

SITUATIONAL AWARENESS IN THE COMMERCIAL AIRCRAFT COCKPIT: A COGNITIVE PERSPECTIVE

Marilyn Jager Adams & Richard W. Pew
Bolt Beranek and Newman Inc
Cambridge, Massachusetts 02138

Abstract

Cognitive theory is presented that has relevance for the definition and assessment of situational awareness in the cockpit. The theory asserts that maintenance of situation awareness is a constructive process that demands mental resources in competition with on-going task performance. Implications of this perspective for assessing and improving situational awareness are discussed.

Introduction

In the design of safe, functional, and productive person-machine systems, situational awareness is an issue of central concern. The potential difficulty is vividly present in the conventional cockpit of commercial aircraft. In the course of managing the navigation, guidance, and control of the aircraft, pilots are simultaneously responsible for the information presented across up to 400 separate gauges and instruments. They are, moreover, responsible for the interpretation of that information and, as necessary, the inferences involved in identifying the root cause of abnormalities and choosing among some set of compensatory or remedial actions. An interruption, an oversight, a hasty inference, or a decision based on incomplete knowledge or information: under conditions of heavy workload or tight temporal pressure -- any crew is vulnerable to each -- could mean disaster.

From this view, the installation of "glass cockpits" in newer commercial aircraft represents a major technological initiative toward managing human workload. In glass cockpits, a half dozen cathode ray tubes replace the hundreds of gauges and instruments while flight management computers manage information presentation and take control of many functions previously assigned to the aircrew. Thus, the memory, discipline, and cognitive effort required for wisely distributing attention around the cockpit is greatly reduced and so, too, is much of the crew's inferential and decision-making onus: When a malfunction occurs, the computer should autonomously correct it or present an identification of the problem to the crew along with a checklist of actions to be taken.

Yet, even now, so soon after their introduction, the design and dynamics of these high-tech cockpits have been questioned. "With old cockpits," explained one pilot, "the workload was high but you were always aware of what's going on." In contrast, the "glass cockpits" are criticized for psychologically distancing pilots from their aircraft and environment. In effect, as the high-tech cockpits have reduced hands-on and eyes-on demands of flight management, they have not eliminated

the workload problem, but shifted its emphasis from one of having too many things to keep track of and do to one of requiring more discipline and thought to oversee functions that have been automated; from one of finding ways of avoiding activity overload to one of supporting the crew's *situation awareness* -- their moment by moment ability to monitor and understand the state of the aircraft, its systems, and its environment.

It is a central thesis of this paper that achieving situational awareness requires active attentional and inferential processes and involves significant perceptual and cognitive resources. As such, acquiring and maintaining situational awareness must be appreciated as an integral part of the crew's mental workload.

Although this discussion is directed to the process of achieving situational awareness, its goal is to establish means of assessing and improving the product of situational awareness: having the right kind of information in one's head when needed. Because the process of achieving situational awareness is constructive--that is, because it depends on expectations, hypotheses, and basic familiarity with the presenting situation--its ease and successfulness both determines and is determined by its product. For example, because the expert pilot differs from the novice in both knowledge and experience, his or her ability to anticipate needs and engage in information-processing activities are both better tuned and better managed. In short, while recognizing their interdependence, we distinguish the process of achieving situational awareness from its product; if we wish to measure and improve the product, we must understand the mental dynamics and constraints through which it is achieved.

In the sections that follow we present an overview of the cognitive processes involved in the construction of situational awareness. As a part of a NASA project, this work is directed toward developing methods and means to enable objective and informative assessment of it in the context of commercial aircraft operations. In later reports we will focus on the measurement procedures per se.

What is Situational Awareness?

In order to respond appropriately to any given situation and to anticipate the next, Nagel suggests, "the pilot must maintain awareness, for example, of where the aircraft is geographically, what its current performance state is, its altitude, power setting, configuration, and relationship to other aircraft. The pilot must (1) be mindful of verbal communications and instructions from air traffic control personnel, (2) be aware of the environment, and (3) systematically monitor for and an-

ticipate changes in weather that might affect his or her performance and that of the aircraft. Information to maintain an appropriate awareness of the flight situation is obtained visually, through the windows of the aircraft and from instruments, aurally from other crew members within the cockpit and from air traffic controllers over radio channels, and through the vestibular senses.¹¹

The foregoing is a top-level and, therefore, oversimplified view of pilots' in-flight management responsibilities. Despite that, it points out several salient characteristics of their mental processing environment. First, there are, at any point in time, a number of different tasks demanding attention. Second, each of these tasks can be both knowledge intensive and procedurally complex. Third, the demands for the various tasks and the information each requires for execution do not arrive in coherent, task-by-task packages but are interleaved in time such that the urgency of executing one or more tasks is liable to peak at the very moment when information triggering, enabling, or urging completion of others is arriving. Fourth, the relevance of any available information is not always consistently signaled either by sources or by sensory modality; it depends, in addition, on its significance to some particular aspect of the overall flight situation and, thus, may be identifiable only to the extent that the pilot recognizes that significance. The pilot's ability to attend selectively to information and to interpret it properly depends, in other words, on her or his ability to call forth the local and long-term knowledge and understanding on which such judgments must be based.^{2,4}

In our study of *situational awareness*, these four characteristics of the pilot's mental processing environment have been given central importance. Additionally we would like to clarify one other aspect of the term as we will use it. In both everyday usage and in the technical literature, the words *situational awareness* are alternately used to define the *ideal* and the *actual*. While the *ideal* corresponds to the externally available information and internal model of the flight with which the pilot *should* work (or, in accident analyses, *should have* worked), the *actual* corresponds to the externally available information and internal model of the world with which a pilot actually *does* (or apparently *did*) work.

We have made a rule of using the term *situational awareness only* with reference to the actual active or working knowledge of the pilot. Beyond respecting the fact that *awareness* belongs to people and not the environment, this usage was motivated by research-oriented considerations. Most critically, in stipulating that it is the *working* knowledge of the pilot that is of interest, we acknowledge that at any moment in time, a person's effective or active knowledge base consists of but a fraction of her or his total knowledge. Because that active subset of knowledge shifts with attention and inasmuch as it determines the thoughts and conditions that are readily available to consciousness, it strongly

influences a person's potential awareness and interpretation of any given environmental event. Moreover, by separating the pilot's actual awareness about a situation from potential or ideal awareness, we create a space for assessing and toward supporting situational awareness. What are the differences? What do they indicate about the capacity and dynamics of human attention and information-processing relative to the demands of flight management and current cockpit technologies? And how can such understanding of the capacity and dynamics of human attention and information-processing be translated into cockpit design so as to bring actual situational awareness as close as possible to the ideal?

Situational Awareness and Workload

From our perspective, achieving situational awareness is but a special class of mental workload. While there is a large literature on the *time-sharing* aspects of workload, there is relatively little empirical literature on how people *manage* more than two concurrent tasks; second, and more importantly, there is even less on how people cope with concurrent tasks that are cognitively complex.

Even in the domain of single tasks, workload management is a complex problem, extending across all phases of processing from perception and interpretation through response selection and execution. For any single task, person, and situation, however, the components of the total workload involved stand as synergistic states and subprocesses of the same goal. The case of multiple, simultaneously demanding tasks is very different. Ensuring that each individual task is manageable in isolation in no way guarantees the concurrent manageability of the ensemble. And even having demonstrated that difficulties exist, the identification of their cause and cure is vastly complicated.

Some of these complications have to do with the competition between tasks for the operator's services. It is obvious that human operators have limited physical attributes--one pair of eyes, one mouth, two hands, etc.--and any multitask design must be expected to take these limitations into account. It may be equally obvious that human operators have limited attentional and cognitive resources, but the considerations involved in designing around these constraints are central to achieving situational awareness and are far more subtle.

Beyond direct conflicts between tasks, there are issues that arise from the very fact that the operator has more than one on the agenda. It has been shown that even when an operator is faced with as few as two tasks and the "tasks" consist of nothing more than the detection or recognition of simple signals, the requirement to divide or switch attention between them may result in a significant loss in sensitivity or time that can be allocated to either.^{5,6}

The Memory Demands of Cognitively Complex Tasks

By cognitively complex, we mean tasks that interact and involve stimulus or response uncertainty of the sort that requires interpretive reasoning, situation assess-

ment, or thoughtful response selection—that is, we mean precisely those types of tasks for which it is generally deemed necessary to keep the human in the loop. A characteristic of cognitively complex tasks is that their execution requires the operator to access a mental model of the situation replete with knowledge of the history of past experiences, the specifics of the most recent occurrences, procedural policies, special conditions of the aircraft or flight situation, best bets, and caveats. The information in this model must be reconciled with environmental information and be consistent with the models built into the system's aiding procedures as well as those governing all broader concerns of protocol and operation.

To maintain adequate situational awareness, the pilot is obliged to determine the relevance, procedural implications, and urgency of each incoming datum or event as it occurs. Yet, there are two major classes of difficulties associated with the successfulness of this endeavor. One of these classes of difficulties derives from the fact that the temporal arrival of real-world information is unpredictable. Whenever an event occurs, the operator may be obliged to interrupt any ongoing activity at least to assess its significance—we will return to this problem later.

At present we focus on the second class of difficulties: The significance of any given event can be assessed only by accessing background knowledge of the conditions that it signals and the operations to which it pertains. It can be assessed, in other words, only by activating one's memories and understandings of the situations and activities.⁷ However, and this is extremely important, thoughtful attention is modular. People can consciously think about only one thing at a time.

The Role of Long-Term Memory in Perception, Interpretation, and Action Selection. Empirical and theoretical work on people's ability to learn and interpret extended information indicates that human long-term memory is nothing like a piecemeal catalog of information. Instead, the substance of each object, concept, or event is reduced in memory to an extended array of primitive units while its structure is preserved through interconnections among them. More specifically, the interconnections capture all of the various relations among the constituents of the experience that one has learned, noticed, or understood, consciously or unconsciously.⁸

Within current learning theory, much of the power of these memory structures derives from three basic assumptions.⁹⁻¹¹ The first of these is that the primitive units are not duplicated and are relatively small in number. In this way, the representation of any given concept is essentially hierarchical, and the whole of any idea or percept—however complicated it may be—is defined by its interrelated parts. The second is that these long-term memory structures serve not only as the internal records of one's experiences but also as the

medium for perceiving and (thereby) interpreting experiences. The third is that the connections that hold the units and clusters of units together in memory serve also to pass excitation and inhibition among them. More specifically, the more frequently any set of units has been associated in past experience, the more positively the activation of one will promote the activation of the others.

In themselves, these three assumptions go a long way toward explaining how it is possible for people to interpret cognitively complex situations. Specifically, if the perception of a signal or event consists of activating the representations of its familiar parts in memory, and if the directly activated parts in turn evoke their most frequent associates, then it is to be expected that any highly familiar and categorically significant signal or symptom will—in the course of its perception—call forth related background information, contingencies, exceptions, and conditional responses. This is the essence of situational assessment.

Of note, these assumptions also explain the continuum on which people respond to situations as a function of their novelty. A highly familiar array of information may pass through interpretation to response with near automaticity. A novel situation that is piecewise familiar may evoke a number of piecewise responses -- that may or may not be compatible with one another. A wholly novel signal or situation may evoke none at all. More generally, to the degree that a situation is novel, the operator is obliged to construct rather than recall its significance and appropriate response; to do so requires conscious, methodical, time-consuming, and attentionally demanding thought. The Danish engineer, Jens Rasmussen, has documented this relationship in a variety of industrial scenarios: The situational continuum of familiar-to-novel translates into a response continuum of automaticity and success versus effortfulness and error. The point is that it does so as the direct result of our mental architecture.¹²

The Role of Active Memory in Directing Attention. From a theoretical perspective, the major problem with this conception of human information-processing is that it presumes conscious access to bodies of information that are potentially very large and highly complex. In particular, the significance of an event in real-world situations requires full sensitivity to the immediate context in which it occurs as well as the time-marked stream of events that preceded it. Yet the capacity of human short-term memory is classically held to be to 7 ± 2 items at a time, and herein lies a paradox.¹³ The resolution to this paradox comes with the recognition that, in the experimental measurement of short-term memory, the to-be-remembered items are typically physically discrete, semantically unrelated, and often meaningless. Although the number of "items" that are simultaneously considered in a complex and familiar, real-world scenario may be very large, they have, as

described above, become richly interassociated and mutually supportive through experience.

Through experimentation in the domain of text comprehension, Sanford and Garrod have theorized that an individual's active memory consists of two bins, *explicit* and *implicit* focus.¹⁴ *Explicit focus* corresponds roughly to what is conventionally labeled as the "short-term store." It is the working space within which readers integrate the new information from the text with information that they know or have read about already. At any given moment, explicit focus contains a tightly-limited number of interrelated tokens of (or pointers to) larger knowledge structures in long-term memory. Although the contents of explicit focus are regulated more or less like a push-down stack, its maintenance of any given token depends not only on the recency with which it has been activated by the text but also on its implicit relevance to the current interpretive stream. As an example from the text processing domain, the prominence of a character in explicit focus may be boosted either by mention (direct or indirect) or because, though unmentioned, the character is of central relevance to the event being described.

Implicit focus, in contrast, subtends the full-blown representation of the situation that is partially represented in explicit focus. Information relevant to the knowledge in implicit focus can be brought to the interpreter's attention with neither the speed nor the obliqueness of reference that suffices for information in explicit focus. On the other hand, such information can be interpreted far more quickly and with far less support than information that is unrelated to the contents of explicit focus.

To support these active memories, Sanford and Garrod suggest that the reader's latent (currently inactive) memory is also sectioned into two bins. The first, *long-term episodic memory*, contains a complete record of the knowledge structures that have been built or accessed in the course of reading the current text--or, for our purposes, in the context of the current flight. Meanwhile, *long-term semantic memory* contains a person's lifetime accumulation of knowledge in general. Knowledge in either of these latent memories can be brought to consciousness only given considerable effort or strong cueing.

Extrapolating from the literature on text comprehension, we can anticipate some of the parameters that will control the ease or probability with which a given event is properly processed by the manager of real-world information. Events that are relevant to those aspects of a task on which a person is currently working should be readily assimilated because they will map themselves onto the knowledge currently in explicit focus. Thus, for example, the pilot will readily notice and respond to changes in glideslope indication that he is tracking in the course of landing. Events that pertain to the task but not to the particular aspect of the task with which a person is engaged are also expected to be inter-

preted relatively quickly and cogently as they will map onto knowledge in implicit focus. Thus, for example, even while tracking the glideslope, the pilot may be readily alerted to changes in engine noise that are consistent or inconsistent with landing experience. In contrast, when the interpretation of an event requires consideration of knowledge in latent memory, the probability or effort associated with its proper processing will depend on such factors as the transparency of its significance and the time available for working on it: When very close to touchdown, for example, the pilot will be relatively unprepared to receive and interpret unrelated communications.

While we may borrow the base theories from the literature on text comprehension, there are issues related to the structure and dynamics of processing real-world scenarios that require substantial modification and elaboration of those theories. In particular, whereas the episode structures of text are well planned and organized by their authors, the arrival of real-world information may follow no organizational principles save order of occurrence or discovery. Typically, some portion of the available, real-world data is relatively unimportant or even irrelevant. For another portion, a response is urgently required. For the rest, the procedural implications are important but deferrable within limits--and those limits must also be computed and registered.

In the real-world, moreover, it's up to the person to figure out the relevance of any bit of data. Yet, perceived significance depends entirely on how it is interpreted, and the very nature of complex dynamic scenarios render the operator highly vulnerable to errors of misinterpretation. At any point in time, the pilot may be occupied with such tasks as resolving contradictions; inducing large fragments of missing structure; monitoring, adapting and correcting the hypotheses thus induced; or selecting and initiating action alternatives and monitoring their consequences. Each such activity may involve significant processing effort and must require the presence of its own supporting knowledge in explicit focus. As a consequence, each such activity must be expected to distort, diminish, or even block one's capacity to attend to unrelated issues and events. Alternatively, when external events do succeed in interrupting an on-going train of thought, proper completion of the latter is threatened.

An example of these vulnerabilities is provided by a recent incident at the Detroit Metropolitan Airport.¹⁵ The aircraft took off without lowering its flaps and crashed. Although the crew had begun the pre-flight checklist properly, they were interrupted by Air Traffic Control before verifying the status of the flaps. Although they might still have resumed the checklist routine prior to take-off, other issues usurped their attention: They were confused as to which taxi-way to use, the runway direction had just been changed, and weather and runway conditions were not provided until the taxi itself. Although, with proper handling, the air-

craft could have become airborne without flaps, the crew had been given a windshear alert: When the problem with the flaps expressed itself during take-off, the symptoms were interpreted -- and responded to -- as though they were caused by windshear.

The Memory Requirements of Coordinating Multiple Tasks. Cognitive activity is modular. Though we may be able to do several things at once, we can devote thoughtful, conscious attention to only one at a time. So here is yet another problem: the crew of a complex system is not merely required to shift attention among tasks, they must also perform each in subordination to the broader constraints of the system and its mission.

The implications are, first, that the management of multiple cognitively complex tasks must consist essentially in working on one while queuing some number of others. Second, the queue of to-be-attended tasks cannot be worked through any simple first-in first-out heuristic. Instead, the tasks in the queue must be prioritized with deference to both the temporal requirements on their execution and their overall importance to the management of the situation as a whole. Within actual systems, moreover, the nominal set of tasks in the queue as well as their relative priorities change dynamically as a function of events and of changes in the status of the subsystems involved. Grappling with these issues is forcing us to recognize a variety of questions that are in dire need of research.

We know that memory is limited. List maintenance is effortful and fallible; more so, if the list must be ordered; even more, if the membership of the list must be dynamically reordered and modified during retention.¹⁶ It follows that maintenance of such a queue would require considerable cognitive effort even if it consisted of nothing more than a list of things to do. Yet the queue cannot be productively conceived as a list: Like the tokens in explicit focus, those in the queue must correspond to pointers to knowledge structures in memory--structures that detail the procedural and declarative knowledge about each task and that must be accessed in its prioritization, reprioritization, scheduling, and execution.

Because, more than anything else, it may be maintenance and prioritization of this queue that must determine the pilot's capacity to respond appropriately to the individual demands of the scenario, it is worth examining its cognitive requirements more closely. For purposes of discussion, these requirements can be divided into two (nonindependent) subsets: (1) How does the pilot prioritize the pending tasks? and (2) What are the factors that determine when she or he will shift attention to any particular task in the queue?

The pilot's ability to update or reset task priorities appropriately depends on her or his ability to direct attention-- to notice events, to keep track of systems status, and to evaluate the requisite variables in a timely and appropriate manner. Some events may pass entirely unnoticed by the aircrew. In the interest of supporting

situational awareness, models for predicting and methods for minimizing such oversights are of obvious importance. However, the cognitive costs of noticing events should be of equal concern.

To notice the occurrence of an event in any useful way, the pilot must immediately interrupt ongoing activities, at least to evaluate its significance and establish the priority of its response implications. Resumption of the interrupted task must require thoughtful review of its status and may require repetition or reinitiation of one or more of its procedural components. Thus, the very reception of unanticipated data must always introduce an additional and disruptive element of workload. The design implications, especially for time-critical systems, should not be ignored.

Beyond that, research indicates that mental shifts between topics or semantic domains require measurable time and effort and are prone toward certain classes of biases and errors.¹⁷⁻²⁰ To the extent that incoming information is unrelated to the task in which the pilot is concurrently engaged, its interpretation must involve considerable mental workload. The more time and effort the pilot invests in its interpretation, the greater must be both its disruptiveness to ongoing activities and its potential for blocking notice or proper interpretation of other available data. The less time and effort he invests in its interpretation, the greater the likelihood of misconstruing the nature and priority of its implications.

The issues of how the operator will shift attention between tasks can also be broken into subsets. First, there is the question of *when* attention will be shifted. An answer to this question necessarily involves enumeration of the conditions that will disincline the operator towards shifts, and chief among these must be the intensity with which the operator is attending to the task at hand. Second, it must also involve enumeration of the conditions that will predispose a shift. Most obvious among these are superficially salient events, including system alarms. In addition, theory suggests that the operator's attention will be more easily captured on completion of goals and subgoals, when the current contents of explicit focus are being closed up and replaced.²¹ Yet another possibility is that tasks on the to-be-serviced queue elicit attention in proportion to their computed urgency or criticality. As adaptive as this possibility sounds, it too carries a danger: In moments of greatest temporal stress--i.e., in moments where more than one task is urgently demanding attention--pilots may be least able to concentrate on and, thus, least efficient at completing any one.

Coupled with questions of when the pilot will shift attention are questions of *which* task the operator will shift attention. Whenever the relative priority of any single task is distinctly higher than that of any other, the task selection process should be trivial; the pilot should immediately and without cogitation turn directly to that highest ranked task. Yet it is less clear

what will happen when two or more are close in value and clustered at the top of the stack.

Clearly the pilot is responsible for auditing the status of the tasks in the queue so as to respond to each task, ideally, at the most efficient possible moment or, minimally, before it is too late. But when the priorities are very high or equivalently, the constraints are very tight on two tasks at once, the pilot is faced with a dilemma: To attend to either task is to defer the other. The only escape from this dilemma would seem to involve reviewing the constraints on each task in search of a method for compromising some subset of them in a way that best protects the overall integrity of the flight. But note: the implication of pursuing this escape is that the amount of workload required for optimal task selection will be greatest when the amount of time available for it is least.

Assessing Situational Awareness

In order to measure situational awareness and to evaluate the ways in which it is supported or impeded by particular interface designs or situations an operational definition is required. The definition must allow assessment in ways that are methodical and reliable and yield outcomes that are accessible to system analysts, designers and training specialists.

As a first step, we divide the pilot's activities into top-level *flight management categories*, as shown in Table 1.

Table 1. Flight Management Categories
1. Local Navigation, Guidance, and Control of the Aircraft
2. Macro-planning and -Navigation, including route planning, replanning and programming.
3. Remote Communications
4. Flight Crew Resource Management
5. Cabin Management
6. Management of Physical Equipment, Resources and Systems
7. Management of Flight Management Computer and Related Systems
8. Bridging Activities (activities not directly related to flight management)

Because the categories are to correspond to frames of memory and thought, their contents are defined by flight management goals rather than by types of equipment or procedures.²² As an example, the pilot's use of the Flight Management Computer in the service of reprogramming a descent belongs to category 2 and is distinctly different from attention in the service of monitoring current flight status (category 1). In contrast, episodes requiring attention to be devoted to the computer per se (as to figure out how to interact with it or to trouble-shoot it), belong to category 7. Indeed, whenever (whether due to poor interface, programming

difficulty, slow response time, bugs, etc.) the computer becomes, in itself, the object of attention rather than a means for attending to flight-integral variables and procedures, it must be considered diversionary. Best use of such analysis depends on pushing further into the hierarchies of goals that underlie flight management. Even at this very highest level analysis, however, we hope we have illustrated how these categories may usefully support the measurement of situational awareness.

We are currently devising experimental paradigms and measurement parameters to enhance our understanding of the dynamics and limitations of situational awareness as well as objective, replicable techniques for measuring it.

Conclusions

The goal of inserting advanced technology into any system is that it result in an increase in the effectiveness, timeliness, and safety with which the system's activities can be accomplished. Intentions notwithstanding, the inherent difficulties of the multitask situation are very often compounded by the introduction of automation. To maximize situational awareness, the dynamics and capabilities of such technologies must be designed with thorough respect for the dynamics and capabilities of human information-processing.

Acknowledgments

This research is being supported by the NASA-Langley Research Center (Contract No. NAS1-18788). We thank Dr. Randall Harris, our task monitor at NASA-LARC, as well as William Rogers and Sheldon Baron for helpful comments on the manuscript.

References

- [1] Nagel, DC (1988) Human error in aviation operations. In EL Weiner & DC Nagel (Eds.), *Human factors in aviation*. New York: Academic Press, p. 215.
- [2] Endsley, MR (1988). A methodology for the objective measurement of situation awareness. *AGARD Symposium on Situation Awareness in Aerospace Operations*. Copenhagen.
- [3] Regal, DM, Rogers, WH, & Boucek, GP (1988). Situational awareness in the commercial flight deck. *7th Aerospace Behavioral Technology Conference and Exposition*. Warrendale, PA: Society of Automotive Engineers, 65-69.
- [4] Wickens, CD, & Flach, JM (1988). Information-processing. In E. L. Weiner & D. C. Nagel (Eds.), *Human factors in aviation*. New York: Academic Press.
- [5] Swets, JA (1984). Mathematical models of attention. In R Parasuraman & R Davies (Eds.), *Varieties of attention*. New York: Academic Press.
- [6] Broadbent, DE (1957). A mechanical model for human attention and immediate memory. *Psychological Review*, 64, 205-215.

- [7] Fracker, ML (1988). A theory of situation assessment. *Proceedings of the 32nd Annual Meeting of the Human Factors Society*, Anaheim, CA.
- [8] Adams, MJ (1989) *Beginning to read: Thinking and learning about print*. Cambridge, MA: MIT Press.
- [9] Hintzman, DL (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, 93, 411-428.
- [10] Rumelhart, DE, & McClelland, JL (1986) *Parallel distributed processing, Vol 1*. Cambridge, MA: MIT Press.
- [11] MacKay, DG (1987). *The organization of perception and action*. New York: Springer-Verlag.
- [12] Rasmussen, J (1986) *Information processing and human-machine interaction*. New York: North-Holland.
- [13] See, e.g., Miller, GA (1956). The magical number seven, plus or minus two. *Psychological Review*, 63, 81-97.
- [14] Sanford, AJ, & Garrod, SC (1981) *Understanding written language*. New York: Wiley.
- [15] NTSB (1988) *Aircraft accident report--Northwest Airlines, McDonnell Douglas DC-9-82, N312RC, Detroit Metropolitan Wayne Co. Airport, Romulus MI, August 16, 1987*. NTSB: Washington, D.C., 1988.
- [16] Bower, GH (1970). Organizational factors in memory. *Cognitive Psychology*, 1, 18-46.
- [17] Anderson, RC & Pitchard, JW (1978) Recall of previously unrecallable information following a shift in perspective. *Journal of Verbal Learning and Verbal Behavior*, 17, 1-12.
- [18] Bower, GH (1982) Plans and goals in understanding episodes. In A Flammer & W Kintsch (eds.), *Discourse processing*. New York: North-Holland.
- [19] Sanford & Garrod, *op. cit.*
- [20] Schank, RW et al. (1982) What's the point? *Cognitive Science*, 6, 255-275.
- [21] Miller, GA, Galanter, E, & Pribram, KH (1960). *Plans and the structure of behavior*. New York: Holt, Rinehart, & Winston.
- [22] Pew, RW. et al. (1986). *Cockpit automation technology* (BBN Rpt #6133). Cambridge, MA: BBN.