

Distant Supervision–Local Action Given the Potential for Surprise

D. D. Woods¹ and L. G. Shattuck²

¹Cognitive Systems Engineering Laboratory, Institute for Ergonomics, Ohio State University, Columbus, OH, USA;

²Department of Behavioral Sciences and Leadership, United States Military Academy, West Point, NY, USA

1. PLANS, SURPRISE AND ADAPTATION

Process control (Woods et al 1987), military command and control (van Creveld 1985; Shattuck and Woods 1997), and air traffic management (Smith et al 1998) are examples of distributed supervisory control systems. Distributed supervisory control systems are hierarchical and cooperative. They include remote supervisors who work through intelligent local actors to control some process.

With this framework, human supervisors, designers and procedure writers can all be viewed as remote supervisors who provide plans and procedures to multiple local actors. The distant supervisors have a broader scope and a better understanding of the overarching goals and constraints for the larger distributed system. The local actors have privileged access to the monitored process and what is actually happening ‘on the ground’ within their field of view and narrower scope.

These plans and procedures often are inadequate to cope with the *potential for surprise* in specific situations (Woods and Roth 1988; Woods et al 1994). Given the potential for surprise and privileged access to data about the evolving situation, local actors must adapt the plans and procedures to the situation based on their understanding of the intent and goals behind the detailed steps in the plan (this is often referred to as the intent behind the plan).

The potential for surprise in a field of practice, as well as other factors, lead plans and procedures to be *underspecified* (Suchman 1990; Woods et al 1990). For example, in one study of device troubleshooting, a rule-based expert system’s directions to human technicians functioned as a kind of plan linking the expectations, analyses and heuristics of a remote designer to the local situation – an actual broken system to be restored to service. Inevitably,

in both anticipated and unanticipated ways, complicating factors arose which challenged execution of the troubleshooting plan. For example, impasses arose where a diagnostic test requested by the machine expert could not be carried out given other circumstances. The plans and procedures were underspecified in the face of the potential for surprise, requiring technicians to ‘supply knowledge and act outside of the scope and direction’ of the expert system (Roth et al 1987). As Suchman (1987) summarised, ‘instructions must be interpreted with respect to a collection of actions and circumstances that they never fully specify’; in other words, ‘plans are resources for action’.

Woods (1984) found in studies of simulated and actual nuclear power plant emergency operations that ‘good operations require more than rote rule following’. Two types of failures can occur when ‘events demanded a relatively variable sequence of component actions and extensive feedback from the environment in order to adapt to unpredictable constraints or disturbances’ (Woods et al 1987):

- ‘Type A problems where rote rule following persisted in the face of changing circumstances that demanded adaptable responses.’
- ‘Type B problems where adaptation to unanticipated conditions was attempted without the complete knowledge or guidance needed to manage resources successfully to meet recovery goals.’

In these studies, 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, 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 1988; Woods et al 1994).

Shattuck and Woods (1997) found the same 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 organisational 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.

The very high potential for surprise in military operations leads organisations 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 practise communicating commander's intent and using intent information to develop skill at adapting to surprise (Klein 1993; Shattuck and Woods 1997). Similarly, the potential for surprise is high in space mission operations, and here too we see an organisation that has developed a means to balance distant supervision with local adaptation.

Watts-Perotti (2000) observed that space shuttle mission control teams wrestle with the relationship between contingency plans and the unique characteristics of anomalous situations. During a space shuttle mission an anomaly occurred during the ascent phase. The anomaly disrupted plans for the mission. Among the reverberations of the disrupting event, a previously written contingency plan directed the controllers to drastically shorten the mission. Was this plan relevant to the situation they faced? What was the intent and assumptions behind that rule? How did this situation relate to that intent and assumed situation? How did this contingency plan apply to the trade-offs and risks in the situation they faced? How should the plans be adapted for this mission, given the anomaly, and how does this situation teach us to revise the plans for future missions?

Watts-Perotti analysed the cooperative interactions across multiple teams in the replanning process as they coped with the consequences of the disrupting event. The teams engaged in a sophisticated process of considering the implications of the anomaly for future plans, evaluating possible contingencies in light of the anomaly, and revising mission plans using the previous procedures as a resource. The study revealed a variety of mechanisms by which

mission control balanced distant plans as a guide to action with the need to adapt to surprising events.

2. RESILIENCE

These findings, from observations across multiple domains, illustrate an inherent and fundamental trade-off in the relationship between remote supervision and local action in establishing the framework for adaptation (Hollnagel 1993; Ashby 1956). We will call skill at this trade-off the *resilience function* of a distributed system.

Supervisors and the larger organisational 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 (Hirschhorn 1993). Supervision that establishes centralised 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 synchronised 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 trade-off function (Shattuck and Woods 1997; Woods and Patterson 2000).

3. PATTERNS IN THE JCO CRITICALITY ACCIDENT

With this conceptual backdrop from past research, we can see the operation of essential variables and generic patterns in the details and unique sequence leading to the JCO criticality accident and in stakeholder reactions in the aftermath of the accident as described by Furuta et al (2000). The accident is a case of the type B failure noted by Woods (1984). The team adapted plans and procedures to local constraints without considering all of the larger goals and constraints relevant to the situation. In particular, the adaptation process involved missing side effects of the changes.

The standard reaction to accidents like this is for stakeholders biased by hindsight to miss the underlying trade-off and resilience function completely. These stakeholders usually see the lesson of the accident as the need to impose new organisational pressures to 'just follow the rules'. The actual effects of this response are quite different from those imagined by the managers and regulators who impose this pressure (Woods et al 1994, ch. 4).

Increasing pressure to follow the rules is simply changing the weighting on the trade-off function, shifting the

criterion away from adapting plans to local conditions toward following plans even when local conditions challenge the plan. This organisational response just shifts the response of the distributed system from risking one type of failure to risking another type of failure – a criterion shift on a trade-off. The effectiveness of shifting the criterion in this direction depends on the potential for surprise and the consequences of failing to adapt to surprises.

It is critical to see that this organisational response does not enhance skill at handling the inherent trade-off – resilience. In fact it may even undermine skill because it creates a double bind for workers (Hirschhorn 1993): if they fail to adapt to handle surprises they will be blamed, but they also will be blamed if they adapt unsuccessfully. Understanding patterns in how workers adapt to such double binds is one part of exploring the trade-off at the heart of distributed supervisory control.

The lesson from this accident converges with many other sources of evidence about success and failure in distributed supervisory control. We need to explore and innovate new ways to support and enhance the resilience function in the relationship of distant supervision and local action given the potential for surprise and changing circumstances.

4. APPRECIATING POTENTIAL AND CHANGING ‘HAZARDS’ IN A SAFETY CULTURE

The organisations involved in the accident sequence failed to provide mechanisms to keep salient larger goals and constraints so that the local actors could recognise that they were relevant and how they were relevant in specific situations. As a result, the local actors considered adaptations to plans based on a set of local constraints without any appreciation of other constraints that were critically relevant in this specific case. The constraints they considered were related to factors and events that arose in operations in the experience of these workers, while the constraints missed were distant factors that had faded from view at all levels of the organisation and regulatory process.

A critical part of a culture of safety is how the organisation appreciates potential hazards and paths that can lead to those hazards in an accident sequence (Rochlin 1999; Woods and Cook 2000). A basic hazard had faded from view so that the accident struck as a fundamental surprise to the organisations involved, that is, as a path to failure that was outside of the organisations’ perceptions about how accidents could happen (Wagenaar and Groeneweg 1987; Woods et al 1994). In part, the organisations’ sensitivity to the hazard was reduced as a side effect of success in managing some of the paths that could lead to that failure (see Feynman’s comments on how

this same process contributed to the *Challenger* accident, 1986).

One can look at ‘perceived hazards’ as *potential* paths to or forms of failure. A path to failure is a story about how multiple contributors can come together to create the conditions for future failure. The stories of failures (actual or possible) we celebrate and how we tell each other these stories define that community’s view of the hazards as potential paths to failure (Westrum 1993; Rochlin 1999; Woods and Cook 2000). The community develops failure-sensitive strategies to forestall those potential paths and forms of failure of which they are aware.

Updating and calibrating our awareness of the *potential* paths is essential for avoiding failures because we are only partially aware of these paths, and, since the world is constantly changing, the paths are changing. The effort to escape or avoid stale, limited views of the changing potential for failure is one portion of the process of building a safety culture. The culture of safety observed in high-reliability organisations anticipates and plans for possible failures in ‘the continuing expectation of future surprise’. (Rochlin 1999, p. 1549).

The organisations involved in this accident took a background of apparent success as a basis for confidence on safety, especially given the backdrop of resource pressures and performance demands on production. They failed to sustain learning about the potential paths to failure as the basis for testing and updating their failure sensitive strategies. This is not unique to these organisations, but rather commonplace. Our challenge is to devise means for organisations to continually examine and reflect on the changing potential for failure and the effectiveness of their failure-sensitive strategies, before significant accidents like this occur.

References

- Asby WR (1956). An introduction to cybernetics. Wiley, New York.
- Fischhoff B, Johnson S (1990). The possibility of distributed decision making. In *Distributed decision making*. National Academy Press, Washington, DC.
- Furuta K, Sasou K, Kubota R, Ujita H, Shuto Y, Yagi E (2000). Human factor analysis of JCO criticality accident. *Cognition, Technology & Work* 2(4):182–203.
- Hirschhorn L (1993). Hierarchy vs. bureaucracy: the case of a nuclear reactor. In Roberts KH (ed). *New challenges to understanding organizations*. Macmillan, New York.
- Hollnagel E. (1993). *Human reliability analysis: context and control*. Academic Press, London.
- Klein GA (1993). Characteristics of commander’s intent statements. In *Proceedings of the 1993 BRG symposium on command and control research*.
- Rochlin G (1999). Safe operation as a social construct. *Ergonomics* 42:1549–1560.
- Rogers WP, et al (1986). Report of the Presidential Commission on the space shuttle Challenger accident. Government Printing Agency, Washington, DC.
- Roth EM, Bennett KB, Woods DD (1987). *Human interaction with an*

- 'intelligent' machine. *International Journal of Man–Machine Studies* 27:479–525.
- Shattuck L, Woods DD (1997). Communication of intent in distributed supervisory control systems. In Proceedings of the 41st annual meeting of the Human Factors and Ergonomics Society, September 1997.
- Smith P, Woods D, McCoy E, Billings C, Sarter N, Denning R, Dekker S (1998). Using forecasts of future incidents to evaluate future ATM system designs. *Air Traffic Control Quarterly* 6(1):71–85.
- Suchman L (1987). *Plans and situated actions: the problem of human–machine communication*. Cambridge University Press, Cambridge, UK.
- Suchman L (1990). What is human–machine interaction? In Robertson S, Zachary W, Black J (eds). *Cognition, computing and cooperation*. Ablex, Norwood, NJ.
- Van Creveld M (1985). *Command in war*. Harvard University Press, Cambridge, MA.
- Wagenaar W, Groeneweg J (1987). Accidents at sea: multiple causes and impossible consequences. *International Journal of Man–Machine Studies* 27:587–598.
- Watts-Perotti J (2000). How anomaly response is distributed across functionally distinct teams in space shuttle mission control. Manuscript submitted for publication.
- Westrum R (1993). Cultures with requisite imagination. In Wise JA, Hopkin VD, Stager P (eds). *Verification and validation of complex systems: human factors issues*. NATO ASI Series. Springer, Berlin.
- Woods DD (1984). Some results on operator performance in emergency events. In Whitfield D (ed). *Ergonomic problems in process operations*. Institute of Chemical Engineering Symposium Series 90.
- Woods DD (1988). Coping with complexity the psychology of human behavior in complex systems. In Goodstein LP, Andersen HB, Olsen SE (eds). *Mental models, tasks and errors*. Taylor & Francis, London, pp 128–148.
- Woods DD, Cook RI (2000). Characterizing the multiple landscapes of safety. Institute for Ergonomics, Ohio State University, working paper, May 2000.
- Woods DD, Patterson ES (2000). How unexpected events produce an escalation of cognitive and coordinative demands. In Hancock PA, Desmond P (eds). *Stress workload and fatigue*. Erlbaum, Hillsdale, NJ (in press).
- Woods DD, Roth EM (1988). Cognitive systems engineering. In Helander M (ed). *Handbook of human–computer interaction*. Elsevier Science, Amsterdam.
- Woods DD, O'Brien J, Hanes LF (1987). Human factors challenges in process control: the case of nuclear power plants. In Salvendy G (ed). *Handbook of human factors/ergonomics*. Wiley, New York.
- Woods DD, Roth EM, Bennett KB (1990). Explorations in joint human–machine cognitive systems. In Robertson S, Zachary W, Black J (eds). *Cognition, computing and cooperation*. Ablex, Norwood, NJ.
- Woods DD, Johannesen L, Cook RI, Sarter NB (1994). Behind human error: cognitive systems, computers, and hindsight. Crew Systems Ergonomic Information and Analysis Center, WPAFB, Dayton OH.
-
- Correspondence and offprint requests to:* D. D. Woods, Cognitive Systems Engineering Laboratory, Institute for Ergonomics, Ohio State University, Columbus, OH 43210, USA.