*White Paper*

# Human-Centered Design of Automated Agents and Human-Automation Team Play

David Woods
Cognitive Systems Engineering Laboratory
Ohio State University
woods.2@osu.edu

July 11, 2001

## Summary

Autonomous capabilities and the ability to coordinate one's activities with others as part of a team or larger ensemble are two different skills sets. In very few places do we value humans with high autonomous capability but low coordinative capability, and, in such cases as this is unavoidable, we adapt other roles and the larger organization to compensate for the missing coordinative skills. Yet technologists continue to pursue increased autonomous capability assuming that coordinative skills can be grafted on subsequently with minor impact and at minor cost. Evidence from studies of expert human-human coordination and from breakdowns in existing human-automation systems indicate this assumption is untenable.

To create successful human-automation ensembles we need data, models and innovation on distributed cognition. Agents, be they human or machine, are always partially autonomous with differential capabilities. Robust performance derives from how different roles are coordinated given the inherent variability of the world, given irreducible uncertainty about the situation and future, given finite resources of any agent and set of agents tasked to achieve control in the pursuit of goals, and given that there are always multiple but potentially conflicting goals. This means that modeling, studying and designing how cognitive work is distributed is as fundamental as increasing machine autonomy and that balance across these lines of work is necessary if we are to avoid repeating 'classic' design errors.

## Background and Past Results

There is a need to coordinate the assessments and activities of autonomous computer agents with human monitors and supervisors. Past research on automation surprises have revealed that these computer agents are *semi*-autonomous and need to exhibit properties of effective team work that are seen in expert human-human cooperative systems (Sarter, Woods and Billings, 1997; Woods and Sarter, 2000). This finding has been captured in various summaries as the problem of: how to make automated agents team players (Malin et al., 1991), human centered automation (Billings, 1996), distributed cognitive systems (Hutchins 1995), and architectures for distributed supervisory control (Shattuck and Woods, 1996).

A variety of 'standard' approaches have been attempted to couple human supervisor and automated systems which have produced limited results or failed repeatedly. These include the human supervisor in the role of critiquing machine decisions. For example, Layton, Smith and McCoy (1996) demonstrated conclusively that this architecture created framing effects that undermined people's ability to assess the situation and cross-check the machine's proposed solution.

Management by exception architectures have been proposed as a way to minimize workload for human supervisory controllers--intervene only when the automation is having trouble solving the problem. These proposals have not been successful because they place the human in an untenable double bind (Dekker and Woods, 1999).

More fruitful directions have been identified based on studies of expert human-human collaboration in dynamic worlds. In one line of work, Woods and colleagues have studied mission control as an example of expertise at distributed collaborative systems. Another line of work has examined human-automation interaction on commercial aircraft flight decks (Sarter and Amalberti, 2000).

Among the critical issues for designers of cooperative architectures are:
- Overcoming brittleness in machine problem solving (Woods, Roth and Bennett, 1990),
- Handling the dynamics of how problems evolve and cascade (Woods, 1994; Woods and Patterson, 2000),
- Creating sources of resilience in the face of the potential for surprise, unanticipated variability, and complicating factors (Woods and Shattuck, 2000).
- Directing and re-directing the attention of collaborating agents as circumstances change (Woods, 1995; Patterson et al., 1999).

The studies of successful human-human distributed collaboration have suggested ways to solve these problems (Sarter, Woods, and Billings, 1997; Christoffersen and Woods, in press).

First, increasing the level of autonomy and authority of automated systems requires an *increase in the levels and kinds of feedback between agents* about their current and future activities as system state varies. Field studies, incidents and simulation results all reinforce this as a basic finding or 'law that governs cognitive system.' Ignoring this relationship produces coordination surprises between agents as events occur in the world.

The research challenge is to define the levels and forms of feedback needed to achieve coordination across partially autonomous human and computer agents. Critical to the new forms of feedback are representations of automation activity not simply process state. Previous work outlines some of the criteria solutions will meet: event-based, future-oriented, and pattern-based.

Second, increasing the level of autonomy and authority of automated systems requires mechanisms to manage or re-direct automated systems as resources -- directability. Giving human agents the ability to observe the automation's reasoning processes and activities is only one side of the coin in shaping machine agents into team players. Human supervisors also need the ability to substantively re-direct the machine agent's activities.

Giving humans control over how problems are solved entails that designers view the automation as a resource which exists to assist human agents in the process of their problem solving efforts. Past human-automation designs have generally tried to finesse this design problem by allowing manual takeover. The human operators can always decide to interrupt the automation and take over 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 and complexity demands and without many potentially useful functions which the automation provides. What is required are intermediate, coordinative 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. Re-directability criteria are not met if the human is put in the role of micro-managing machine agents. Instead, 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 to adapt responses to change and surprise in the world.

## Objectives: Seven Challenges for Research on Human-Automation Teamplay

Past work on effective distributed collaboration point to 7 critical challenges if automated systems are to be team players and if the distributed human-automation architecture will be capable of fluently adapting to complicating factors.

Challenge 1: Designing for observability of automation activities. New forms of feedback balance new levels of autonomy to achieve coordinated activity.

Challenge 2: Designing for re-directability of automated resources. Means to re-direct or constrain automated system activities as situations fluctuate in workload and criticality are necessary.

Challenge 3: Designing for resilience in adapting to surprises. The critical test of a joint or distributed cognitive system is resilience; the critical limit on a cognitive system is brittleness. Means to assess the resilience of a system are needed to spur design of human-computer collaborative architectures

Challenge 4: Designing for control of attention in a multi-threaded situations. Multiple lines of work on human-human collaboration have established that support for control of attention, support for judging interruptibility, support for re-directing the attention of others are all central and critical parts of coordinated multi-agent activity.

Challenge 5: Designing for building a common ground across multiple agents.

Challenge 6: Increasing levels of automation create on-call collaborative architectures which require support for the cognitive functions of orientation for intervention and coming up to speed.

Challenge 7: Designing for synchronization of activities.

**Challenge 1: Designing for observability of automation activities.**
To meet this goal techniques for pattern-based representations will be needed to make automation activities visible. New work on how to develop event pattern representations (e.g., Christoffersen, 1999; Chow, Christoffersen, Woods, Watts-Perotti, and Patterson, 2000) will be needed, that is, displays that support the recognition, monitoring, and anticipation of *events* (i.e., meaningful changes or patterns of changes in data over time). To shift communication to events rather than instantaneous values of base data elements will require intelligent displays that provide mechanisms for the generation, revision, and comparison of *expectations* to events and that use knowledge of plans and activities to highlight potential upcoming events.

Previous work outlines some of the criteria solutions will meet. The new forms of feedback must 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. |
| • Support Attention: | orienting perceptual functions such as audition, peripheral visual display, and tactile inputs need to be used to help operators re-orient attention as new events occur. |

**Challenge 2: Designing for re-directability of automated resources.**
To meet this goal research is needed on how agents in supervisory roles can change the delegation of tasks to lower order autonomous agents as a situation escalates.

One critical area for progress is development of techniques for communicating a machine's intent to human users and stakeholders. Different answers to the question -- What is a machine agent's intent – are different models of human-automation collaborative architectures. A machine agent's intent could be defined simply as the target(s) it is currently trying to achieve and the constraints it is trying to meet. This is a useful definition for design from the point of view of making automation's current activities observable. It is a straightforward, bedrock principle with few obstacles to use in the design of today's and future systems. Obstacles to using this principle in development lie in coordination failures between cognitive engineering and software engineering (Woods, Christoffersen, and Tinapple, 2000; Woods and Christoffersen, 2001).

Another useful answer is to think of a machine agent's intent as what it will do next depending on how the situation is developing or could develop. This is a contingency analysis point of view (Woods, 2001). Pursuing this line of modeling human-automation team play leads to potential new display innovations and potential new directions for applying machine intelligence.

A third useful answer from the point of view of design of human-automation team play is to see a machine agent as a stand in for distant human groups and organizations and their models of situations and how to act in local conditions. In this view (Woods, 2002), a machine agent's intent corresponds to distant groups' models of the situations actors on the scene will face and how actors on the scene should behave in those situations as constrained or supported by the design of the artifact. This dovetails with commander's intent when the distant commander's presence is mediated by computer artifacts. But as the military services increasingly use partially autonomous computer agents and computer based systems the same issue will arise relative to the design organization's intent for local decision and action. This level of analysis is particular useful in designing human-automation team play to handle the potential for surprises, both disrupting events and new opportunities (Woods and Shattuck, 2000).

Another critical area for progress is exploration of techniques for human supervisors to exert constraint based control of automated agent activities. This direction has been explored by NASA and the FAA in the context of new models of Air Traffic Management and new roles for ATC-aircraft coordination (Smith et al., in press).

### Challenge 3: Designing for resilience in adapting to surprises.

The very high potential for surprise in military operations leads organizations to develop a means to support skill at adapting to surprise within the context of larger plans. In command and control, supervisor-local actors teams practice communicating commander's intent and using intent information to develop skill at adapting to surprise (Klein. 1993; Shattuck and Woods, 2000). Similarly, the potential for surprise is high in space mission operations, and here too we see an organization that has developed a means to balance distant supervision with local adaptation.

Two types of failures are observed in these situations: either local actors failed to adapt plans and procedures to local conditions, often because they failed to understand that the plans might not fit actual circumstances (Type A problems), or they adapted plans and procedures without considering the larger goals and constraints in the situation. In the latter type B problems the failures to adapt often involved *missing side effects* of the changes in the replanning process (Woods and Shattuck, 2000).

Shattuck and Woods (2000) found this pattern in a study of how local actors adapted when surprises occurred in simulated command and control scenarios and how they used their commander's statement of intent behind the plan in adapting to unexpected events. At one extreme, practitioners would rotely follow the original plans as described by the supervisor with no regard for the local complicating factors. At the other extreme, practitioners would act completely autonomously, leaving their supervisors 'out of the loop' and failing to coordinate with other local actors toward an organizational target. The results demonstrate the need to strike a cooperative balance between remote supervisors and local actors, where local actors have the knowledge and authority that they need to respond to unanticipated local situations in ways that support achieving higher level goals.

These findings illustrate an inherent and fundamental tradeoff in the relationship between remote supervision and local action in establishing the framework for adaptation. Woods and Shattuck (2000) called skill at this tradeoff the *resilience function* of a distributed system.

Supervisors and the larger organizational context must determine the latitude or flexibility they will give actors to adapt plans and procedures to local situations given the potential for surprise in that field of activity. Supervision that establishes centralized control inhibit local actors' adaptations to variability, increasing the risk of Type A failures. At the other extreme is supervision that provides local actors complete autonomy. In the latter case, the goals and constraints important in remote supervisors' scope are disconnected from the activity and decision making of local actors. As a result, the response across multiple local actors may not be coordinated and synchronized properly, increasing the risk of Type B failures. Skill is a resilience process in distributed cognition that balances the risks across the two types of failure on either side of the tradeoff function.

The resilience function is both a constraint on models of distributed systems and a criterion for evaluation of the performance of a distributed system across a range of situations. Scenarios design should sample both sides of this tradeoff. Models of human-automation coordination need to consider how to maximize resilience without degrading response to textbook situations.

**Challenge 4: Designing for control of attention in multi-threaded situations.** Control of attention has proven to be a critical component of cognitive work and collaborative activity. Inherent in the situations of interest are multiple interwoven threads of reasoning and activity. Switching among these while balancing coherence with sensitivity to new events and information is a critical component of cognitive work. Architectures of distributed collaborative systems can support this function or degrade it. The design of interrupts and alerts are critical as well as the ability to look in on or listen in on ongoing activities without disrupting or competing with primary tasks. The use of alternative perceptual modalities has also proven extremely useful in supporting this collaborative function. Another fundamental part of human-human collaboration is the ability to see where another agent's attention is focused.

New work has begun to explore how people are able to do this well and inspire new ideas for how to support this function in human-automation teams (Woods, 1995; Patterson et al., 1999; Nikolic and Sarter, 2001; Sarter, 1999; Sklar and Sarter, 1999).

## Challenge 5: Designing for building a common ground across multiple agents.

Klein et al (2000) analyzed results from many studies in command and control and other areas to assess how distributed teams build and maintain a common ground.

Following a disrupting event, practitioners must not only adapt their plans to the evolving situation, but also endeavor to maintain a common ground with other agents affected by the event or the modification to ongoing plans. In a software engineering case, this might relate to a proposed funding cut that the project manager learns about during the course of the project. In this case, the manager must not only modify the project plan to accommodate the change, but also keep team members and other parties affected by the changes in-the-loop up dated. And team members need to anticipate how the changes affect processes within their scope of responsibility.

Common ground is essential for collaboration and difficult to establish and maintain for distributed teams. A number of mistakes and accidents can be traced back to the loss of common ground, whereby individuals believe they are talking about the same things or who believe they are making the same inferences, and do not realize how confusions have entered.

Establishing common ground between human and machine agents in part concerns how different agents understand the intent of other agents. It also concerns attributes of what is called an open workspace (Patterson et., 1999). Collaboration in an open workspace allows for unobstrusive assessment of the region of interest, activities, stance, goals of another agent. This has proven critical to effective human-human collaboration.

Analysis of requirements for how to build and maintain common ground provide important input to the design of representations that make the assessments and activities of automated agents visible to human monitors. Similarly, there is potential for the automated systems to adapt to the human supervisor's intent to handle disruptions and new opportunities much as commanders communicate intent to subordinates to help the latter adapt to surprises (Shattuck and Woods, 2000).

## Challenge 6: Design of on-call collaborative architectures require support for the cognitive functions of orientation for intervention and coming up to speed.

In envisioned mission operations environments, a small number of controllers and system technical experts will be responsible for large numbers of systems. Sophisticated monitoring and control automation will perform many of the duties that human agents perform today -- not only basic vigilance and procedural tasks, but also higher level system assessments and complex disturbance management. Not only will this reduce the number of human controllers and experts that need to be directly involved, but it will significantly transform the roles of the "on-duty" personnel that remain. These human agents will act as supervisors, less of the systems themselves than of the high level automation which monitors and interacts with these systems. These agents will have roles that are broader in scope, leading to a shift from personnel who are specialists, highly expert in a specific area, to generalists, broadly, but less deeply expert over a range of systems. Specialists who are deeply familiar with specific systems will still be retained in case of emergencies, but will not be a part of routine operations. This is the essence of the "on-call" operations model and is representative of a paradigm shift in domains such as nuclear power, industrial process control, military operations, and space operations, away from continuous monitoring towards a strategy of applying minimal human staffing until a problem arises. The shift to an "on-call" model of human expertise utilization is expected to maintain reliability while improving cost-effectiveness during nominal operations and most routine anomalous scenarios.

Under this new model, the function of the advanced automation is to independently cope with an ever larger range of nominal and off-nominal operational scenarios. This serves to change the definition of what sorts of operational states will be considered "routine" from the perspective of the supervising human agent. Significant automation activity may be occurring, without the involvement of controllers, as various minor disturbances and anomalies are identified and compensated for. Responsibility for tracking and reacting to these problems is transferred away from human agents to the automation. The function of human agents will increasingly be to ensure that the automation is functioning properly and appropriately, to coordinate across automation systems and with other human agents, and to act as a "safety net" in case a system enters a state that the automation cannot deal with. An important question therefore becomes: how much does the human agent need to know about what is going on?

In these envisioned operations environments, the operators and on-call system experts responsible for the monitored systems will need to remain "oriented" at some level to the past and current status of those systems. In order to effectively assess system states and processes, and be prepared to intervene if necessary, human monitors need to build and maintain a model of recent performance. This *orientation for intervention*, or mission awareness, provides a set of referents and expectancies against which the current system dynamics and automation activities can be evaluated. But at what level will these human agents need to maintain their awareness of the system state and the flow of automation activities in order to effectively perform in their new roles? What will they need to know to ensure that the automation is working properly and responding appropriately to system states? To perform coordination activities? How will human agents know when a system is nearing an unsafe state? If a serious anomaly is encountered, what will human agents need to know, and when, to be *prepared* to react effectively? How will human agents *come up to speed* as they become involved in an evolving situation where the incoming operator should understand what has happened as if he or she had been present and personally engaged in all the events and activities up to that point?

### Challenge 7: Designing for synchronization of activities.

Nyssen and Javaux (1996) results explore how multiple human agents synchronize activities. Pavard and others have explored synchronization constraints in communications between human agents, in particular the need to judge interruptibility. Patterson and Woods (2001) and Watts-Perotti and Woods (1997) found that cooperating human agents communicate information on the stance of groups towards events and activities as a distributed team modifies a plan in progress following a disrupting event. Chow, Christoffersen, and Woods (2000) integrate findings like these in a model of communication during collaborative work. This model provides a starting point for including synchronization constraints in the modeling and design of collaborative work.

## References

Billings, C.E. (1996). Aviation Automation: The Search For A Human-Centered Approach. Hillsdale, N.J.: Lawrence Erlbaum Associates.

Bond, A. et al. (editors). (in press). *Socially Intelligent Agents*. Klumwer Academic Press.

Chow, R., K. Christoffersen, D. D. Woods, J. Watts-Perotti, and E. Patterson. Communication during Distributed Anomaly Response and Replanning. . Institute for Ergonomics/Cognitive Systems Engineering Laboratory Report, ERGO-CSEL 00-TR-02, The Ohio State University, Columbus OH, September, 2000.

Chow, R., Christoffersen, K. and Woods, D.D. A Model of Communication in Support of Distributed Anomaly Response and Replanning. In Proceedings of the 44nd Annual Meeting of the Human Factors and Ergonomics Society/IEA2000, July, 2000.

Christoffersen, K. (1999). Event representations: Communicating change information in supervisory control. In Proceedings of the Human Factors and Ergonomics Society, 43rd Annual Meeting, Houston, TX.

Christoffersen, K. and D.D. Woods. (in press). How to make automated systems team players. In E. Salas (Ed.), *Advances in Human Performance and Cognitive Engineering Research, Volume* 2. JAI Press/Elsevier.

Dekker S.W.A. and D.D. Woods. To Intervene or Not to Intervene: The Dilemma of Management by Exception. *Cognition, Technology and Work*, 1, 86-96, 1999. a

S. Dekker and D.D. Woods. Extracting Data from the Future: Assessment and Certification of Envisioned Systems. In S. Dekker and E. Hollnagel (Eds.), Coping with Computers in the Cockpit., Ashgate, p. 7-27, 1999. b

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.*

Klein, G., A. Armstrong , D. Woods, M. Gokulachandra and H. A. Klein. Cognitive Wavelength: The Role of Common Ground in Distributed Replanning. Prepared for AFRL/HECA, Wright Patterson AFB, September 2000.

Layton, C., Smith, P.J. and McCoy, C.E. (1994). Design of a Cooperative Problem-Solving System for En-Route Flight Planning: An Empirical Evaluation. Human Factors, 36:94--119.

Malin, J., Schreckenghost, D., Woods, D., Potter, S., Johannesen, L., & Holloway, M. (1991). *Making Intelligent Systems Team Players* No. TR-104738). NASA Johnson Space Center.

Malin, J. T. (2000). Preparing For The Unexpected: Making Remote Autonomous Agents Capable Of Interdependent Teamwork. In Proceedings of the 44nd Annual Meeting of the Human Factors and Ergonomics Society/IEA2000, July, 2000.

Mumaw, R. J. & Roth, E. M. How to be more devious with a training simulator: Redefining scenarios to emphasize cognitively difficult situations. *1992 Simulation Multi-Conference: Nuclear Power Plant Simulation and Simulators*, Orlando, FL, April 6-9, 1992.

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.

Nyssen, A. S. and Javaux, D. (1996). Analysis of synchronization constraints and associated errors in collective work environments, Ergonomics, 39, 1249—1264.

Patterson, E.S., Woods, D.D., Sarter, N.B., & Watts-Perotti, J. (1998). Patterns in cooperative cognition. COOP '98, Third International Conference on the Design of Cooperative Systems. Cannes, France, 26-29 May.

Patterson, E. S., Watts-Perotti, J., & Woods, D. D. (1999). Voice loops as coordination aids in space shuttle mission control. *Computer-Supported Cooperative Work*, *8,* 353—371.

Patterson, E. S., & Woods, D. D. (2001). Shift Changes, Updates, and the On-Call Model in Space Shuttle Mission Control. *Computer-Supported Cooperative Work*, in press.

Ranson, D. and Woods, D.D. (1996). Animating Computer Agents. Proceedings of Human Interaction with Complex Systems, IEEE Computer Society Press, Los Alamitos, CA.

Roth, E. M., Malin, J. T. , & 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).

Roth, E. M., Mumaw, R. J., & Lewis, P. M. An Empirical Investigation of Operator Performance in Cognitively Demanding Simulated Emergencies. Washington D. C.: U. S. Nuclear Regulatory Commission, 1994. (NUREG/CR-6208)

Sarter, N.B. (1999). The Need for Multisensory 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.

Sarter N. and D.D. Woods. Team Play with a Powerful and Independent Agent: A Full Mission Simulation. *Human Factors*, 42, 390—402, 2000.

Sarter, N. D.D. Woods and C. Billings. Automation Surprises. In G. Salvendy, editor, *Handbook of Human Factors/Ergonomics*, second edition, Wiley, New York, 1997.

Sarter N. and D.D. Woods. Pilot Interaction with Cockpit Automation II: An Experimental Study of Pilot's Model and Awareness of the Flight Management System. *International Journal of Aviation Psychology*, 4:1-28, 1994.

Sarter N. and R. Amalberti, Eds., (2000). *Cognitive Engineering in the Aviation Domain*, Erlbaum, Mahwah NJ.

Shattuck, L. G. and D. D. Woods. Communication of Intent in Military Command and Control Systems. In Carol McCann and Ross Pigeau (editors), *The Human in Command: Exploring the Modern Military Experience*. New York : Kluwer Academic/Plenum Publishers, 279-292, 2000.

Schreckenghost, D. (1999). Checklists for human-robot collaboration during space operations. *Proceedings of the Human Factors and Ergonomics Society 43rd Annual Meeting*, 46-50.

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.

Smith, P, D. Woods, E. McCoy, C. Billings, N. Sarter, R. Denning and S. Dekker. Using forecasts of future incidents to evaluate future ATM system designs. *Air Traffic Control Quarterly*, 6(1), 71—85, 1998.

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.

Smith-Jentsch, K. A., Johnston, J. H. and Payne, S. (1998). Measuring team-related expertise in complex environments. In Cannon-Bowers, J. A. and Salas, E. (Eds.) *Making decisions under stress: Implications for individual and team training*. Washington, D. C.: American Psychological Association.

Watts-Perotti, J. and Woods, D. D. (1997). A Cognitive Analysis of Functionally-Distributed Anomaly Response in Space Shuttle Mission Control. CSEL Report 1997-TR-02). Cognitive Systems Engineering Laboratory, Institute for Ergonomics, The Ohio State University.

Woods, D.D. and K. Christoffersen. (2001). Balancing Practice-Centered Research and Design. In M. McNeese and M. A. Vidulich (editors), *Cognitive Systems Engineering in Military Aviation Domains*. Wright-Patterson AFB, OH: Human Systems Information Analysis Center, in press.

Woods, D.D., Christoffersen, K. and Tinapple, D. (2000). Complementarity and Synchronization as Strategies for Practice-Centered Research and Design. Plenary Address, 44th Annual Meeting of the Human Factors and Ergonomics Society and International Ergonomic Association, August 1, 2000. Multimedia Production to be available at http://csel.eng.ohio-state.edu/iea2000/

Woods, D.D. and E. S. Patterson. 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, 2000.

Woods D. D. and Sarter, N. (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 L. G. Shattuck. Distant supervision—local action given the potential for surprise *Cognition, Technology and Work*, 2, 86-96, 2000.

Woods, D.D. The price of flexibility in intelligent interfaces. *Knowledge-Based Systems*, 6:1-8, 1993.

Woods, D. D. (1994). Cognitive Demands and Activities in Dynamic Fault Management: Abductive Reasoning and Disturbance Management. In N. Stanton (Eds.), *Human Factors of Alarm Design* London: Taylor & Francis.

Woods, D.D. (1995). The alarm problem and directed attention in dynamic fault management. *Ergonomics*, 38(11), 2371-2393.

Woods, D.D. (2002). Laws that Govern Cognitive Systems. Plenary Address, Annual Meeting of the Cognitive Science Society, August, 2002.

Young, R. M. (2001). Using Grice's Maxim of Quantity to Select the Content of Plan Descriptions. *Artificial Intelligence*, in press.

Young, R. M. (in press). The Co-operative Contract in Interactive Entertainment. In Bond, A. et al. (editors). *Socially Intelligent Agents*. Klumwer Academic Press, in press.