptually complex than today's agents, the characteristics of human perception have proved valuable when designing agent perception systems. The human brain is capable of filtering and processing hundreds of auditory, visual and tactile inputs per second (Potter 1993). Since similar sensors are likely to exist in an artificial intelligence, we can look to the human perceptual attention system for solutions to the input management problem.

Human Perceptual Attention

The methods and structures that people have evolved to handle perceptual input are:

- *Sensory buffers separate from regular memory and having a very short lifetime. According to (Klatzky 1980), evidence exists for perceptual storage areas. The visual and auditory buffers have been named the icon and the echo in (Neisser 1967). Perceptions may be replayed from the sensory buffers, but the inputs are not consciously processed until the information moves to short term memory.* ***Only a small percentage of sensory inputs are permanently retained.***

- *A short term memory distinct from long term memory. Short term memory has a small capacity, approximately seven plus or minus two (Miller 1956), and an exponentially decaying activation lifetime measured in seconds (Klatzky 1980).* ***People don't retain most of the data they receive or inferences that they make.***

- *An attention mechanism to filter and organize incoming data. Selective attention (Broadbent 1958; Treisman 1993) helps to organize and prioritize incoming data both within a modality (e.g. vision) and across modalities.* ***People process important data first and importance depends on the current situation.***

Visual Attention

The task of general computer vision is so computationally complex that it requires an attention mechanism (Tsotsos 1990). Many *visual attention* models (e.g. Ahmad 1991; Milanese *et al.* 1992; Mozer 1988; Sandon 1990) are based on a theory of human perception and attention known as Feature Integration Theory (Treisman 1993). In Feature Integration Theory, perception is divided into two levels, sub-symbolic and symbolic, with corresponding attention mechanisms, known as early and late attention respectively. The first stage of attention is *modal attenuation,* in which one modality (input channel) is given preferential treatment.

The second stage of attention is *feature registration.* This stage recognizes certain basic features of the input, for example the primary colors, contrast and verticality in vision; pitch and timbre in auditory input, etc. Feature registration is performed in parallel by low-level sensor and brain structures. The resulting features are placed into *feature maps* which contain the locations of recognized features.

In the third stage, feature maps are combined into *perceptual structures* representing objects in the external world[1]. A perceptual structure contains the features, along with their values and activation levels, that correspond to an external object. The activation level of a feature comes from the early attention mechanism.

If a new object enters the field of view, a new perceptual structure is created to store information about it. If an existing object moves, its existing perceptual structure is updated. In people, creation of perceptual structures is affected by attention. Under time pressure, people will create only a few important perceptual structures; the rest of the input will not be available for higher-level processing. However, the input may be processed later if recalled from the *visuo-spatial sketchpad* (Baddeley 1993), a replay loop that temporarily stores visuo-spatial information.

Perceptual structures bridge the sub-symbolic/symbolic gap. They are used as patterns to activate concepts from short term and long term memory. Concepts activated by perceptual inputs are placed in short term memory and

[1] Our "perceptual structures" correspond to Treisman's "object files", a term with a different meaning in Computer Science.

used for the current processing. This explains, for example, how seeing a set of keys can remind you of your own keys.

In addition to visual attention, a highly perceptive agent will undoubtedly need an attention module for each input modality. These will need to be cooperatively managed according to the current high-level goals, a process I call *coordinated multimodal attention*. The process of utilizing input from different input modalities at a high level is called *multimodal interaction* (Oviatt 1997).

Figure 1 shows an overview of a perceptive system. External input feeds into a perceptual management subsystem (the black box). From there, percepts move into the Sensor Buffer where they can trigger reactive behavior and/or activate knowledge and reasoning processes used in deeper cognitive behavior. The memory model in Figure 1 is based on a model presented in (Shiffrin and Atkinson 1969).

Characteristics of Attention-based Perception

The higher-level cognitive processes can control attention in several ways: modality selection, feature registration and selection, and semantic attention. Modality selection merely sets a preference for one modality over the others[2]. Feature selection causes a preference for certain features (colors, shapes, sounds, etc.). This is why it is more useful to tell someone "*Bring me the big,* ***orange*** *book with* ***black*** *letters*" rather than "*Bring me the AAAI-87 conference proceedings*".

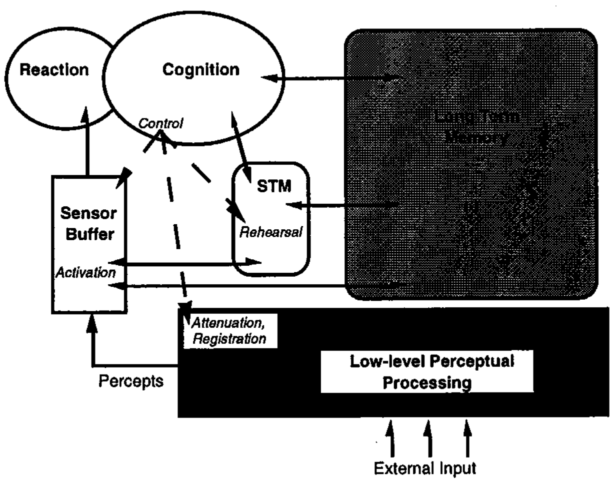

Fig. 1. Symbolic attention, reasoning and memory.

[2] Artificial intelligences may be able to implement a more complex modality preference system. It is as yet unclear whether some characteristics of human attention are limitations or useful features to artificial intelligences.

At the symbolic level, *semantic attention*, also called *late selection*, can either enhance or suppress the creation and/or activation of concepts. Enhancement is called *semantic priming*, while suppression is called *semantic interference*. The effects of semantic priming are relatively easy to detect. Subjects display faster recall of concepts that are semantically related to recently-encountered concepts. The classic illustration of semantic priming is described in (Neely 1977), which shows that conscious visual priming takes at least 700ms to happen and that people rely on sub-symbolic attention when less reaction time is available.

The classic account of *semantic interference* is the Stroop effect (Stroop 1935). In this experiment, subjects took longer to name the color of ink in which a word was printed if the word was the name of a different color, for example the word "yellow" printed in blue ink. The effect is also encountered if the subjects are to name the color of a rectangle but the name of another color is printed nearby.

In summary, human perception has evolved a multistage perceptual attention mechanism that covers both sub-symbolic and symbolic representation levels. The early stages of perceptual attention perform parallel registration of basic input features, while later stages integrate those features into percepts that act as indices into symbolic memory.

3 Architectural Implications

The general-purpose attention process described above combined with a high perceptual data rate will both provide some benefits and cause some problems for knowledge based systems. This section discusses the impact of the attention process on the structure and behavior of a typical knowledge based system.

Symbol grounding

One of the main tenets of AI is that symbolic reasoning forms a basis for intelligent reasoning. Harnad (Harnad 1993) points out that symbolic computation is the "*systematically interpretable manipulation of meaningless symbols*" and that "*symbol meanings must be grounded in something besides their interpretability if they are going to be candidates*" for intelligent reasoning. Harnad proposes that symbols can be grounded in their *sensorimotor projections*.

I propose that the perceptual structures generated by the perceptual attention process qualify as "sensorimotor projections" suitable for use in symbol grounding. A grounded instance such as FIRE-EXTINGUISHER-3 in the knowledge base would have a pointer to one or more perceptual structures that were created when it was sensed. Interpreting (i.e. retrieving the semantics of) the FIRE-

EXTINGUISHER-3 symbol might involve further processing of the perceptual structures. In this way, a semantic aspect of the symbol is retrieved from someplace other than the other symbols in the knowledge base, satisfying Harnad's requirement.

Storing and retrieving information from sensory images is a research area called computational imagery (Narayanan 1993). This is still a young research area and there is considerable debate as to the appropriate storage and intermediate data formats for performing imagery.

Separate short term and long term memory

If large amounts of data are continually asserted into the knowledge base, it will quickly bloat and become useless. One solution is to separate the storage areas for short term and long term memory. Data arrive in short term memory and then automatically decay away unless explicitly transferred into long term memory. This is the model used in this paper, but it is not the typical memory model used in knowledge based systems today.

Human memory exhibits characteristics of having separate short term and long term memories. The usefulness of this differentiation between the two types of memory is not completely known, although it can be postulated that short term memory functions much the same as a computer's memory cache; it provides fast access and temporal grouping. It is probable that the attention process described in this paper is best used with a two-level (or more) memory structure. Since most knowledge based systems have a monolithic memory, this would necessitate a reorganization of the general KB system model.

It is likely that early AI systems intended to incorporate short term and long term memory models similar to those that people use. However, by the time `OPS` (Brownston, *et al.* 1985) was implemented the distinction had been lost. The idea of separate memory areas was revived for `Soar` (Laird, Newell and Rosenbloom 1987). In `Soar`, all of the processing occurs in short term memory. Items must be specifically transferred into long term memory via a decision procedure. Although Prolog is not modeled after human memory, it also distinguishes between working memory and long term memory. Prolog does not retain intermediate results and stores information only when an assert operator is executed.

Reconstructive explanation and recall

Given the decay of short term data in our model, there will be very few situations in which all of the applicable sensor data is available to the reasoning process. For example, an agent may hear a voice behind it, but no longer remember that a person recently walked behind it. This will have an impact on the reasoning process, but even more so on the explanation process. How can a system explain events when information about them is fragmented? Once again, we find a possible answer in human cognition.

People, when faced with incomplete data, often use *reconstructive memory* to fill in the details. For example, if someone were asked what happened when they started their car a week ago, they probably couldn't remember all of the details. However, they could use the generic template (a.k.a. a *schema* or *case*) for starting a car and reply, *"I must have opened the door, sat down, put the key in the ignition, and started the car. Oh, and it was sunny that day so then I put on my sunglasses."* In this example, one known fact (that it was a sunny day) is supplemented by numerous details from a generic template.

The AI term for schema-based explanation and recall is case-based reasoning (CBR) (Riesbeck and Schank 1986; Leake 1996). CBR has proven to be a very capable method of reasoning. Its most difficult problem is *indexing*—selecting and organizing features that are used to retrieve cases from memory.

A modern CBR system contains a case library of abstract cases that are instantiated by relevant data. This method will work well in a reconstructive memory system.

Integrating perception and reasoning

One important question is how perceptual data are incorporated into the reasoning cycle of a knowledge based system. Data arrive asynchronously, but high level operations such as updating a Rete network or activating knowledge sources usually do not allow for asynchronous data. In systems that must do so, several mechanisms have been used. See (Hewett and Hayes-Roth 1989) for an overview and for a description of the method used in Guardian.

The mechanism used in the Guardian system (see Section 6) buffers incoming data until a safe point in the reasoning cycle is reached. It is then incorporated into the knowledge base all at once. No more than one action is performed between safe points, so much of the asynchronous nature of the perceptual input is retained.

Soar takes another approach by making input available as the value of a special frame slot. It is unclear whether unread data eventually disappears or whether it is buffered indefinitely.

4 Applying Perception-Based Architectures

The ultimate measure of the usefulness of an attention module will be how it aids a complex system perform its tasks. Let's look at several perceptually-complex systems to see how the attention module would help them.

The helicopter pilot agents mentioned in the first section of this paper were involved in a disaster because they improperly divided attention between two tasks: flying the helicopter; and observing enemy tanks. In our attention module, the contents of the SACK, which control perception, are derived from components of the current plan. If both tasks were elements of the current plan, then input from their respective sensors would be available for processing. While the attention module would guarantee that relevant percepts would be available, the cognitive component of the agent would still be responsible for ensuring that the sensor input would be processed. The attention module provides an organizational advantage, but is not a total solution.

A second example is MIT's Intelligent Room project (Coen 1997). The room contains a vision system that recognizes gestures, and a speech recognition system. The goal is to have a room, say in a weather bureau, where people can point at map displays and ask questions of the room's computer.

A recent video of the system shows a person gesturing at a map, waiting a few seconds, and then asking a question. Clearly the system is unable to process visual and audio input at the same time (as people can). To be fully functional, the room needs to be able to perceive multiple sources of input at the same time. There is no apparent reason, aside from having a limited implementation of a perception model, that the system could not process both input modes simultaneously.

A third application for a perceptual attention module would be an autonomous vehicle. Imagine a car driving down a street. Most autonomous vehicles concentrate on following the road, but there are many other important aspects to consider. A car backing down a driveway might pull out in front of you. A ball rolling into the street might be followed by a small child. The workers cutting a tree branch at the side of the road might drop it on the road in front of you. Each of these situations involves perceptions that activate reasoning processes or learned schemas. A perceptual attention module integrated with symbolic memory can handle these situations.

5 Related Work

Most of the previous work on handling large amounts of perceptual input has been in the area of real-time AI systems (Garvey and Lesser 1993). The solutions were developed as an approach to dealing with a real-world problem in a perceptually complex environment such as medical diagnosis (Hayes-Roth 1992), fighting forest fires (Howe *et al.* 1990) or pilot agents (Hill 1997).

One of the assumptions of the work described in this report is that *all* intelligent systems should be capable of existing in a perceptually complex environment and, to be useful, should respond reasonably quickly. Although this work doesn't set response-time goals and is therefore not a real-time system, it is clear that the organization of the system has much in common with real-time AI systems.

The Guardian patient-monitoring system

The Guardian system for intensive-care patient monitoring (Hayes-Roth *et al.* 1992) uses a buffer-based *perceptual preprocessing* system to handle the large amount of input it receives from patient monitors (Washington, Boureau and Hayes-Roth 1990). In Guardian, the perceptual system runs on a separate machine. It allocates a small sensory buffer to each input stream. The buffers act as filters and pass only significant information to the monitoring system. The monitoring system tells the perceptual system what "significant" means for each type of input data.

For a patient monitoring system, a significant value is one that goes outside the boundary of expected values. For example, a high or low temperature or blood pressure is significant. Since the buffers are small (they typically hold less than five items) they also implement data decay to make room for new data, and they have a timeout feature so that the monitoring system will occasionally receive a datum even if the input values stay within acceptable parameters. Finally, the boundary parameters are automatically adjusted by the perceptual system to maintain specific data rates. For example, the acceptable range may be widened in times of heavy input load.

This type of perceptual system works well in a monitoring environment. Attention is periodically shifted to each input data stream (via the timeout feature) while important data is immediately passed to the monitoring system (via the parameter range settings). And the perceptual system is self-adjusting in times of heavy load, which reduces the possibility of sensory overload.

However, it lacks some features needed in a general-purpose perceptual system. First, it is heavily dependent on numeric input data. Perceiving squares or colors does not utilize a parameter range with boundary values. Processing verbal input would also appear to be difficult to implement with numerically-bounded buffers. Perceptual preprocessing as used in Guardian does not appear to be a general solution to perceptual input management, but the perceptual preprocessing techniques could be used in a more general system to manage numeric input data.

The RANGER system

RANGER (Kelly 1994) is a system for managing visual attention during high-speed autonomous vehicle navigation. In high-speed navigation, detailed image processing is important when the environment is dynamic and unpredictable. However, detailed image processing becomes untenable at high vehicle speeds due to the impossibly short processing times available.

The RANGER system solves this by computing the relevant sub-window of the visual input that is important at the current speed. Objects below the sub-window are unimportant because the vehicle does not have time to react to them anyway. Objects above the window are not on the ground and thus do not need to be avoided. Using the window, the amount of processing required is significantly reduced, allowing the system to safely navigate at higher speeds.

RANGER is not knowledge based; it is based entirely on mathematical equations involving the vehicle speed, sensor range, and the angle between the camera and the ground. However, it can be seen as a special case of attention where a spatial map is segmented and used to govern the creation and update of perceptual structures for only those objects in the relevant segment of the spatial map.

6 Summary and Future Work

This paper has discussed the use of perceptual attention for input management. Perceptual attention helps to organize, filter, prioritize and serialize input received from any number of sensors. It can be integrated with the cognitive and memory systems of a complex artificial agent. In addition to its benefits, such as symbol grounding, it has some drawbacks. This paper has explored several ways of overcoming limitations imposed by the attention processes.

As we continue to develop perceptual attention and apply it to our project, an autonomous wheelchair, we intend to explore several areas mentioned in this paper. One is the question of whether characteristics of human perception and attention are features or limitations. Should artificial intelligences be subject to the same limitations as people? Is the size of human short term memory relevant to that of artificial agents? Should the inability of people to attend to two conversations at the same time be incorporated into artificial conversationalists?

We will also look into the various ways of controlling the low-level (sub-symbolic) perception mechanism via the control lines to the cognitive component. There may be some useful control regimes that are not present in human perception. For example, people do not have direct access to textual input (written words are recognized at the cognitive level) but computers could receive direct textual input or symbolic input through a knowledge sharing language.

Acknowledgements

Thanks to Ben Kuipers and the UTCS Autonomous Robots group for valuable discussions related to this topic.

References

S. Ahmad, *VISIT: An Efficient Computational Model of Human Visual Attention*. Ph.D. Thesis, University of Illinois at Urbana-Champaign. Also as Technical Report TR-91-049, International Computer Science Institute, Berkeley, CA, 1991.

Anderson, J. R. 1984. Spreading Activation, in: *Tutorials in Learning and Memory: Essays in honor of Gordon Bower*, J. R. Anderson and S. M. Kosslyn, editors. W.H. Freeman and Company, San Francisco.

Baddeley, A. 1993. Working memory or working attention?, in *Attention: Selection, Awareness and Control*, A. Baddeley and L. Weiskrantz, eds. Clarendon Press, Oxford.

Broadbent, D. E. 1958. *Perception and communication*. Pergamon Press, Oxford.

Brownston, L., R. Farrell, E. Kant and N. Martin 1985. *Programming Expert Systems in OPS5: An Introduction to Rule-Based Programming*, Addison-Wesley, Reading, MA.

Coen, M. H. 1998. Design Principles for Intelligent Environments. In: *Proceedings of the AAAI 1998 Spring Symposium on Intelligent Environments*, AAAI Press.

Garvey, A. and V. Lesser 1983. *A Survey of Research in Deliberative Real-Time Artificial Intelligence*. UMass Computer Science Technical Report 93-84, Department of Computer Science, University of Massachusetts, Amherst.

Gribble, W. S., R. L. Browning and B. J. Kuipers 1997. *Dynamic binding of visual percepts for robot control*. Unpublished technical report, Department of Computer Sciences, University of Texas, Austin, TX.

Harnad, S. 1993. Symbol Grounding is an Empirical Problem: Neural Nets are Just a Candidate Component. Proceedings of The *Fifteenth Annual Meeting of the Cognitive Science Society*. NJ: Erlbaum.

Hayes-Roth, B. 1995. An architecture for adaptive intelligent systems, *Artificial Intelligence*, **72**(1-2):329-365.

Hayes-Roth, B., R. Washington, D. Ash, A. Collinot, A. Vina and A. Seiver 1992. A prototype intensive-care monitoring agent. *Artificial Intelligence in medicine*, **4**:165-185, 1992.

Hewett, M. and B. Hayes-Roth 1989. Real-time I/O in Knowledge-Based Systems. In: *Blackboard Architectures and Applications*, Jagannathan, V., R. Dodhiawala and L. Baum, Eds., Academic Press: Boston.

Hewett, M. and B. J. Kuipers 1998. *Helping Perceptive Agents Pay Attention*, submitted to AAAI-98.

Hill, R. W. Jr., J. Chen, J. Gratch, P. Rosenbloom and M. Tambe 1997. Intelligent Agents for the Synthetic Battlefield: A Company of Rotary Wing Aircraft. In: *Proceedings of AAAI-97*, Providence, Rhode Island.

Howe, A. E., D. M. Hart and P. R. Cohen 1990. Addressing real-time constraints in the design of autonomous agents. *The Journal of Real-Time Systems*, **2**(1/2):81-97, 1990.

Kelly, A. 1994. *Adaptive Perception for Autonomous Vehicles*, Technical Report CMU-RI-TR-94-18, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA.

Klatzky, R. L. 1980. *Human Memory: Structure and Processes.* W. H. Freeman and Company, New York.

Laird, J., A. Newell and P.S. Rosenbloom 1987. Soar: An architecture for general intelligence. *Artificial Intelligence* **33**(1):1-64.

Leake, D., ed. 1996. *Case-Based Reasoning: Experiences, Lessons, & Future Directions*, MIT Press, Cambridge, Massachusetts.

Milanese, R., J.-M. Bost and T. Pun 1992. A Bottom-up Attention System for Active Vision. In *Proceedings of the 10th European Conference on Artificial Intelligence*, pp. 808-810. John Wiley and Sons.

Miller, G. A. 1956. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review* **63**:81-97.

Mozer, M. C. *A Connectionist Model of Selective Attention in Visual Perception.* University of Toronto, Technical Report CRG-TR-88-4, 1988.

Narayanan, N. H., ed. (1993). The Imagery Debate Revisited, special issue of *Computational Intelligence* **9**(4), November 1993.

Neely, J. H. 1977. Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General*, **106**:226-254.

Neisser, U. 1967. *Cognitive Psychology.* New York: Appleton-Century-Crofts.

NTIS 1973. *Eastern Air Lines, Inc., L-1011, N310EA, Miami, Florida, December 29, 1972.* NTIS Report Number 19B-222359/2 and NTSB Report Number NTSB-AAR-73-14. See also [WWW document] URL:
http://www.d-n-a.net/users/dnetGOjg/291272.txt

Oviatt, S. 1997. Multimodal interactive maps: Designing for Human Performance. *Human-Computer Interaction* **12**:93-129. Lawrence Erlbaum Associates, Inc.

Pomerleau, D. 1995. RALPH: Rapidly Adapting Lateral Position Handler, in *Proceedings of the IEEE Symposium on Intelligent Vehicles*, Detroit, MI.

Potter, M. C. 1993. Very short-term conceptual memory, *Memory & Cognition* **21**(2):156-161.

Riesbeck, C. K. and R. C. Schank (1989). *Inside case-based reasoning.* L. Erlbaum Associates, Hillsdale, N.J.

Sandon, P. A. 1990. Simulating Visual Attention, *Journal of Cognitive Neuroscience*, **2**(3):213-231.

Shiffrin, R. M. and R. C. Atkinson 1969. Storage and Retrieval Processes in Long-Term Memory. *Psychological Review* **76**:179-193.

Stroop, J. R. 1935. Studies of interference in serial verbal reactions. *Journal of Experimental Psychology* **18**:643-662.

Treisman, A. 1993. Features and Objects, in *Attention: Selection, Awareness and Control*, A. Baddeley and L. Weiskrantz, eds. Clarendon Press, Oxford.

Treisman, A. M. 1969. Strategies and models of selective attention. *Psychology Review* **76**:282-99.

Washington, R., L. Boureau, and B. Hayes-Roth 1990. *Using Knowledge for Real-Time Input Data Management.* Stanford KSL Technical Report KSL-90-14, Knowledge Systems Laboratory, Stanford University.