Observations from Studying Cognitive Systems in Context

David D. Woods
Cognitive Systems Engineering Laboratory
Center for Cognitive Science
The Ohio State University

Keynote Address
Annual Conference of the Cognitive Science Society
August 1994

Cognitive Systems in the Wild

I study cognitive systems in the wild.[1] I study fields of practice where highly trained practitioners do cognitive work (monitor, assess, diagnose, plan and act) under time pressure, uncertainty and stress (Woods, 1994). In particular, I study the people and the technology in control centers. If we look at flightdecks of commercial jet airliners, or control centers that manage space missions, or surgical operating rooms, or control rooms that manage chemical or energy processes, or control centers that monitor telecommunication networks, or many other fields of human activity, what do we see?

First, we do not see cognitive activity isolated in a single individual, but rather cognitive activity goes on distributed across multiple agents (Hutchins, in press). Second, we do not see cognitive activity separated in a thoughtful individual, but rather as a part of a stream of activity (Klein et al., 1992). Third, we see these sets of active agents embedded in a larger group, professional, organizational, institutional context which constrains their activities, sets up rewards and punishments, defines not altogether consistent goals, and provides resources (Woods et al., in press). In the wilds, cooperation and coordination are ubiquitous.

Fourth, we see phases of activity with evolution and transitions. Cognitive and physical activity ebbs and flows, with periods of lower activity and more self paced tasks interspersed with busy, high tempo, externally paced operations where task performance is more critical. Higher tempo situations create greater need for cognitive work and at the same time often create greater constraints on cognitive activity (e.g., time pressure, uncertainty, exceptional circumstances, failures and their associated hazards). Fifth, we see that there are consequences at stake for the individuals and the groups and organizations involved in the field of activity or affected by that field of activity – economic, personal, safety.

Sixth, even a causal glance at these domains reveals that tools of all types are everywhere; almost all activity is aided by something or someone beyond the unit of the individual cognitive agent. Aided information processing is the norm.

Seventh, technology change is rampant in these settings. Ubiquitous computerization has tremendously advanced our ability to collect, transmit and transform data. In all areas of human endeavor, we are bombarded with computer processed data, especially when anomalies occur. User interface technology has allowed us to concentrate this expanding field of data into one physical platform (typically a single VDU) by providing the capability to manage multiple windows and the capability to generate tremendous networks of computer displays as a kind of virtual perceptual field viewable through the narrow aperture of the VDU (Woods, in press). Heuristic and algorithmic technologies expand the range of subtasks and cognitive activities that can be automated. These "intelligent" machines create joint cognitive systems that distribute cognitive work across multiple agents (Woods, 1986; Roth, Bennett and Woods, 1987; Hutchins, 1990).

But despite these possibilities and the claims of technologists, we find that many organizations have experienced significant difficulties in turning AI research and other new developments in computational technology into systems that actually improve performance in the target field of practice (e.g., space, flightdecks, air traffic control, nuclear power plant control rooms, communication network management, ground satellite control stations). In fact, we find that there seems to be an epidemic of failures labeled as "human error" as the complexity of systems grows (see Hollnagel, 1993; Woods et al., in press).

Our ability to understand artifacts, "their use and effects" (Winograd, 1987), has been limited. Some are lost in the details and short term horizon of particular fields of practice. Some are lost in the technology itself, blinded to larger views by the effort required to actually create new systems. Some are lost in their personal visions of what they imagine to be the impact of technology on human performance defined broadly. And some are simply aloof in the pursuit of apparently larger academic game.

Eighth, more in depth observation of the interaction of practitioners and artifacts reveals that the technology is often not well adapted to the needs of the practitioner – that much of the technology is clumsy in that it makes new demands on the practitioner, demands that tend to congregate at the higher tempo or higher criticality periods of activity (Woods, 1993). Our ability to digest and interpret data, despite the promises of the promoters of each wave of technology, has failed to keep pace with our abilities to generate and manipulate greater and greater amounts of data. Systems that automate some aspects of cognitive work are often strong, silent and non-directable. In other words, automation often does not function as a team player within the larger ensemble.

Ninth, close observation reveals that people and systems of people (operators, designers, regulators, etc.) are not passive in the face of the onslaught of clumsy technological artifacts. Rather they are active at adapting the tools and adapting their

[1] This lovely phrase is borrowed from Ed Hutchins; see Hutchins (in press).

activities continuously to respond to indications of trouble or to meet new demands. Furthermore, new machines are not used as the designers intended, but are shaped by practitioners to the contingencies of the field of activity in a *locally* pragmatic way (Woods et al., in press).

It has turned out that using new computational possibilities to create effective human-machine ensembles, what we will refer to as joint or distributed cognitive systems, is a substantive issue at the intersection of cognitive psychology, software engineering, social psychology and artificial intelligence (Hollnagel and Woods, 1983; Woods, 1986). In other words, this is an issue in Cognitive Science.

What is the status of studying cognitive systems in the wild? How do we deal with this conglomerate of technology and behavioral and social science, this conglomerate of individuals, sets of agents and organizational context? How does it relate to other activities and questions in Cognitive Science? Is this merely the application side of Cognitive Science? Are the methods employed in these studies casual or do they reflect the particular constraints of studying complex wholes?

The Joint or Distributed Cognitive Systems Perspective

The reverberations of technology change and observations of cognitive work in the wild lead us to an idea that can serve as a unifying theme. This is the idea suggested by Hollnagel and Woods (1983) and Hutchins (1991), among others, that one can look at operational systems – the individual people, the organization both formal and informal, the high technology artifacts (AI, automation, intelligent tutoring systems, computer-based visualization) and the low technology artifacts (displays, alarms, procedures, paper notes, training programs) intended to support human practitioners – that one can look at all of these things as a single cognitive system.

Operational systems can be thought of as joint and distributed human- machine cognitive systems in that:

- one can describe and study and design these systems in terms of cognitive concepts such as information flow, knowledge activation, control of attention, etc.,
- cognitive systems are distributed over multiple agents both multiple people and mixtures of people and apparently animate, agent-like machines,
- external artifacts come to function as cognitive tools through use and properties of these artifacts modify the activities of agents within the cognitive system,

There is a reciprocal relationship or mutual shaping between properties of external artifacts (e.g., how they represent aspects of the field of activity, e.g., Zhang and Norman, 1994; Woods, in press) and the cognitive activities distributed within the cognitive system. Properties of these artifacts and representations shape practitioner cognitive strategies and in turn these artifacts are shaped by practitioners to function as tools within the field of activity.

- cognitive systems adapt to the demands of the field of practice.

Hence, the cognitive systems perspective can be summarized by the triad – cognition in context, cooperation, and tools.

How To Make Automated Systems Team Players

One example of this intimate and intertangled triad can be seen when intelligent machine agents are introduced into a field of practice. We have observed this process across different domains, with different interventions, across different specific technological systems.

Heuristic and algorithmic technologies expand the range of subtasks and cognitive activities that can be automated. Automated resources can in principle offload practitioner tasks. Computerized systems can be developed that assess or diagnose the situation at hand, alerting practitioners to various concerns and advising practitioners on possible responses.

Our image of these machine capabilities is that of a machine alone rapt in thought or action. But the reality is that automated subtasks exist in a larger context of interconnected tasks and multiple actors. Introducing automated and intelligent agents into a larger system changes the composition of the distributed system of monitors and managers and shifts the human's role within that cooperative ensemble (Hutchins, in press). In effect, these 'intelligent' machines create joint cognitive systems that distribute cognitive work across multiple agents (Woods, 1986; Roth, Bennett and Woods, 1987; Hutchins, 1990; Billings, 1991). It seems paradoxical but studies of the impact of automation reveal that design of automated systems is really the design of a new human-machine cooperative system (contrast many of the discussions about machine abductive reasoning with the observations about human and cooperative abduction in Woods, 1994).

The behavior and capabilities of the machine agent in human-machine systems is changing. In simpler devices, each system activity was dependent upon operator input; consequently, the operator had to act to evoke undesired system behavior. In more automated systems, machine agents are capable of apparently "autonomous" activities. Once they are instructed and activated, systems are capable of carrying out long sequences of tasks without user interventions. These capabilities create new monitoring and coordination demands for humans in the system (Wiener, 1989; Norman, 1990; Sarter and Woods, in press). As observers of this change in technology have put it, the most common questions people ask about their automated partners are: what is it doing? why is it doing that? what will it do next? how in the world did we get into that mode?

These questions are asked in the context of automation surprises. Automation surprises are situations where the automated systems act in some way outside of the expectations of their human supervisors. Introducing more "autonomous" machines into a system increase the need for mechanisms to coordinate the activities of the multiple agents. Automation surprises are one class of symptoms of a deeper problem – the automated systems are not working as team players (Malin et al., 1991).

For example, consider the diagnostic situation in a multi-agent environment, when one notices an anomaly in the process they monitor (Woods, 1994). Is the anomaly an indication of an underlying fault, or does the anomaly indicate some activity by another agent in the system, unexpected by this monitor? In fact, in a number of different settings, we observe human practitioners respond to anomalies by first checking for what other agents have been or are doing to the process jointly managed. Data from studies of these surprises in aviation and medicine (Norman, 1990; Sarter and Woods, 1993; Moll van Charante et al., 1993) indicate that poor feedback about the activities of automated systems to their human partners is an important contributor to these problems.

Designing automated systems is more than getting that machine to function autonomously. It is also making provisions for that automated agent to coordinate its activity with other agents. Or, perhaps more realistically, it is making provisions so that other human agents can see the assessments and activity of the automated agent so that these human practitioners can perform the coordination function by managing a set of partially autonomous subordinate agents (see Billings, 1991; Sarter and Woods, 1994).

Cognitive Systems And Context

In studying cognitive systems in the wild we are concerned with cognitive work within complex fields of practice. We are context bound. Does this mean we are building an applied side to cognitive science? After all, the domain of Cognitive Science is the universal with respect to mind, brain and language? Is cognition in context simply the laboratory where one studies tool creation and skilled use as another in the pantheon of capabilities which are truly human? Is it an alternative paradigm, one based on anthropology where the study of cognition can only progress through study the study of the situation in which cognitive activity occurs? For me, progress is based on a creative tension between and complementarity among these possibilities rather than dominance of one view or another.

"*It is, ..., the fundamental principle of cognition that the universal can be perceived only in the particular, while the particular can be thought of only in reference to the universal*" (Cassirer, 1953, p. 86). As Hutchins puts it, "There are powerful regularities to be described at a level of analysis that transcends the details of the specific domain. It is not possible to discover these regularities without understanding the details of the domain, but the regularities are not about the domain specific details, they are about the nature of human cognition in human activity."[2] To be context bound in the study of cognitive systems is not simply to do "applied" studies in particular domains (though context bound studies when done well should influence short term change and immediate problems in the host "natural" laboratory).

It is in the tension between the particular and the universal in cognitive science that we can see the proper complementarity between so called basic and applied work where the experimenter functions as designer and the designer as experimenter. "New technology is a kind of experimental investigation into fields of ongoing activity. If we truly understand cognitive systems, then we must be able to develop designs that enhance the performance of operational systems; if we are to enhance the performance of operational systems, we need conceptual looking glasses that enable us to see past the unending variety of technology and particular domains" (Woods and Sarter, 1993).

The experimenter as designer? Cognitive tools are ubiquitous; technology change implicitly changes cognitive systems through the introduction of agent-like machines and through the introduction of artifacts that constrain cognitive work. This means new technology is a kind of experimental manipulation that can be exploited to help understand human cognition as expressed in meaningful fields of practice.

The designer as experimenter? The possibilities of technology seem to afford designers great degrees of freedom. The possibilities seem less constrained by questions of feasibility and more by concepts about how to use the possibilities skillfully to meet operational and other goals. In other words, in order for designs to be developed in a problem-driven manner, as opposed to the more typical technology-driven fashion, in order for designs to provide real and not illusory benefits for real operational systems, the designer must adopt the attitude of an experimenter trying to understand and model of the dynamics of joint cognitive systems.

To conceive of the experimenter as designer and designer as experimenter requires a drastic shift in the normal attitude of researchers and developers towards the subjects of their work – people and technology. Instead of separate and independent topics, they are intimately interconnected as parts of a larger and more useful system boundary – a joint cognitive system. To conceive of the experimenter as designer and designer as experimenter shifts the relationship between 'basic' and 'applied' research. The above concepts mean that these activities are complimentary where growing the research base and developing effective applications are mutually inter-dependent (Woods, 1993).

If we lose our balance and pursue the universal disconnected from the particular, we can miss the vary phenomena we claim to be studying and be aloof from the potential for new ideas to stimulate or modulate change. If we lose our balance and are governed only by the particular, we are lost in the short term horizon of today's hot buttons. By discarding barriers, barriers between technological and behavioral sciences, between individual and social perspectives, between the laboratory and the field, we will develop new insights about cognition and we will help steer technology change into more human-centered channels; in other words, we can "ascend to the particular."

Acknowledgment

The research on cognitive systems in the wild has been supported by grants from NASA Ames Research Center (Dr. Everett Palmer), the FAA (John Zalenchak), the Anesthesia Patient Safety Foundation, and NASA Johnson Space Center (Dr. Jane Malin). My colleagues at the Cognitive Systems Engineering Laboratory are integral to this work and the ideas it has generated – Richard Cook, Nadine Sarter, Scott Potter, and Leila Johannesen. Jane Malin has been particularly inspiring in discussions, critiques, and joint work and is responsible for the title which summarizes a great deal of this work – how to male intelligent systems teal players.

References

Billings, C.E. (1991). *Human-centered aircraft automation philosophy.* (NASA Technical Memorandum 103885). Moffett Field, CA: NASA-Ames Research Center.

Cook, R.I., Woods, D.D. and Howie, M.B. (1990). The natural history of introducing new information technology into a

[2]Hutchins, 1992, personal communication.

high risk environment. *Proceedings of the 35th Annual Meeting of the Human Factors Society.*

Flores, F., Graves, M., Hartfield, B. and Winograd, T. (1988). Computer systems and the design of organizational interaction. *ACM Transactions on Office Information Systems*, 6, 153-172.

Hollnagel. E. (1993). *Human Reliability Analysis: Context and Control.* London: Academic Press.

Hutchins, E. (1990). The technology of team navigation. In J. Galegher, R. Kraut, and C. Egido (Eds.), *Intellectual teamwork: Social and technical bases of cooperative work.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Hutchins, E. (in press). *Cognition in the Wild.* MIT Press.

Moll van Charante, E., Cook, R.I., Woods, D.D., Yue, Y., Howie, M.B. (1993). Human-computer interaction in context: Physician interaction with automated intravenous controllers in the heart room. In H. G. Stassen (Ed.), *Analysis, design and evaluation of man-machine systems 1992*, Pergamon Press.

Norman, D.A. (1988). *The psychology of everyday things.* New York, NY: Basic Books.

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.

Perrow, C. (1984). *Normal accidents. Living with high-risk technologies.* New York: Basic Books.

Reason, J. (1990). *Human error.* Cambridge, England: Cambridge University Press.

Roth, E.M., Bennett, K.B., and Woods, D.D. (1987). Human interaction with an "intelligent" machine. *International Journal of Man-Machine Studies*, 27, 479-525.

Sarter, N.B. and Woods, D.D. (1992). Pilot Interaction with cockpit automation: Operational Experiences with the Flight Management System. *International Journal of Aviation Psychology*, 2(4), 303-322.

Sarter N.B. and Woods D.D. (1994). Pilot interaction with cockpit automation II: An experimental study of pilots' mental model and awareness of the Flight Management System (FMS). *International Journal of Aviation Psychology*, 4, 1–28.

Sarter N.B. and Woods D.D. (1994). Decomposing Automation: Autonomy, Authority, Observability and Perceived Animacy. *Proceedings of Conference on Automation and Human Performance*, Washington DC, April, 1994.

Sarter, N.B. and Woods, D.D. (in press). "How in the world did we get into that mode?" Mode error and awareness in supervisory control. *Human Factors* (Special Issue on Situation Awareness).

Wiener, E.L. (1989). *Human factors of advanced technology ('glass cockpit') transport aircraft.* Technical Report 117528. Moffett Field, CA: NASA Ames Research Center.

Winograd, T. (1987). Three responses to situation theory. Technical Report CSLI-87-106, Center for the Study of Language and Information, Stanford University.

Winograd, T. and Flores, F. (1987). *Understanding computers and cognition.* Reading, MA: Addison-Wesley.

Woods, D.D. (1986). Paradigms for intelligent decision support. In E. Hollnagel G. Mancini, D. D. Woods (Eds.), *Intelligent decision support in process environments.* New York: Springer-Verlag.

Woods, D.D. (1988). Coping with complexity: The psychology of human behavior in complex systems. In L. P. Goodstein, H. B. Andersen and S. E. Olsen (Eds.), *Tasks, errors, and mental models.* New York: Taylor and Francis.

Woods, D.D. (1993). Price of flexibility in intelligent interfaces. *Knowledge-Based Systems*, 6(4), 189-196.

Woods, D.D. (1994). Cognitive demands and activities in dynamic fault management: Abductive reasoning and disturbance management. In Stanton, N. (Ed.), *The human factors of alarm design.* London: Taylor and Francis.

Woods, D.D. (in press). Towards a theoretical base for representation design in the computer medium: Ecological perception and aiding human cognition. In J. Flach, P. Hancock, J. Caird and K. Vicente (Eds.), *An ecological approach to human-machine systems I: A global perspective.* Hillsdale NJ: Lawrence Erlbaum Associates.

Woods, D.D., Johannesen, L.J., Cook, R.I., and Sarter, N.B. (in press). *Behind Human Error: Cognitive Systems, Computers, and Hindsight.* Crew Systems Ergonomic Information and Analysis Center (CSERIAC), Dayton, OH (State of the Art Report).

Zhang, J. and Norman, D.A. (1994). Representations in Distributed Cognitive Tasks *Cognitive Science*, 18(1), 87-122