

How to Make Automated Systems Team Players

Klaus Christoffersen and David D. Woods
Institute for Ergonomics
The Ohio State University

To appear in
E. Salas (Ed.)
Advances in Human Performance and Cognitive Engineering Research
Volume 2.
JAI Press/Elsevier

Interface (noun): an arbitrary line of demarcation set up in order to apportion the blame for malfunctions.
(Kelly-Bootle, 1995, p. 101).

Human-Automation Cooperation: What have we learned?

Advances in technology and new levels of automation have had many effects in operational settings. There have been positive effects from both an economic and a safety point of view. Unfortunately, operational experience, field research, simulation studies, incidents, and occasionally accidents have shown that new and surprising problems have arisen as well. Breakdowns that involve the interaction of operators and computer-based automated systems are a notable and dreadful path to failure in these complex work environments.

Over the years, Human Factors investigators have studied many of the "natural experiments" in human-automation cooperation - observing the consequences in cases where an organization or industry shifted levels and kinds of automation. One notable example has been the many studies of the consequences of new levels and types of automation on the flight deck in commercial transport aircraft (from Wiener and Curry, 1980 to Billings, 1996). These studies have traced how episodes of technology change have produced many surprising effects on many aspects of the systems in question.

New settings are headed into the same terrain (e.g., free flight in air traffic management, unmanned aerial vehicles, aero-medical evacuation, naval operations, space mission control centers, medication use in hospitals). What

can we offer to jump start these cases of organizational and technological change from more than 30 years of investigations on human-automation cooperation (from supervisory control studies in the 1970's to intelligent software agents in the 1990's)?

Ironically, despite the numerous past studies and attempts to synthesize the research, a variety of myths, misperceptions, and debates continue. Furthermore, some stakeholders, aghast at the apparent implications of the research on human-automation problems, contest interpretations of the results and demand even more studies to replicate the sources of the problems.

Escaping from Attributions of Human Error versus Over-Automation

Generally, reactions to evidence of problems in human-automation cooperation have taken one of two directions (cf. Norman, 1990). There are those who argue that these failures are due to inherent human limitations and that with just a little more automation we can eliminate the "human error problem" (e.g., "clear misuse of automation ... contributed to crashes of trouble free aircraft", La Burthe, 1997). Others argue that our reach has exceeded our grasp -- that the problem is over-automation and that the proper response is to revert to lesser degrees of automated control (often this position is attributed to researchers by stakeholders who misunderstand the research results -- e.g., "... statements made by ... Human Factors specialists against automation 'per se' ", La Burthe, 1997). We seem to be locked into a mindset of thinking that technology and people are independent components --- either this electronic box failed or that human box failed.

This opposition is a profound misunderstanding of the factors that influence human performance (hence, the commentator's quip quoted in the epigraph). The primary lesson from careful analysis of incidents and disasters in a large number of industries is that many accidents represent a breakdown in **coordination** between people and technology (Woods and Sarter, 2000). People cannot be thought about separately from the technological devices that are supposed to assist them. Technological artifacts can enhance human expertise or degrade it, "make us smart" or "make us dumb" (Norman, 1993).

The bottom line of the research is that technology cannot be considered in isolation from the people who use and adapt it (e.g., Hutchins, 1995). Automation and people have to coordinate as a joint system, a single team (Hutchins, 1995; Billings, 1996). Breakdowns in this team's coordination is an important path towards disaster. The real lessons of this type of scenario and the potential for constructive progress comes from developing better ways to coordinate the human and machine team -- human-centered design (Winograd and Woods, 1997).

The overarching point from the research is that for any non-trivial level of automation to be successful, the key requirement is to design for fluent, coordinated interaction between the human and machine elements of the system. In other words, automation and intelligent systems must be designed to participate in team play (Malin et al., 1991; Malin, 1999).

The Substitution Myth

One of the reasons the introduction of automated technologies into complex work environments can fail or have surprising effects is an implicit belief on the part of designers that automation activities simply can be substituted for human activities without otherwise affecting the operation of the system. This belief is predicated on an assumption that the tasks performed within the system are basically independent. However, when we look closely at these environments, what we actually see is a network of interdependent and mutually adapted activities and artifacts (e.g., Hutchins, 1995). The cognitive demands of the work domain are not met simply by the sum of the efforts of individual agents working in isolation, but are met through the interaction and coordinated efforts of multiple people and machine agents.

Adding or expanding the role of automation changes the nature of the interactions in the system, often affecting the humans' role in profound ways (one summary is in Woods and Dekker, 2000). For example, the introduction of a partially autonomous machine agent to assist a human operator in a high workload environment is, in many respects, like adding a new team member. This entails new coordination demands for the operator - they must ensure that their own actions and those of the automated agent are synchronized and consistent. Designing to support this type of coordination is a post-condition of more capable, more autonomous automated systems. However meeting this post-condition receives relatively little attention in development projects. The result can be automation which leaves its human partners perplexed, asking Wiener's (1989) now familiar questions: what is it doing? why is it doing that? what is it going to do next?

As designers, we clearly want to take advantage of the power of computational technologies to automate certain kinds of cognitive work. However, we must realize that the introduction of automation into a complex work environment is equivalent to the creation of a new cognitive system of distributed human and machine agents and new artifacts. We must also realize the coordination across agents in the system is at least as important as the performance of the individual agents taken in isolation, especially when situations deviate from textbook cases. The attention we give to designing support for this coordination as incidents evolve and

escalate can be the determining factor in the success or failure of the human-machine system (Woods and Patterson, 2000).

How to Design for Coordination: Observability And Directability

More sophisticated automated systems or suites of automation represent an increase in autonomy and authority (Woods, 1996). Increasing the autonomy and authority of machine agents is not good or bad in itself. The research results indicate that increases in this capability create the demand for greater coordination. The kinds of interfaces and displays sufficient to support human performance for systems with lower levels of autonomy or authority are no longer sufficient to support effective coordination among people and more autonomous machine agents. When automated systems increase autonomy or authority without new tools for coordination, we find **automation surprises** contributing to incidents and accidents (for summaries see Woods, 1993; Woods, Sarter and Billings, 1997; Woods and Sarter, 2000).

The field research results are clear -- the issue is not the level of autonomy or authority, but rather the degree of coordination. However, the design implications of this result are less clear. What do research results tell us about how to achieve high levels of coordination between people and machine agents? What is necessary for automated systems to function as cooperative partners rather than as mysterious and obstinate black boxes? The answer, in part, can be stated simply as -- Cooperating automation is both observable and directable.

Observability: Opening up the Black Box

One of the foundations of any type of cooperative work is a shared representation of the problem situation (e.g., Grosz, 1981; McCarthy et al, 1991). In human-human cooperative work, a common finding is that people continually work to build and maintain a "common ground" of understanding in order to support coordination of their problem solving efforts (e.g., Patterson et al., 1999).

We can break the concept of a shared representation into two basic (although interdependent) parts: (1) a shared representation of the problem state, and (2) representations of the activities of other agents. The first part, shared representation of the problem situation, means that the agents need to maintain a common understanding of the nature of the problem to be solved. What type of problem is it? Is it a difficult problem or a routine problem? Is it high priority or low priority? What types of solution strategies are appropriate? How is the problem state evolving? The second part, shared representation of other agents' activities, involves access to information about what other agents are working on, which solution strategies they are pursuing, why they chose a particular strategy, the

status of their efforts (e.g., are they having difficulties? why? how long will they be occupied?), and their intentions about what to do next.

Together with a set of stable expectations about the general strategies and behavior of other agents across contexts, mutual knowledge about the current situation supports efficient and effective coordination among problem solving agents (Patterson et al., 1999). Agents can anticipate and track the problem solving efforts of others in light of the problem status and thus coordinate their own actions accordingly. The communicative effort required to correctly interpret others' actions can be greatly reduced (e.g., short updates can replace lengthy explanations). The ability to understand changes in the state of the monitored process is facilitated (e.g., discerning whether changes are due to a new problem or to the compensatory actions of others). An up to date awareness of the situation also prepares agents to assist one another if they require help.

Notice how much of the knowledge discussed here is available at relatively low cost in "open" work environments involving multiple human agents. For example, in older, hardwired control centers, individual controllers can often infer what other controllers are working on just by observing which displays or control panels they are attending to. In the operating room, surgical team members can observe the activities of other team members and have relatively direct, common access to information about the problem (patient) state. The open nature of these environments allows agents to make intelligent judgments about what actions are necessary and when they should be taken, often without any explicit communication. However, when we consider automated team members, this information no longer comes for free - we have to actively design representations to generate the shared understandings which are needed to support cooperative work.

Data Availability Does Not Equal Informativeness

Creating observable machine agents requires more than just making data about their activities available (e.g., O'Regan, 1992). As machine agents increase in complexity and autonomy, simple presentations of low-level data become insufficient to support effective interaction with human operators. For example, many early expert systems "explained" their behavior by providing lists of the individual rules which had fired while working through a problem. While the data necessary to interpret the system's behavior was, in a literal sense, available to operators, the amount of cognitive work required to extract a useful, integrated assessment from such a representation was often prohibitive. A more useful strategy was to provide access to the intermediate computations and partial conclusions that the machine agent generated as it worked on a problem. These were valuable because they summarized the machine agent's conception of the problem and the bases for its decisions at various points during the solution process.

In general, increases in the complexity and autonomy of machine agents requires a proportionate increase in the feedback they provide to their human partners about their activities. Representations to support this feedback process must emphasize an integrated, dynamic picture of the current situation, agent activities, and how these may evolve in the future. Otherwise, mis-assessments and miscommunications may persist between the human and machine agents until they become apparent through resulting abnormal behavior in the process being controlled. For example, the relatively crude mode indicators in the current generation of airliner cockpits have been implicated in at least one major air disaster. It is clearly unacceptable if the first feedback pilots receive about a miscommunication with automation is the activation of the ground proximity alarm (or worse).

Human agents need to be able to maintain an understanding of the problem from the machine agent's perspective. For instance, it can be very valuable to provide a representation of how hard the machine agent is having to work to solve a problem. Is a problem proving especially difficult? Why? If the automated agent has a fixed repertoire of solution tactics, which have been tried? Why did they fail? What other options are being considered? How close is the automation to the limits of its competence? Having this sort of information at hand can be extremely important to allow a human agent to intervene appropriately in an escalating critical situation.

Providing effective feedback to operators in complex, highly automated environments represents a significant challenge to which there are no ready-made solutions. Answering this challenge for the current and future generations of automation will require fundamentally new approaches to designing representations of automation activity (e.g., Sarter, 1999; Sklar and Sarter, 1999; Nikolic and Sarter, 2001). While the development of these approaches remains to be completed, we can at least sketch some of the characteristics of these representation strategies (Woods and Sarter, 2000). The new concepts will need to be:

- Event-based: representations will need to highlight changes and events in ways that the current generation of state-oriented display techniques do not.
- Future-oriented: in addition to historical information, new techniques will need to include explicit support for anticipatory reasoning, revealing information about what should/will happen next and when.
- Pattern-based: operators must be able to quickly scan displays and pick up possible abnormalities or unexpected conditions at a glance rather than having to read and mentally integrate many individual pieces of data

Directability: Who Owns the Problem?

Giving human agents the ability to observe the automation's reasoning processes is only one side of the coin in shaping machine agents into team players. Without also giving the users the ability to substantively influence the machine agent's activities, their position is not significantly improved. One of the key issues which quickly emerges in trying to design a cooperative human-machine system is the question of control. Who is really in charge of how problems are solved? As Billings (1996) pointed out, as long as some humans remain responsible for the outcomes, they must also be granted effective authority and therefore ultimate control over how problems are solved. Giving humans control over how problems are solved entails that we, as designers, view the automation as a resource which exists to assist human agents in the process of their problem solving efforts.

While automation and human activities may integrate smoothly during routine situations, unanticipated problems are a fact of life in complex work environments such as those where we typically find advanced automation. It is impossible in practice, if not in principle, to design automated systems which account for every situation they might encounter. While entirely novel problems may be quite rare, a more common and potentially more troublesome class of situations are those which present complicating factors on top of typical, "textbook" cases (cf. studies of brittleness of automated systems include Roth et al., 1987; Guerlain et al., 1996; Smith et al., 1997). These cases challenge the assumptions on which the pre-defined responses are based, calling for strategic and tactical choices which are, by definition, outside the scope of the automation's repertoire. The relevant question is, when these sorts of problems or **surprises** arise, can the joint system **adapt** successfully?

Traditionally, one response to this need has been to allow human operators to interrupt the automation and take over a problem manually. Conceiving of control in this way, an all-or-nothing fashion, means that the system is limited to operating in essentially one of two modes - fully manual or fully automatic. This forces people to buy control of the problem at the price of the considerable computational power and many potentially useful functions which the automation affords. What is required are intermediate, cooperative modes of interaction which allow human operators to focus the power of the automation on particular sub-problems, or to specify solution methods that account for unique aspects of the situation which the automated agent may be unaware of. In simple terms, automated agents need to be flexible and they need to be good at taking direction.

Part of the reason that directability is so important is that the penalties for its absence tend to accrue during those critical, rapidly deteriorating

situations where the consequences can be most severe. One of the patterns that we see in the dynamic behavior of complex human-machine systems during abnormal situations is an *escalation* in the cognitive and coordinative demands placed on human operators (Woods and Patterson, 2000). When a suspicious or anomalous state develops, monitoring and attentional demands increase; diagnostic activities may need to be initiated; actions to protect the integrity of the process may have to be undertaken and monitored for success; coordination demands increase as additional personnel/experts are called upon to assist with the problem; others may need to be informed about impacts to processes under their control; plans must be modified, contingencies considered; critical decisions need to be formulated and executed in synchronization with other activities. All of this can occur under time pressure (Klein et al., 2000).

These results do not imply that automation work only as a passive adjunct to the human agent. This is to fall right back into the false dichotomy of people versus automation. Clearly, it would be a waste of both humans' and automation's potential to put the human in the role of micro-managing the machine agent. At the same time however, we need to preserve the ability of human agents to act in a strategic role, managing the activities of automation in ways that support the overall effectiveness of the joint system. As was found for the case of observability, one of the main challenges is to determine what levels and modes of interaction will be meaningful and useful to practitioners. In some cases human agents may want to take very detailed control of some portion of a problem, specifying exactly what decisions are made and in what sequence, while in others they may want only to make very general, high level corrections to the course of the solution in progress. Accommodating all of these possibilities is difficult and requires careful iterative analysis of the interactions between system goals, situational factors, and the nature of the machine agent. However, this process is crucial if the joint system is to perform effectively in the broadest possible range of scenarios (Roth et al., 1997; Dekker and Woods, 1999; Guerlain et al., 1999; Smith et al., 2000; Smith, in press).

In contrast to this, technology-driven designs tend to isolate the activities of humans and automation in the attempt to create neatly encapsulated, pseudo-independent machine agents. This philosophy assumes that the locus of expertise in the joint human-machine system lies with the machine agent, and that the human's role is (or ought to be¹) largely peripheral. Such designs give de facto control over how problems are solved to the machine agent. However, experience has shown that when human agents are ultimately responsible for the performance of the system, they will actively devise means to influence it. For example, pilots in highly automated

¹ Recall that intelligent automation has often been introduced as an attempt to replace "inefficient" or "error-prone" human problem solvers.

commercial aircraft have been known to simply switch off some automated systems in critical situations because they have either lost track of what the automation is doing, or cannot reconcile the automation's activities with their own perception of the problem situation. Rather than trying to sort out the state of the automation, they revert to manual or direct control as a way to reclaim understanding of and control over the situation. The uncooperative nature of the automated systems forces the pilots to buy this awareness and control at the price of abandoning the potentially useful functions that the automation performs, thus leaving them to face the situation unaided.

Whither Automated Agents? Invest in Design for Team Play

Repeatedly, performance demands and resource pressures lead mission organizations to invest in increasing the autonomy and authority of automated systems. Because of unquestioned assumptions that people and automated systems are independent and inter-changeable, organizations fail to make parallel investments in design for observability and directability. Often in the process of recruiting resources for new levels of automation, advocates vigorously promote the claim that the more autonomous the machine, the less the required investment in team play and the greater the savings for the organization.

The operational effects of this pattern of thinking are strikingly consistent. Inevitably, situations arise requiring team play; inevitably, the automation is brittle at the boundaries of its capabilities; inevitably, coordination breakdowns occur when designs fail to support collaborative interplay; and inevitably, operational personnel must scramble to work around clumsy automation which is ill-adapted to the full range of problems or to working smoothly with other agents. Meanwhile, cycling in the background, commentators from various perspectives bicker about crediting one or another agent as the sole cause of system failures (Woods and Sarter, 2000).

We have no need to witness or document more of these natural experiments in strong, silent, difficult to direct automation. Experience has provided us with ample evidence for the shallowness, error, and sterility of these conventional beliefs. If we simply drop the blinders of the Substitution Myth, the scene comes into clear focus (Woods and Tinapple, 1999). The analysis of past natural experiments reveals ways to go forward. Because of increasing capabilities of automated systems, the design issue is collaboration within the joint human-machine system as this joint system copes with the variety and dynamics of situations that can occur. For this joint human-machine system to operate successfully, automated agents need to be conceived and designed as "team players". Two of the key elements needed to support this coordinated cognitive work are observability and directability.

Summary

When designing a joint system for a complex, dynamic, open environment, where the consequences of poor performance by the joint system are potentially grave, the need to shape the machine agents into team players is critical. Traditionally, the assumption has been that if a joint system fails to perform adequately, the cause can be traced to so-called "human error". However, if one digs a little deeper, they find that the only reason many of these joint systems perform adequately at all is because of the resourcefulness and adaptability that the human agents display in the face of uncommunicative and uncooperative machine agents. The ability of a joint system to perform effectively in the face of difficult problems depends intimately on the ability of the human and machine agents to coordinate and capitalize upon the unique abilities and information to which each agent has access.

For automated agents to become team players, there are two fundamental characteristics which need to be designed in from the beginning: observability and directability. In other words, users need to be able to see what the automated agents are doing and what they will do next relative to the state of the process, and users need to be able to re-direct machine activities fluently in instances where they recognize a need to intervene. These two basic capabilities are the keys to fostering a cooperative relationship between the human and machine agents in any joint system.

References

- Billings, C. E. (1996). *Aviation Automation: The Search for a Human-Centered Approach*. Hillsdale, NJ: Erlbaum.
- Dekker, S.W.A. and Woods, D.D. (1999). To Intervene or Not to Intervene: The Dilemma of Management by Exception. *Cognition, Technology and Work*, 1, 86-96, 1999.
- Grosz, B. J. (1981). Focusing and description in natural language dialogues. In A. K. Joshi, B. L. Webber, and I. A. Sag (Eds.), *Elements of Discourse Understanding*. Cambridge, MA: Cambridge University Press.
- Guerlain, S., Smith, P.J., Obradovich, J. Heintz, Rudmann, S., Strohm, P. Smith, J.W., Svirbely, J. and Sachs, L. (1999). Interactive critiquing as a form of decision support: An empirical evaluation. *Human Factors*, 41, 72-89.

- Guerlain, S., Smith, P.J., Obradovich, J.H., Rudmann, S., Strohm, P., Smith, J., & Svirbely, J. (1996). Dealing with brittleness in the design of expert systems for immunohematology. *Immunohematology*, 12 (3), p.101-107.
- Hutchins, E. (1995). *Cognition in the Wild*. Cambridge, MA: MIT Press.
- Kelly-Bootle, S. (1995). *The Computer Contradictionary* (2nd edition). MIT Press, Cambridge MA.
- Klein, G., Armstrong A., Woods, D., Gokulachandra M. and Klein, H. A. (2000). *Cognitive Wavelength: The Role of Common Ground in Distributed Replanning*. Prepared for AFRL/HECA, Wright Patterson AFB, September 2000.
- La Burthe, C. (1997). Human Factors perspective at Airbus Industrie. Presentation at International Conference on Aviation Safety and Security in the 21st Century. January 13-16, Washington DC.
- Malin, J. T., Schreckenghost, D. L., Woods, D. D., Potter, S. S., Johannesen, L., Holloway, M., and Forbus, K. D. (1991). *Making Intelligent Systems Team Players: Case Studies and Design Issues*. (NASA Technical Memorandum 104738). NASA Johnson Space Center: Houston, TX.
- Malin, J. T. (1999). Preparing for the Unexpected: Making Remote Autonomous Agents Capable of Interdependent Teamwork. In Proceedings of
- McCarthy, J. C., Miles, V. C., and Monk, A. F. (1991). An experimental study of common ground in text-based communication. In *Proceedings of the 1991 Conference on Human Factors in Computing Systems (CHI'91)*. New York, NY: ACM Press.
- Nikolic, M. I. and Sarter, N. B. (2001). Peripheral Visual Feedback: A Powerful Means of Supporting Attention Allocation and Human-Automation Coordination In Highly Dynamic Data-Rich Environments. *Human Factors*, in press.
- Norman, D. A. (1990). The 'problem' of automation: Inappropriate feedback and interaction, not 'over-automation.' *Philosophical Transactions of the Royal Society of London*, B 327:585--593.
- Norman, D. A. (1993). *Things that Make us Smart*. Addison-Wesley: Reading, MA.

- O'Regan, J. K. (1992). Solving the "real" mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology*, 46, 461-488.
- Patterson, E. S., Watts-Perotti, J. C. and Woods, D. D. (1999). Voice Loops as Coordination Aids in Space Shuttle Mission Control. *Computer Supported Cooperative Work*, 8, 353–371.
- Roth, E. M., Bennett, K. and Woods, D. D. (1987). Human interaction with an 'intelligent' machine. *International Journal of Man-Machine Studies*, 27:479--525.
- Roth, E. M., Malin, J. T., and Schreckenghost, D. L. (1997). Paradigms for Intelligent Interface Design. In M. Helander, T. Landauer & P. Prabhu (Eds.) *Handbook of Human-Computer Interaction* (2nd edition), Amsterdam: North-Holland. (pp. 1177-1201).
- Sklar, A. E. and Sarter, N. B. (1999). "Good Vibrations": The Use of Tactile Feedback in Support of Mode Awareness on Advanced Technology Aircraft. *Human Factors*, 41(4), 543-552.
- Sarter, N. B. (1999). The Need for Multi-sensory Feedback in Support of Effective Attention Allocation in Highly Dynamic Event-Driven Environments: The Case of Cockpit Automation. *International Journal of Aviation Psychology*, 10(3), 231-245.
- Smith, P.J., McCoy, E. and Layton, C. (1997). Brittleness in the design of cooperative problem-solving systems: The effects on user performance. *IEEE Transactions on Systems, Man and Cybernetics*, 27, 360-371.
- Smith, P. J., Billings, C., Chapman, R. J., Obradovich, J. H., McCoy, E., and Orasanu, J. (2000). Alternative architectures for distributed cooperative problem solving in the national airspace system. *Proceedings of the 5th International Conference on Human Interaction with Complex Systems*, Urbana, IL, 203-207.
- Smith, P.J., McCoy, E. and Orasanu, J. (in press). Distributed cooperative problem-solving in the air traffic management system. In Klein, G. and Salas, E. *Naturalistic Decision Making*. Mahwah, NJ: Erlbaum, 369-384.
- Wiener, E. L. (1989). *Human factors of advanced technology, ("Glass Cockpit") transport aircraft*. (NASA Contractor Report No. 177528). Moffett Field, CA: NASA Ames Research Center.
- Wiener, E. L. and Curry, R. E. (1980). Flight-deck automation: Promises and pitfalls. *Ergonomics*, 23, 995-1011.

- Winograd, T. and Woods, D. D. (1997). Challenges for Human-Centered Design. In *Human-Centered Systems: Information, Interactivity, and Intelligence*. J. Flanagan, T. Huang, P. Jones, S. Kasif, (Eds.), National Science Foundation, Washington DC, July, 1997.
- Woods, D. D. (1993). Price of flexibility in intelligent interfaces. *Knowledge-Based Systems*, (6), 4, 189-196.
- Woods, D. D. (1996). Decomposing Automation: Apparent Simplicity, Real Complexity, In R. Parasuraman and M. Mouloua, editors, *Automation Technology and Human Performance*, Erlbaum.
- Woods, D. D. and Dekker, S. W. A. (2000). Anticipating the Effects of Technological Change: A New Era of Dynamics for Human Factors. *Theoretical Issues in Ergonomic Science*.
- Woods, D. D. and Patterson, E. S. (2000). How Unexpected Events Produce an Escalation of Cognitive and Coordinative Demands. In *Stress Workload and Fatigue*. P. A. Hancock and P. Desmond (eds.) Lawrence Erlbaum, Hillsdale NJ.
- Woods, D. D. and Sarter, N. B. (2000). Learning from Automation Surprises and Going Sour Accidents. In N. Sarter and R. Amalberti (Eds.), *Cognitive Engineering in the Aviation Domain*, Erlbaum, Hillsdale NJ.
- Woods, D. D. and Tinapple, D. (1999). *W³: Watching Human Factors Watch People at Work*. Presidential Address, 43rd Annual Meeting of the Human Factors and Ergonomics Society, September 28, 1999. Multimedia Production at <http://csel.eng.ohio-state.edu/hf99/>
- Woods, D. D., Sarter, N. B. and Billings, C. E. (1997). Automation Surprises. In G. Salvendy (Ed.), *Handbook of Human Factors/Ergonomics*, (2nd ed.). New York, NY: Wiley.