Chapter 8

Anomaly Response

Cascading Effects in an Episode of Anomaly Response

Several minutes into the ascent phase of a space shuttle mission (Figure 7), one of the flight controllers responsible for monitoring the health and safety of the mechanical systems in mission control centre noticed an anomaly – an unexpected drop in hydraulic fluid in an auxiliary power unit (APU). The personnel monitoring immediately recognized that the anomalous reading pointed to a hydraulic leak. Did this anomaly require an immediate abort of the ascent? In other words, she needed to assess—How bad was the leak? Was it a threat to the safety of the mission? What were the relevant criteria (and who knew them, and where did they reside)? The mechanical systems controllers did a quick calculation that indicated the leak rate was below the predetermined abort limit – the mission could proceed to orbit. The analysis of the event relative to an abort decision occurred very quickly, in part because the nature of the disturbance was clear and because of the serious potential consequences to the safety of the astronauts of any anomaly at this stage of a shuttle flight.

As the ascent continued, a second collection of demands and activities was intertwined and went on in parallel. The controllers for the affected system informed the Flight Director and the other members of the mission control team of the existence of the hydraulic leak and its severity. Because of the nature of the artifacts for supporting coordination across controllers and teams (voice loops, see Patterson et al., 1999), this occurred quickly and with little overhead costs in work and attention. The teams also had to consider how to respond to the anomaly before the transition from the ascent to the orbit phase was completed. Replanning was aimed both at how to obtain more information to diagnose the problem as well as how to protect the affected systems. It also required finding a way to resolve the conflicting goals of maximizing the safety of the systems as well as determining as confidently as possible the source of the anomaly. The team decided to alter the order in which the auxiliary power units (APUs) were shut down to obtain more diagnostic information. This change in the mission plan was then communicated to the astronauts.

After the initial assessment, responses, and communications, the new assessments of the situation and the changed plans were communicated to other controllers who were or might be affected by these changes. This happened in parallel because other personnel could listen in on the voice loops and overhear the previous updates provided to the flight director and the astronauts about the hydraulic leak. After the changes in immediate plans were communicated to the astronauts, the controllers responsible for other subsystems affected by the leak and the engineers who designed the auxiliary power units contacted the mechanical systems controllers to gain further information. In this process, new issues arose, some were settled, and sometimes issues that were previously handled needed to be re-visited or they re-emerged as new information arose.

For example, a series of meetings between the mechanical systems controllers and the engineering group were set up. This series of meetings served to assess contingencies and to decide how to modify mission plans such as a planned docking with the MIR space station, as well as re-entry. In addition, they provided opportunities to detect and correct errors in the assessment of the situation, to calibrate the assessments and expectations of differing groups, and to anticipate possible side effects of changing plans.

Additional personnel were called in and integrated with others to help with the new workload demands and to provide specialized knowledge and expertise. In this process, the team expanded to include a large number of agents in different places acting in a variety of roles and teams, all coordinating their efforts to produce a new mission and re-entry plan (from Woods & Patterson, 2000; see Watts et al., 1996 for more on the case).

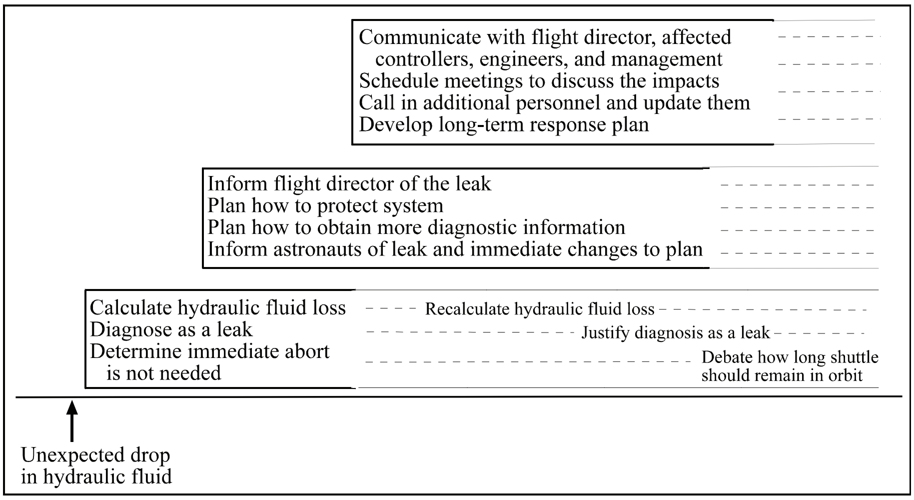

Figure 7. Schematic of an episode of anomaly response in space shuttle mission control (From Woods & Patterson, 2000).

CONTROL CENTERS IN ACTION

CSE began with studies of control centers in action (Figure 8). These studies focused on process control settings and, in particular, nuclear power plant control centers (Figure 4). The Three Mile Island accident in 1979 presented a "fundamental surprise" (Woods et al., 1994) that triggered international efforts to re-examine how teams of operators handle these kinds of plant emergencies (Rasmussen & Rouse, 1981) and to re-design analog control centers using digital technologies in the hope of providing better support for operator work (Woods et al., 1987).

These studies examined one major class of work performed by (JCSs)—*anomaly response*. In anomaly response, there is some underlying process, an engineered or physiological process which will be referred to as the *monitored process*, whose state changes over time. Faults disturb the functions that go on in the monitored process and generate the demand for practitioners to act to compensate for these disturbances in order to maintain process integrity—what is sometimes referred to as "safing" activities. In parallel, practitioners carry out diagnostic activities to determine the source of the disturbances in order to correct the underlying problem.

Anomaly response situations frequently involve time pressure, multiple interacting goals, high consequences of failure, and multiple interleaved tasks (Woods, 1988; 1994). Typical examples of fields of practice where dynamic fault management occurs include flight deck operations in commercial aviation (Abbott, 1990), control of space systems (Patterson et al., 1999; Mark, 2002), anesthetic management under surgery (Gaba et al., 1987), terrestrial process control (Roth, Woods & Pople, 1992), and response to natural disasters.

Figure 8. Spaceflight controllers in the front room at Mission Control (Courtesy of NASA).

When a researcher first encounters a control center and considers abnormal situations, they often may say the topic under study is "diagnosis" and imply that diagnosis involves information processing that links sets of symptoms to categories of faults (Woods, 1984; Rasmussen, 1986). This view mistakes one very specialized type of diagnostic setting—troubleshooting a broken device which has been removed from service—as if it instantiated a comprehensive model of anomaly response. It mistakes one strategy for diagnostic search (symptom-fault matching) as if it were a comprehensive view of the challenges of anomaly response as a general form of work in JCSs. As a result, studies and designs based on these assumptions have proved unable to penetrate the mask of an information processing sequence and have failed to recognize the deeper patterns of adaptation and dynamic control in anomaly response (*JCS-Foundations*, Chapter 1).

Work based on CSE approached the questions about how to improve operator response in abnormal and emergency situations by following the process schematicized in Figure 1:

- Making diverse observations of the interplay of people and technology in demanding work across different settings,
- Abstracting patterns across specific settings and situations to escape from surface variability—even to recognize that anomaly response is a general type of work for JCSs,
- Generating models of what makes anomaly response difficult (vulnerabilities) and what is the basis for expert and successful performance,
- Innovating promising hypotheses about what would be useful based on understanding how artifacts provide affordances relative to these demands.

What did these processes of functional synthesis reveal about anomaly response?

Cascading Effects

The driving constraint in anomaly response is that faults produce a time series of disturbances along lines of functional and physical coupling in the monitored process, or a **cascade of disturbances** (e.g., Abbott, 1990). The cascade of effects takes on different time courses in different circumstances. For example, an event may manifest itself immediately or may develop more slowly; a triggering event may occur pristine against a quiet background or in combination with other events in a changing context.

The cascade of disturbances captures the process by which situations move from canonical or textbook to non-routine to exceptional. Different domains may have different gradients in this process depending on the kinds of complicating factors that occur, the rhythms of the process, and consequences that may follow from poor performance.

Operational settings are typically data rich and noisy. Many data elements are present that could be relevant to the problem solver (Woods, 1995b). There are a large number of data channels (either through sensors, human reports, or direct perception) and the signals on these channels usually are changing and some changes are not significant. The raw values are rarely constant even when the system is stable and normal. The default case is detecting emerging signs of trouble against a dynamic background of signals rather than detecting a change from a quiescent, stable, or static background. The noisiness of the background can mask symptoms and can provide plausible alternatives to explaining changes and events away as part of a different pattern (e.g., Roth et al., 1992).

Note how basic characteristics of anomaly response can be used to guide problem-design for staged world studies. For example, a basic heuristic is to develop a scenario where the events and disturbances of interest occur against a moving or turbulent background where practitioners are already active in handling or adjusting to another event.

Interventions

The nature of the responses by practitioners affects how the incident progresses—less appropriate or timely actions (or too quick a reaction in some cases) may sharpen difficulties, push the tempo in the future, or create new challenges. The operational responses by automatic systems and by operators are part of the process of disturbance propagation and development over time and, therefore, a part of the episode itself.

Given the consequences of disturbances for safety and other pressing goals, operators cannot wait for a definitive diagnosis before acting to counter or correct the situation. "Safing" actions begin relatively early in the sequence of events. For example, while monitoring telemetry data during one space shuttle mission, flight controllers recognized high and increasing pressure in a fuel line connected to an auxiliary power unit (APU). At first, the controllers thought about several possible sources that could be producing the anomaly (a blockage in the line which in turn could be due to different sources—contamination blocking the line or a small leak could expose the hydrazine fuel to the freezing conditions of space and the resulting fuel freeze could block the line—with different responses). However, they saw that the pressure continued to increase and approach the limit for the fuel line. To keep the line from bursting, flight controllers asked the shuttle crew to cycle a relief valve to lower the pressure. This response action served as a **therapeutic intervention** (e.g., safing), which carried the goal of relieving the pressure regardless of what was producing the anomaly, so that the line would not burst under the growing pressure.

As illustrated in this example, actions to manage the consequences of disturbances are highly intertwined with diagnostic search. Critically, interventions taken for therapeutic purposes also function as **diagnostic interventions**; that is, the response of the monitored process to the intervention provides information about the nature and source of the anomaly (therapeutic interventions sometimes serve quite simply to "buy time" for diagnostic search). To continue the space shuttle example, practitioners followed up the action of cycling a relief valve by monitoring for the effect of this therapeutic intervention. The intervention was expected to produce a decrease in pressure. Instead, the flight controllers were surprised to observe a quite different response to opening of the relief valves. The telemetry value of the fuel line pressure continued to gradually rise. Since opening the relief valves had no effect on the pressure value, the practitioners began to hypothesize that the telemetry signature must have been due to a failed transducer or pressure sensor. In this case, the intervention that was intended to be therapeutic also became a diagnostic intervention because it provided practitioners with more evidence about possible explanations (note that actions taken only for diagnostic value, as interventions, may still have risks relative to the monitored process; e.g., checking a difficult to access area of a process may introduce new problems, or in cardiovascular anesthesiology, different actions to gain information pose varying risks of other injuries to the patient).

Revision

The time course of disturbances and interventions produces a series of changes in the data available about the state of the monitored process. Evidence about the anomaly and what produced it comes in over time. The high fuel line pressure case illustrates the process: eventually the flight controllers observed that the fuel line pressure leveled off and stayed constant (later, pressure began to drift slowly back toward normal). As they pondered what would account for this additional and unexpected behavior, the practitioners also began to consider how to modify plans for future mission events and evaluate contingencies that might arise in light of the anomaly. In working through this thread in anomaly response, the flight controllers spoke with others at the shuttle launch site and learned that the shuttle had been exposed to a rainstorm while it was waiting to launch. Therefore, the insulation surrounding the anomalous fuel line might have gotten wet. The new information, combined with the new events in pressure behavior (stabilize, then drift lower) led some practitioners to hypothesize that the high pressure was caused by wet insulation. They proposed that when the shuttle reached outer space, the insulation around the line froze and, therefore, caused the hydrazine fuel to freeze. As the ice sublimated from the insulation and the heaters warmed the line, the frozen hydrazine began to melt and return to its normal state. This new hypothesis accounted for all of the findings to be explained and therefore many practitioners involved believed that this provided the best explanation for the anomaly.

This example highlights the critical demand factor and basic vulnerability in anomaly response—**revision or failures to revise** of assessments. The evidence across studies of anomaly response show that initial explanatory hypotheses tend to be correct or plausible given the evidence available at that point. Problems arise later as the situation changes and as new evidence comes in, if the practitioners fail to revise their assessments in light of these changes. In other words, the initial assessments are basically correct (given what information is available) but practitioners can get stuck in their previous assessment and fail to revise these assessments as the situation changes. To see specific results, check the synthesis from multiple studies of operator performance in nuclear power emergencies, both observation during high fidelity simulations and retrospective analyses of actual accidents in Woods et al., 1987 and also Roth et al., 1992; or see Paul Johnson's line of research across Johnson et al., 1988; 1991; 1992; 2001).

Garden path problems are a specific class of problems where revision is inherently difficult since "early cues strongly suggest [plausible but] incorrect answers, and later, usually weaker cues suggest answers that are correct" (Johnson, Moen & Thompson, 1988). The following is an example of a garden path situation from nuclear power plant failures (from Roth, Woods & Pople, 1992).

An ISLOCA is an interfacing system loss of cooling accident in one kind of nuclear power plant. It can happen if a double valve failure occurs in the pipes that connect the high pressure cooling system to the low pressure cooling system. The most salient early symptoms of an ISLOCA can lead one down a garden path because they suggest another fault: a primary system break inside containment. The operational context encourages the garden path aspect in that once the symptoms consistent with a primary system break are observed, the goal becomes to maintain primary system integrity, placing a significant demand on operator attention and activities (potentially diverting resources from situation assessment and additional hypothesis generation and evaluation).

Since the evidence strongly indicates a primary break inside containment, during emergencies, crews are supposed to follow emergency procedures that are specific to the particular diagnostic conclusion about the type of anomaly or fault present. It is interesting to note that once the operators begin to follow the primary "break inside containment" procedure, there is no provision within the structure of the procedure system to redirect them to the ISLOCA procedure (this is a hole in how the procedure set was designed to accommodate the general risk of misdiagnosis of fault type). It is only over time that the complete set of symptoms that indicate the actual nature of the fault build up, making the primary system break to containment diagnosis increasingly less plausible.

Garden path problems are one category of problems that illustrates how anomaly response is fraught with inherent uncertainties and trade-offs. Incidents rarely spring full blown and complete; incidents evolve. Practitioners make provisional assessments and form expectancies based on partial and uncertain data. These assessments are incrementally updated and revised as more evidence comes in. Furthermore, situation assessment and plan revision are not distinct sequential stages, but rather they are closely interwoven processes with partial and provisional plan development and feedback that lead to revised situation assessments (Woods, 1994; see Figure 11, p. 86).

As a result, it may be necessary for practitioners to entertain and evaluate those assessments that later turn out to be erroneous. In the garden path problem above, it would be considered quite poor performance if operators had not seriously considered the primary break to containment hypothesis or if they had been too quick to abandon it since it was highly plausible as well as a critical anomaly if present. The vulnerability lies not in the initial assessment, but in whether or not the revision process breaks down and practitioners become stuck in one mindset or even become fixated on an erroneous assessment, thus missing, discounting, or re-interpreting discrepant evidence. Unfortunately, this pattern of **failures to revise** situation assessment as new evidence comes in has been a part of several major accidents (e.g., the Three Mile Island accident; Kemeny, 1979).

Fixation

Cases of fixation begin, as is typical in anomaly response, with a plausible or appropriate initial situation assessment, in the sense of being consistent with the partial information available at that early stage of the incident. As the incident evolves, however, people fail to revise their assessments in response to new evidence, evidence that indicates an evolution away from the expected path. In fixations, the practitioners appear to be stuck on the previous assessment. They fail to revise their situation assessment and plans in a manner appropriate to the data now present (for data on fixations see De Keyser & Woods; 1990; Johnson et al., 1981; 1988; 2001; Gaba et al., 1987; Rudolph, 2003). A failure to revise can be considered a **fixation** when practitioners fail to revise their situation assessment or course of action and **persist** in an inappropriate judgment or course of action **in the face of opportunities to revise**.

One critical requirement in order to describe an episode as a fixation is that there is some form of **persistence** over time in the behavior of the fixated person or team. This means cases of fixation include opportunities to revise, that is, cues, available or potentially available to the practitioners, that could have started the revision process if observed and interpreted properly. In part, this feature distinguishes fixations from simple cases of inexperience, lack of knowledge, or other problems that can impair detection and recovery from mis-assessments. The basic defining characteristic of fixations is that the immediate problem-solving context has framed the practitioners' mindset in some direction that is inappropriate given the actual evolution of the episode. In naturally occurring problems, the context in which the incident occurs and the way the incident evolves activates certain kinds of knowledge as relevant to the evolving incident (as described in the perceptual cycle; *JCS-Foundations*, p. 20). This knowledge, in turn, affects how new incoming information is interpreted. In psychology of cognition literature, this is often referred to as a *framing effect.* Another distinguishing mark in cases of fixation is that, after the fact, or after the correct diagnosis has been pointed out, the solution seems obvious, even to the fixated person or team.

Analyses of failures to revise that have occurred in field studies suggest several different forms of persistence (see De Keyser & Woods, 1990). In some cases the operators seem to have many hypotheses in mind, but never the correct one. The external behavior exhibited appears incoherent as they jump from one action to another one, but without any progress. Ongoing persistence on a single line of thought and action is seen in other cases, usually taking the form of multiple repetitions of the same action or the same monitoring check despite an absence of results or feedback. The operators seem to detect the absence of any effect or change, yet continue in the same vein. In a third pattern the operators seem insensitive to new cues and evidence. Cues discrepant with their model of the situation seem to be missed, discounted, or reasoned away (sensor failures and false alarms may be quite plausible reasons given past experiences).

Feltovich et al. (2001) studied pediatric cardiologists in a garden path scenario and abstracted a set of knowledge shields that people use to explain away discrepancies that could signal that their assessment or framing of the situation is mistaken. Expertise at revision appears to be quite special and seems to depend on taking a “suspicious stance” toward any data that fails to fit the current assessment, despite the relative prevalence of “red herrings” (data changes sometimes are not significant or due to other minor factors).

There is a complementary vulnerability to be avoided too, which has been called *vagabonding* (Dörner, 1983). Not every change is important; not every signal is meaningful. If the team pursues every signal and theme, they can jump incoherently from one item to the next, treating each superficially and never developing a coherent coordinated response. Shifting attention constantly from one change, item, or thread to another undermines the ability to formulate a complete and coherent picture of the trajectory of the system. Conversely, not shifting attention to new signals can result in missing cues that are critical to revising situation assessment. The contrast captures the potential trade-off between the need to revise assessments and the need to maintain coherence.

When adopting a suspicious stance, practitioners recognize the fragility and tentative nature of assessments and actively invest in continuing to check for and exercise sensitivity to new cues that could update or modify their assessment, *despite* the need (or pressures) to carry out other necessary activities consistent with their roles (Klein, Pliske, et al., 2005). For dramatic contrasting examples of the difficulty, examine the history of how mission control handled anomalies during the Apollo program as described in Murray & Cox, 1989 (especially examine Apollo 13 with regards to what mission control was doing to make sense of the many anomalous readings during the first 40 minutes following the initial presentation of the anomaly. Consider how the team avoided becoming fixated on the interpretation that the multiple anomalies were due to an instrumentation problem); then contrast this with how the Mission Management Team quickly discounted uncertain evidence of trouble rather than following up vigorously (see CAIB, 2003, Chapter 6; Woods, 2005b).

Generating Hypotheses

Performance at diagnostic reasoning depends on the ability to generate multiple plausible alternative hypotheses that might account for the findings to be explained. Observing anomaly response in action reveals that people do not call to mind all plausible alternatives at once. Rather, context, new evidence, and the response of the monitored process to interventions all serve as cues that suggest new possible explanations. Additionally, they change the relative plausibility of different candidates in ways that shift which ones guide the diagnostic search for additional evidence. In the space shuttle high fuel line pressure anomaly case, note how the set of candidate explanations shifted over time.

This part of anomaly response is referred to as hypothesis generation. Since diagnostic reasoning involves evaluation of hypotheses relative to available competitors, the generation of possible alternative hypotheses is an area that can influence the quality of reasoning. One strategy for supporting expert performance is to aid hypothesis generation, i.e., to **broaden** the set of hypotheses under consideration as candidate explanations for the pattern of findings. Research strongly suggests that (1) multiple people—in the right collaborative interplay—can generate a broader set of candidates than individuals, regardless of level of individual expertise (Gettys et al., 1979; Gettys et al., 1986; Gettys et al., 1987; Hong & Page, 2002 and (2) using automation to remind or critique during human practitioners' anomaly response process—with the right interaction design—produces a broader exploration of possibilities (and perhaps deeper evaluation as well) compared to human-machine architectures where people are asked to monitor the automation's assessments and activities (Woods, 1986a; Layton et al., 1994; Guerlain et al., 1996; Smith et al., 1997; Guerlain et al., 1999). The critical performance attribute is broader or narrower exploration in pace with the evolving situation. Supporting this is inherently a collaborative, multi-agent issue.

As we will see in exploring breakdowns in coordination (automation surprises in Chapter 10), the conventional and common belief that people should be placed in the role of monitoring the automation's assessments and activities produces surprising negative effects. Having the machine first suggest an explanation or provide an assessment can narrow the range of alternative accounts considered by the human monitor (and the total set explored by the collaborative ensemble). Studies have shown that when their evaluation is framed by the machine's suggestion, people miss weaknesses or limits in the machine's account and to accept accounts/plans that they would not find appropriate when doing the task themselves (Layton et al., 1994; Smith et al., 1997). This finding strongly illustrates the advantage of taking a JCS perspective over evaluating and designing each component in isolation. Supporting broadening and avoiding premature narrowing depends on designing the coordination between agents. These joint system issues cannot be addressed if one develops a machine advisor or automation first and then attempts to graft on a subsequent design of a human interface.

The issues of hypothesis generation and revision of assessments are all complicated by the fact that multiple faults or factors may be present and producing the observed time varying course of disturbances. In complex, highly coupled monitored processes, multi-fault or multi-factor cases may be relatively likely or disproportionately risky. For example, faulty sensor readings may be relatively likely to occur in combination with another fault when processes are heavily instrumented or prone to interference. Diagnosis difficulty increases when there is greater coupling or interactions within the monitored process, as this produces more complicated disturbance chains, which exacerbate the difficulties inherent in distinguishing when multiple factors are present (Woods, 1994). For example, disturbances arising from separate underlying conditions may interact and appear to be due to single failure. The question—Is a series of disturbances the effect of a single fault, or are they due to multiple factors?—is omnipresent in anomaly response. Other related diagnostic difficulties in anomaly response include *effects at a distance* where the symptoms that are visible or salient arise from disturbances physically or functionally distant from the actual fault.

Again, note how different aspects of multiple faults and disturbance chains can be used in problem-design for staged world studies. For example, one heuristic is to interject a smaller fault that masks a portion of the second more significant fault of interest and the disturbance chains it produces.

Recognizing Anomalies

Formal accounts of diagnosis often overlook the demands of recognizing, out of many changing data channels, which changes or lack of change are anomalies, and which are less significant (or are even distractions). The inherent variability of the physical world means that changes constantly occur in the monitored process that could be relevant, in principle.

When artificial intelligence algorithms first attempted to automate diagnosis, the programs relied on people to feed them the findings to be explained (Woods et al., 1990). Historically, system developers usually labeled these as "consultants" or as "optional advice-givers"; but they were developed in fact to move toward autonomous systems (Woods, 1986a). The initial attempts to take autonomy further required connecting these programs to sensor and telemetry data from real-time processes, but these attempts quickly broke down because, without the human mediator, the programs could not distinguish which, out of the large number of data changes, were findings-to-be-explained. In other words, the programs immediately collapsed in the face of data overload (*JCS-Foundations*, pp. 79-80). Modeling anomaly response required inclusion of a mechanism for anomaly recognition as a basic demand factor given the large number of data channels and the inherent variability of the physical and human processes being monitored (Roth et al., 1992; Woods, 1994). Discovering the critical role of anomaly recognition then provided a guide in the search for new affordances to support JCSs.

Making progress on anomaly recognition first required a distinction (Woods, 1994; Woods, Patterson, & Roth, 2002): anomalies could be about discrepancies between what is observed and what is expected, or about discrepancies between observed and desired states. The second type—abnormalities—indicated the need to act and to modify courses of action (safing, contingency evaluation, replanning). The first type are violations of expectation, and it is these which act as findings-to-be-explained and trigger lines of diagnostic reasoning.

Second, studying anomaly recognition led to the insight that control of attention was a critical part of anomaly response (Woods, 1994). Attention in multi-threaded situations is inherently about balancing a trade-off (Woods, 1995b): too many irrelevant changes can be pursued (data overload), but one can discard too many potentially relevant changes as irrelevant. In the former case, performance degrades because there are too many findings to be pursued and integrated into a coherent explanation of process state. Symptoms include loss of coherent situation assessment and vagabonding (Dörner, 1983). In the latter case, performance degrades because too many potentially relevant changes are missed, rejected as irrelevant, or discounted, increasing the danger of failures to revise assessments.

Note how in developing a functional synthesis of JCSs at work in anomaly response we continue to stumble onto surprising findings. This section adds to our list: (a) diagnosis is effected as much by deciding what is to be explained as it is about the mechanisms for generating and evaluating alternative candidates, and (b) recognizing anomalies depends on violations of expectations.

The Puzzle of Expectancies

> … readiness to mark the unusual and to leave the usual unmarked—to concentrate attention and information processing on the offbeat.
>
> J. Bruner, 1990, p. 78

Our attention flows to unexpected events or departures from typicality (Chinn & Brewer, 1993). This observation reveals how meaning lies in contrasts—some departure from a reference or expected course, or differences against backgrounds (Woods, 1995a).

An event may be expected in one context and, therefore, go apparently unnoticed, but the same event will be focused on when it is anomalous relative to another context. An event may be an expected part or consequence of a quite abnormal situation, and, therefore, draw little attention. But in another context, the absence of change may be quite unexpected and capture attention because reference conditions are changing. Again, the simple story of the serendipitous auditory display in the nuclear power control room (see the case on pp. 15-16) illustrates the process. The audible clicks provided observability about the behavior of the automated system, which allowed operators to develop and demonstrate skilled control of attention. When the behavior of the automated system was expected in context, the operators did not orient to the signals (or lack of clicking sounds). Initially, outside observers could not tell that the practitioners were even noticing the sounds. Only when the auditory display indicated activity that was unexpected did the observers see explicit indications that the practitioners were using the clicking sounds as an auditory display.

A series of studies of highly skilled teams of practitioners in NASA mission control provide some insights about what practitioners find informative both before and during anomaly response (Patterson et al., 1999; Patterson & Woods, 2001; Chow et al., 2000). These studies found that events are highly prominent in the different forms of communication exchange between spaceflight controllers during missions.

For example, in an analysis of the contents of flight controllers' shift logs, Chow (2000) found that references to change and behaviors outnumbered references to states by about 3/1, and outnumbered references to base data values by nearly 20/1. Similarly, Patterson & Woods (2001), in a study of shift handovers between flight controllers noted that "practitioners rarely discussed base data values (e.g., "the pressure is 82 psi"), but rather described data in terms of event words and phrases that signified a temporal sequence of behaviors (e.g., "there was a water spray boiler freeze up")." Patterson et al. (1999) analyzed the role of the "voice loops" network over which operational groups communicate in mission control. They noted that the communication occurring over the most heavily monitored channels typically consisted of integrated descriptions of conditions, behaviors, and activities in the space shuttle systems (e.g., on the flight director's loop). By simply *listening in*, flight controllers could learn what was going on in systems outside their scope of responsibility and, thus, *anticipate* impacts on their own systems and activities. With respect to the behavior of the shuttle systems, the key point was that the work of integrating patterns of base data into descriptions of operationally significant behaviors and sequences had already been done by the controllers responsible for monitoring those sub-systems. The voice loops allowed flight controllers to absorb the context of operations at an integrated, semantic, and temporal level, i.e., in terms of events.

Note the trick the study team used to see, when they didn't know exactly what to look for. The studies were built on the assumption that experienced practitioners generally are competent meaning-recognition systems (Flach et al., 2003). This led investigators to focus the conditions of observation on points and times when practitioners need to exchange information (handovers, annotations on logs, how they monitor voice loops during anomalies, communications between groups during meetings to re-plan following an anomaly). The data capture and integration then focused on what the practitioners find significant out of large and changing data sets, either in terms of what was informative to them when monitoring their peers' activities or in terms of what they found worth communicating to their peers (in handovers or as annotations on shift logs).

Given these results from natural history studies (direct observation), Christoffersen, Woods & Blike (2006) followed up with a Staged World study of what anesthesiologists found "interesting" as they monitored a standard real-time display of sensor data during a simulated surgical scenario (note how the line of research switched natural laboratories—mission control to operating room—as well as shifted from direct observation to observing in staged world simulations). Christoffersen et al. traced how multiple attending physicians and residents recognized events from the continuous flow of telemetry data and obtained verbal protocols about what the practitioners found informative at these points. The results illustrated that experts are highly sensitive to event patterns in the underlying process, such as the general deteriorate/recovery event pattern. Similar to the studies of space mission control, Christoffersen et al. found that, by a conservative measure, a dominant proportion (fully two-thirds) of the informative features identified by the participants could be classified as event-related, especially assessments of the character of events in progress and event patterns that required integration over the behavior of multiple variables.

To see the role of expectation in anomaly recognition, consider the example in Figure 9. How does one interpret the observation of a decrease in pressure (in event words—"falling")? Is this fall an event in need of explanation? The answer depends on what has preceded and what is expected to happen next. And both of those depend on the observer's model of the influences impinging on the process (i.e., the inputs driving change), and their model of the process dynamics (i.e., the constraints on how the process should behave in response to various inputs). Changing the influence model changes the interpretation of what behavior fits and what does not, and what can happen next.

The space of variations in just this simple example in Figure 9 is huge. If, prior to interval **t,** practitioners recognized that an anomaly was in progress based in part on falling pressure, and that no actions had been taken to counteract the fall, then the decrease is expected and the continued fall is expected. If the normal response to pressure drops is for an automatic system to begin to respond in order to stabilize pressure, then the same decrease can be interpreted as part of a signature of pressure leveling off or beginning to turn around—the automatic system responded as expected and is compensating for the disturbance in pressure. If the decrease has been going on and an automatic system typically responds quickly, the same decrease becomes unexpected since the typical response would have produced a change in pressure behavior. In this case, the same behavior is unexpected and would trigger further investigation: Did the automatic system fail (and can I still restart it or is it unavailable)? Is the fault different from those the automatic system can compensate for? Is the fault larger than the capacity of the automatic system? This example illustrates how what is meaningful depends on context (Woods, 1995b; Woods et al., 2002; Theureau, 2003).

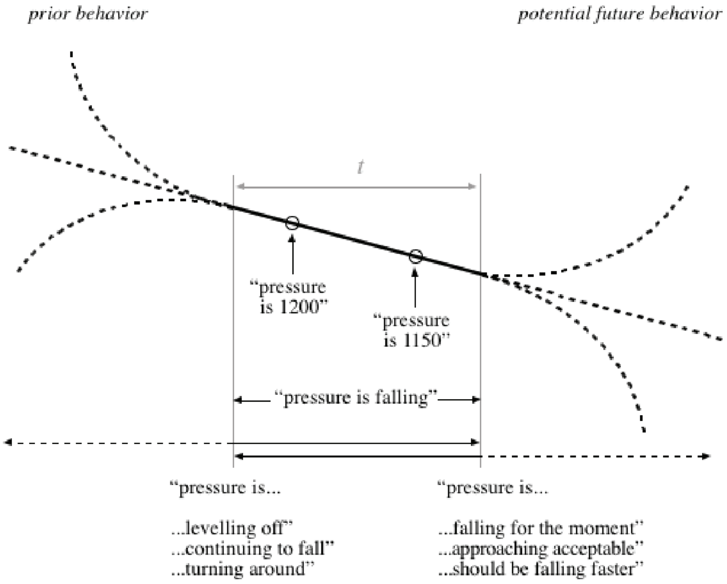

Figure 9: The same behavior in some interval **t** can be interpreted in multiple ways depending on context and expectations of the observer (From Christoffersen & Woods, 2003).

First, notice how expectations provide a way to determine what changes (or absence of changes) are informative and therefore worth practitioners' attention or worth their colleagues' attention. Using expectations to define a kind of anomalous event is essential to avoiding data overload (Woods et al., 2002). Second, expectations are defined relative to a changing model of the active influences thought to be driving the behavior of the monitored process (Figure 10).

The model of active influences may include items such as what actions have been performed on the process or what disturbances are currently at work (e.g., a malfunction in an engineered system or disease process in a medical context).

Knowledge of system dynamics serves to specify how the process ought to behave in response to the active influences. The content of these models will depend on the specific system involved, but together they combine to produce a set of more or less specific expectations and reference states against which the actual behavior of the process can be contrasted.

It is particularly important to note the critical test cases defined by the above model of anomaly recognition. One test case focuses on how practitioners would not explicitly follow-up an abnormal event with diagnostic search when the event is part of the expected sequence that follows from a fault thought to be active or part of a disturbance chain that is thought to be underway. Rather, the expected though abnormal behavior reinforces the practitioners' current assessment.

The second and more difficult test case is to recognize the absence of change as an unexpected event. Because these models of influences allow people to understand how the process should change in certain situations, it can sometimes be the case that an absence of change is itself an informative event. For example, if an action is taken to alter the trajectory of a process, the model of system dynamics can be used to determine how the process should change. If the process continues to behave in the same way subsequent to the action, this signature may define a significant event.

Christoffersen et al. (2006) found evidence for this type of event in an episode involving a lack of change in the simulated patient's blood pressure, even after several doses of medication had been administered to bring pressure down. The normal onset time for the effect of the medication is relatively short, which produced an expectation that blood pressure would quickly start to decrease. The physician then could assess the size of the decrease produced by the initial amount of medication given and make further adjustments until pressure returned to normal range. However, the study participants reacted strongly when blood pressure continued to increase (no response to the intervention), even leading some to update their model of the nature or severity of the patient's underlying problem.

The above results (cf. also Teigen & Keren, 2003) have led Christoffersen & Woods (2003) to propose a functional synthesis of anomaly and event recognition from a stream of data (Figure 10). This model helps frame a number of points about why anomaly recognition is particularly challenging (cf., also Theureau, 2003). First, the synthesis shows how anomaly and event recognition is tuned to the future—what can happen next. For the high blood pressure example, multiple interventions occurred before pressure stabilized and began to turn around. Given the extra control interventions needed to overcome the disturbances, the issue shifted to the danger of overshooting the target range and heading toward low blood pressure with the need for additional interventions to bring the parameter sufficiently under control.

Second, note how, by knowing what to look for, one notices what one was not expecting to see. This is a key criterion that tests for *observability*, or feedback that provides insight into a process—the ability to notice what one was not expecting, otherwise new technology or new designs merely make data available while leaving all of the work to find and integrate the important bits to the observer. Tukey (1977, p. vi) referring to statistical graphics, expressed the target memorably:[7]

> "The greatest value of a picture is when it forces us to notice what we never expected to see."

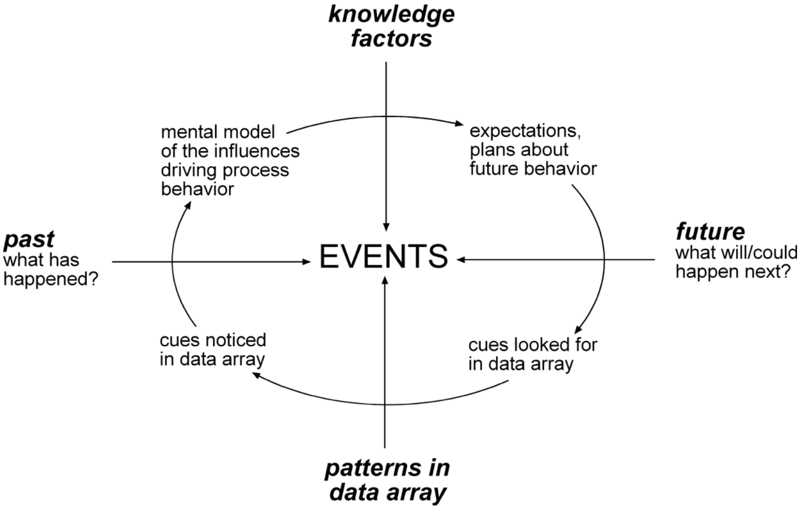

Figure 10. A model of factors underlying anomaly recognition from event patterns in a dynamic telemetry stream (From Christoffersen & Woods, 2003).

Third, expectations depend on experience by building up a highly differentiated set of what is typical behavior for different contexts (and this capability probably has more to do with processes of perceptual learning—elaborate differentiation—than with the learning mechanisms implemented in symbolic computer processing models of an individual's working memory).

In addition, what is informative goes beyond the literal surface properties of the data. As Woods has emphasized (1995a; Woods et al., 2002), what is informative depends on the behavior of the data array relative to the mindset of some observer—goals, knowledge, expectations (cf. the concept of mutuality from Gibson, 1979, discussed on p. 7). Christoffersen et al. (2006) demonstrated this in their event recognition study as the events marked by experts often did not track with simple changes in the display of a single parameter (though for observers with less expertise, the events noted more closely tracked to simple surface changes).

---

[7] The same concept is also central in Rasmussen's (1986) notion of topographic search in diagnosis.

Since the definition of what is informative depends on properties of the observer and not just properties of the world, in any given interval, people with different mindsets may isolate entirely different sets of events and focus on different anomalies from the same stimulus stream. As Ginsburg & Smith (1993) have observed in the context of social perception, the potential for divergent (but still valid) interpretations of the same stimulus stream tends to increase as the level of analysis shifts from low-level physical transitions (is pressure increasing) to event patterns (absence of a response to an intervention) to larger event sequences (a disruption to a plan in progress). Different observers may agree on the set of low-level physical changes that have occurred, but vary drastically in terms of the higher-level events that they perceive, depending on their particular mindset. Hence, the process tracing described in Chapter 5 centers on tracing how mindset is created, then shifts, or fails to shift, as situations evolve.

The temporal scale of events varies as events come nested at multiple levels of analysis (Warren & Shaw, 1985). That is, changes in stimuli defined over any given temporal interval may provide information about different events that are occurring simultaneously at widely varying timescales. For example, data indicating deteriorating performance in a subsystem onboard the space shuttle may have immediate operational significance (e.g., indicating a need to switch to a backup system to compensate for the short term). But the same event may be part of a larger pattern on the scale of the mission as a whole (e.g., at this scale the anomaly might be a loss of capability which creates an impasse that blocks execution of the current mission plan and requires replanning), or for the entire program (e.g., at this scale the anomaly might represent a long term re-design issue across the fleet that challenges the overall shuttle launch and payload schedule for months).

Typically, there is no single privileged timescale; the relevant level of analysis is a function of the observer and the context. In other words, events and anomalies are inherently systems issues (and thus, there is grave doubt whether events can be processed solely through mere data manipulation by automata) in all three of the senses that are critical to distinguishing any systems engineering (and therefore in systems engineering on cognitive systems): emergent relationships, cross-scale interactions, and sensitivity to perspective (see p. 7 in Chapter 1).

Control of Attention

> Everyone knows what attention is. It is the taking possession by the mind, in a clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought.
>
> William James, 1890, I, p. 403-404

The functional synthesis of anomaly response so far has revealed how practitioners make provisional assessments and form expectancies based on partial and uncertain data. These assessments are incrementally updated and revised as more evidence comes in, as the disturbances cascade, and as interventions occur to contain the effects of these disturbances. Situation assessment and plan revision/action are not distinct sequential stages, but rather they are closely interwoven processes with partial and provisional plan development and feedback leading to revised situation assessments. This means anomaly response inherently demands a multi-threaded process as indicated in Figure 11 (which is another elaboration on Neisser's perceptual cycle; *JCS-Foundations*, p. 20).

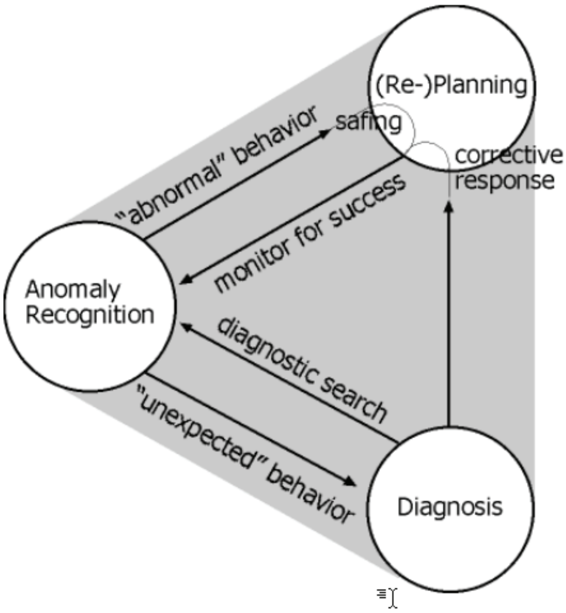

Figure 11. The multiple intermingled lines of reasoning in anomaly response (from Woods, 1994).

Practitioners need to continue to monitor for new changes; they need to look for specific information to support or disconfirm possible explanations of the unexpected findings; they need to follow up checking whether interventions produce the desired and expected results. Abnormalities require interventions, contingency evaluation and replanning activities. Anomaly response involves skillful juggling of these multiple threads. As a multi-threaded activity, anomaly response places significant demands on control of attention.

Focus of attention is not fixed, but shifts to explore the world and to track relevant changes in the world. On flight decks, in operating rooms, and in everyday work activities, attention must flow from object to object and topic to topic. In other words, one re-orients attentional focus to a newly relevant object or event from a previous state where attention was focused on other objects or on other cognitive activities (such as diagnostic search, response planning, and communication to other agents). New stimuli are occurring constantly. Sometimes such new stimuli are distractions. But other times, any of these could serve as a signal that we should interrupt ongoing lines of thought and re-orient attention.

This re-orientation involves disengagement from a previous focus and movement of attention to a new focus (Gopher, 1992).

We are able to focus, temporarily, on some objects, events, or actions in the world or on some of our goals, expectations or trains of thought *while remaining sensitive to new objects or new events that may occur*. It is the last clause that is the critical feature for control of attention—the ability to refocus without unnecessary disruptions to ongoing lines of work.

Thus, a basic challenge for any agent at work is *where to focus attention next in a changing world*. Which object, event, goal or line of thought we focus on depends on the interaction of two sets of activity as specified in Neisser's perceptual cycle (*JCS-Foundations*, p. 20). One of these is goal or knowledge directed, endogenous processes that depend on the observer's current knowledge, goals and expectations about the task at hand. The other set of processes are stimulus or data driven where attributes of the stimulus world (unique features, transients, new objects) elicit attentional capture or shifts of the observer's focus. These salient changes in the world help guide shifts in focus of attention or mindset to relevant new events, objects, or tasks.

The ability to notice potentially interesting events and know where to look next (where to focus attention next) in natural perceptual fields depends on the *coordination* between orienting perceptual systems (i.e., the auditory system and peripheral vision) and focal perception and attention (e.g., foveal vision). The coordination between these mechanisms allows us to achieve a "balance between the rigidity necessary to ensure that potentially important environmental events do not go unprocessed and the flexibility to adapt to changing behavioral goals and circumstances" (Folk et al., 1992, p. 1043).

The orienting perceptual systems function to pick up changes or conditions that are potentially interesting and play a critical role in supporting how we know where to look next. To intuitively grasp the power of orienting perceptual functions, try this thought experiment suggested by Woods & Watts (1997): put on goggles that block peripheral vision, allowing a view of only a few degrees of visual angle. Now think of what it would be like to function and move about in your physical environment with this handicap. Perceptual scientists have tried this experimentally through a movable aperture that limits the observer's view of a scene (e.g., Hochberg, 1986). Although these experiments were done for other purposes, the difficulty in performing various visual tasks under these conditions is indicative of the power of the perceptual orienting mechanisms.

Alarms and Directed Attention

Alarm Overload

"1202." Astronaut announcing that an alarm buzzer and light had gone off and the code 1202 was indicated on the computer display.
"What's a 1202?"
"1202, what's that?"
"12...1202 alarm."

> Mission control dialog over voice loops as the LEM descended to the moon during Apollo 11 (Murray & Cox, 1989).

"The whole place just lit up. I mean, all the lights came on. So instead of being able to tell you what went wrong, the lights were absolutely no help at all."

> Comment by one space controller in mission control after the Apollo 12 spacecraft was struck by lightning (Murray & Cox, 1989).

"I would have liked to have thrown away the alarm panel. It wasn't giving us any useful information."

> Comment by one operator at the Three Mile Island nuclear power plant to the official inquiry following the TMI accident (Kemeny, 1979).

"When the alarm kept going off then we kept shutting it [the device] off [and on] and when the alarm would go off [again], we'd shut it off." "... so I just reset it [a device control] to a higher temperature. So I kinda fooled it [the alarm] ... "

> Physicians explaining how they respond to a nuisance alarm on a computerized operating room device (Cook, Potter, Woods & McDonald, 1991).

The present lunar module weight and descent trajectory is such that this light will always come on prior to touchdown [on the moon]. This signal, it turns out, is connected to the master alarm—how about that! In other words, just at the most critical time in the most critical operation of a perfectly nominal lunar landing, the master alarm with all its lights, bells and whistles will go off. This sounds right lousy to me. ... If this is not fixed, I predict the first words uttered by the first astronaut to land on the moon will be "Gee whiz, that master alarm certainly startled me."

> Internal memo trying to modify a fuel warning alarm computation and how it was linked to the alarm handling system for the lunar landing craft (Murray & Cox, 1989).

A [computer] program alarm could be triggered by trivial problems that could be ignored altogether. Or it could be triggered by problems that called for an immediate abort [of the lunar landing]. How to decide which was which? It wasn't enough to memorize what the program alarm numbers stood for, because even within a single number the alarm might signify many different things. "We wrote ourselves little rules like 'If this alarm happens and it only happens once, don't worry about it. If it happens repeatedly, but other indicators are okay, don't worry about it.' " And of course, if some alarms happen even once, or if other alarms happen repeatedly and the other indicators are not okay, then they should get the LEM [lunar module] the hell out of there.

> Response to discovery of a set of computer alarms linked to the astronauts displays shortly before the Apollo 11 mission (Murray & Cox, 1989).

"I know exactly what it [an alarm] is—it's because the patient has been, hasn't taken enough breaths or—I'm not sure exactly why."

> Physician explaining one alarm on a computerized operating room device that commonly occurred at a particular stage of surgery (Cook et al., 1991).

The above are short samples of problems with alarm systems (Norros & Nuutinen, 2005), each of which points to families of story lines about affordance and clumsiness given the demands on JCSs. This set illustrates many of the patterns on how alarms can be designed that fail to support skill at control of attention in a JCS (Woods, 1995b). Most straightforwardly, alarms are mis-designed in terms of perceptual functions: Are the perceptual signals that mark abnormal conditions in a monitored process discriminable from the background and each other (e.g., Patterson 1990)? Another perceptual function of alarms is the capacity of the signal itself to attract people's attention (exogenous control of attention or alerting). This alerting potential refers to characteristics of the signal (e.g., sudden motion or sudden sound bursts against a quiescent background) to force, to a greater or lesser degree, an observer to orient to the signal itself. Apprehending these perceptual signals conveys the message that some event has occurred which may be of interest.

As support for control of attention, alarms often are mis-designed in terms of *informativeness* of the signal that something is wrong in the monitored process: Do the semantics of alarm messages inform observers about the kinds of abnormalities present or help discriminate between normal and abnormal states or between different kinds of abnormal conditions? There are several ways that alarms can be uninformative. Alarm signals and messages may be underspecified and ambiguous as in the case of the "1202" message and in several other of the above examples.

Another reason alarms may be uninformative is the context sensitivity problem. Often, what is called an "alarm" actually indicates the status of a parameter, subsystem, or component. However, current status may or may not be actually abnormal without reference to additional data which specifies the relevant context. The discussion of what the various computer alarms mean in the case above from Apollo 11 serves as a vivid account of the general problem of context sensitivity for finding what is meaningful and escaping from data overload (also see Figure 9; Woods et al., 2002).

Sometimes alarms occur frequently and seem so familiar that observers miss the cases where something truly unusual is happening (Xiao et al., 2004). In other words, the number of false alarms is high relative to the number of times the alarm signals that the abnormal condition it monitors actually is present (Sorkin & Woods 1985). The counter-intuitive result, established through these mathematical models (and follow-up empirical studies) of the impact of false alarm rates, is that increases in false alarm rates rapidly reduce the inherent informativeness of alerting signals. For example, the statistic "positive predictive value" (PPV = [True positive] / [True + False positives]) is one measure that can begin to establish whether an alarm signal has any information to convey (Getty et al., 1994).

Alarms often are mis-designed in a third way: as a collaborative agent that attempts to re-direct the attention of another—the human observer (Woods, 1995b). Directed attention is a kind of coordination across agents where one agent can perceive and direct the attentional focus of other agents to particular parts, conditions, or events in the monitored process. Directing another's attention and perceiving where another agent's attention is directed are basic aspects of human function as social and cognitive agents (an example of how cognition is public and distributed; see Hutchins 1995a). Alarms can be seen as messages from one agent, a first stage monitor, to another, a second stage supervisory agent who monitors multiple channels and whose cognitive and physical resources can be under moderate to severe workload constraints. The alarm message is an interrupt signal intended to re-direct the attention of the supervisor from their ongoing activity to examine some particular area or topic or condition in the monitored process. In effect, the attention directing signal says, “There is something that I think that you will find interesting or important; I think you should look at this.” Thus, alarm systems participate in processes of *joint reference* as signals refer to events in the monitored process in order to direct another’s attention to an event (Bruner 1986).

The receiver must use some partial information about the attention directing signal and the condition that it refers to, in order to “decide” whether or not to interrupt ongoing activities and lines of reasoning. Some attention directing signals should be ignored or deferred; similarly, some attention directing signals should re-direct attention. The quality of the control of attention is related to the skill with which one evaluates interrupt signals without disrupting ongoing lines of reasoning—knowing when the attention directing event signals “important” information and when the attention directing event can be safely ignored or deferred, given the current context.

Overall, the situation can be expressed in a signal detection theory framework (e.g., Sorkin & Woods 1985) where one can err by excessive false shifts or excessive missed shifts. Thus, in principle, two parameters are needed to describe attentional control: a *sensitivity* parameter that captures variations in skill at control of attention and a *criterion* parameter that captures tradeoffs about the relative costs and benefits of under-shifts versus over-shifts of attention. Note that framing the problem in terms of signal detection points out that even very sensitive systems for control of attention will show errors of false shifts or missed shifts of attention. Even more difficult, though, is the paradox at the heart of directed attention. Given that the receiving agent is loaded by various other demands, how does one interpret information about the potential need to switch attentional focus without interrupting or interfering with current tasks or lines of reasoning already under attentional control: How can one skillfully ignore a signal that should not shift attention within the current context, without first processing it—in which case it hasn't been ignored?

Alarms systems are rarely explicitly designed to function effectively as attention re-directors or as part of a distributed system of agents that coordinate activities as situations evolve. But go back to the case of the serendipitous auditory display in the nuclear power control room (pp. 15-16) for a successful example of the required affordance. Note the key contrast in this case that illustrates the paradox in directed attention—when the auditory display provided information that was unexpected the operators shifted attention to that part of the monitored process, but when the auditory signals were expected given ongoing conditions there were no explicit signs that the operators heard anything at all.

A study by Dugdale et al. (2000) of human–human coordination in an emergency call center illustrates the processes involved in directed attention across agents under workload. They found that directing another's attention depended on being able to see what the other agent is doing in order for one agent to be able to judge when another was interruptible. In other words, interruptibility is a joint function of the new message and the ongoing activity. This requires that one agent is able to see the activity of the other in enough detail to characterize the state of the other's activities—what line of reasoning are they on? Are they having trouble? Does their activity conform to your expectations about what the other should be doing at this stage of the task? Are they interruptible? This kind of coordination illustrates how joint reference works in two directions. In one direction an agent signals another by "referring to something with the intent of directing another's attention to it" in a mentally economical way (Bruner 1986, p. 63). In the other direction, one agent can perceive where and to what another is directing their attention, without requiring any explicit communication on the part of either agent (and the associated workload burdens). Hutchins' analysis of the joint cognitive system involved in speed control during a descent in air transport aircraft illustrates this aspect of joint reference (Hutchins 1995b). The "openness" of the representations of aircraft behavior in terms of speed management allow the physical activities associated with tasks carried out by one agent to be available for the other pilot to pick up without requiring explicit intent to communicate and without disrupting either's ongoing lines of reasoning (see also, Norman 1990).

To coordinate in processes of joint reference depends on the external representations of the monitored process by which agents assess the state of and interact with the monitored process. One refers to some part, condition, event, or task in the referent world through some shared external representation of the monitored process. As a result, one agent can direct or pick up another's focus/activity on some part, condition, event or task in the referent world in mentally economical ways.

While examining the role of alarms in anomaly response reveals stories of affordances for control of attention in multi-threaded situations, in the telling of these stories we find that the storyline naturally shifts revealing the various aspects of coordination across agents that goes on in JCSs at work.

Updating Common Ground When a Team Member Returns during Anomaly Response

The following episode of anomaly response illustrates many of the processes discussed above: cascading disturbances, interventions, expectations, events, revision, reframing, control of attention, multiple threads. In addition, these processes explicitly occur in the context of multiple agents (in this case, the interactions of two individuals in the anesthesia team; the surgical team and the anesthesiology team connections). The case is taken from an observational study of expertise in anesthesia in the operating room (Johannesen, Cook & Woods, 1994). The example also is valuable as a sample of data to be analyzed in process tracing studies of JCSs at work. As a result, we provide a complete transcript of the interactions during the episode in Appendix A.

The Case

An Attending Anesthesiologist and a senior Resident Anesthesiologist together begin a neurosurgery case to clip a cerebral aneurysm. The episode of interest begins after this case is in a maintenance phase (this type of surgery is a long process with long quiet periods for the anesthesiologists). It is about an hour after induction and before the surgeons have exposed the aneurysm. The senior resident is the only anesthesiologist present; the attending has been away for about half an hour to check on another case.

An anomaly occurs and is quickly recognized against the quiet background. The senior resident sees heart rate fall (bradycardia; the opposite of this state is rapid heart rate or tachycardia), and takes corrective action by administering atropine, a drug that raises heart rate. He has the attending paged. He mentions the event to the surgeons and enquires whether they might have been doing anything related. They answer no.

To a practitioner, the bradycardia event is quite observable. The pulse rate as indicated by the beeping of the pulse oximeter suddenly slows down. The resident, who has bent down (apparently to check the urine output or to begin a cardiac output measurement), immediately gets up to look at the vital signs on the computer monitor.

Because of its severity it is critical to treat the bradycardia immediately, before its consequences begin to propagate. Five seconds later he injects the atropine. The basic action to manage the disturbance is quick and direct. It is also important to find the source of the disturbance because the base problem may lead to other disturbances (fighting each disturbance one at a time is not likely to be the best way to manage the patient's cardiovascular system).

The bradycardia is quite unexpected given what would be typical in this situation and context, though there are a variety of situations where it would not be surprising. For example, some anesthetic medications that are used for anesthetic management during surgery (i.e., during the maintenance phase) can result in a lower than normal heart rate, but the resident knows none of these influences is ongoing. Indications of low heart rate can be due to problems with the monitoring equipment, but the anomalous event in this case is too severe.

> Upon the attending's return, this exchange occurs:
>
> Attending Anesthesiologist: {re-enters operating room mid-case}
> "Nice and tachycardic."
>
> Resident Anesthesiologist:
> "Yeah, well, better than nice and bradycardic."

Updating a Shared Frame of Reference

The situation calls for an update to the shared model of the case and its likely trajectory in the future. The background is a shared expectation that this is a quiet period of the case (about an hour after induction and before the surgeons have exposed the aneurysm). The workspace and information displays are "open," that is, open for view to any operationally knowledgeable party standing in the physical space. The attending, on entering the operating room, looks at the displays and notices the anomalous parameter—heart rate is high (tachycardia) and states the observed anomaly aloud (the statement is an invitation for an explanation of how the monitored process arrived in this situation). The response summarizes the physiological storyline—the patient became bradycardic, the resident noticed and responded, though the intervention produced an over-swing in the other direction (the comments also contain the shared model of the situation—slow heart rate is worse for this patient in this surgery). The exchange is very compact and highly coded, yet it serves to update the common ground previously established at the start of the case.

The resident goes on retelling the story of the event (including an event prior to the anomaly), context, and follow-up activities (actions taken and how rate responded). The update to the attending physician is full of event words: "down to and then back up", "dropped down to …", "so I kicked her up to …" (cf., the similar data in Patterson & Woods, 2001). The story provides information about the dynamics of the antecedent event, of the temporal flow of the event itself, and the behavior of another potentially relevant parameter (severe hypertension could produce bradycardia by a reflex pathway, but the update points out the absence of high blood pressure which rules out this mechanism for both). He mentions what action he was taking and what actions others (the surgical team) were doing ("nothing") while the event occurred, again, relating interventions (or their absence) and responses as part of the event description.

Note that this update seems to meet one general criterion for successful coordination (Patterson & Woods, 2001)—that the incoming agent can act subsequently as if they had been present throughout the evolution of the situation.

The update is also the opening to hypothesis exploration as the two anesthesiologists jointly seek explanations for the unexpected disturbance in heart rate, while at the same time the team manages cardiovascular functions to keep them within target ranges. Interestingly, the resident and attending after the update appear to be without candidate explanations as several possibilities have been dismissed given other findings (the resident is quite explicit in this case. After describing the unexpected event, he also adds—"but no explanation"). The attending tries to generate a range of hypotheses, but none seems promising—Surgeon activities? Other medications being given? Patient physiology (contractility)? Consequences of injuries from the interventions needed to place sensors for monitoring cardiovascular function ("Lines went in perfectly normal")? Low potassium?

The attending considers his experiences and points out sources that usually produce this kind of event signature (it is also important to point out that heart rate has stayed stable now with no subsequent changes)—reflexes triggered by activities associated with the surgical steps. The resident then re-examines the context the anomaly occurred in, “revisiting” what, based on the Attending's knowledge, seems to be an important time frame. As he goes back in detail over what was occurring then—they both recognize that the surgeons were engaged in an activity that could have given rise to the event—the anomaly arose at a point where the surgeons can place traction on the outermost of the three membranes (meninges) covering the brain (the Dura mater). This resolves the unexpected anomaly by providing an explanation consistent with the findings, and this explanation updates their model of influences ongoing (the source being the reflex mechanism triggered by normal surgical activities, which is an isolated factor).

While the anesthesia attending-resident coordination illustrates building common ground during an update, the anesthesia-surgical interaction illustrates a breakdown in common ground (Clark & Brennan, 1991). When the resident recognizes the anomaly, the surgeons’ activities do not appear related, and when explicitly asked (“… been doing anything”), the surgeons had answered that they had not (“No”). The surgical team’s response was to say that, from their perspective, nothing “unusual” was going on, which was true. The surgical team did not understand the perspective behind the anesthesia resident’s question—responding to an unexpected anomaly, and his line of reasoning—searching for candidate explanation. Hence, the cross-check failed (Patterson et al., 2004).

This case illustrates the interplay of multiple parties, each with partial information, overlapping knowledge and different kinds of expertise relative to the situation at hand. The case also illustrates how different perspectives play roles in revising situation assessments. The case captures the processes of *updating* as the resident calls in additional expertise (including knowing when and how to bring additional expertise in) and is able to provide concise reconstruction of events and responses that led up to the current state during the update (Patterson & Woods, 2001).

PATTERNS IN ANOMALY RESPONSE

At this point we have introduced most of the dynamic processes that intermingle in anomaly response (Figure 11 provides a composite diagram). By grounding research on observations of the phenomena of interest, anomaly response is revealed to be quite different in almost all ways from classic assumptions that diagnosis can be isolated in terms of mapping symptoms to diagnostic categories.

For example, readers familiar with abduction as a form of reasoning will recognize that the above description of anomaly response seems consistent with models of abduction (Peirce, 1955; Josephson & Josephson, 1994). Formalist approaches to abduction debate different set covering algorithms—how to evaluate the mapping between different findings and different hypotheses—and focus on criteria for judging good coverage or “best” explanation (e.g., how to operationalize criteria that define what is parsimony in an explanation). While the results on anomaly response indicate a relationship to abductive reasoning, note the surprises that result when one starts with patterns abstracted from observation. Our tour of anomaly response has revealed a variety of critical constraints on this form of work:

- Properties of how situations evolve and present themselves,
- The knowledge that interventions occur in parallel with diagnostic search, both simplifying and complicating situation assessment,
- Avoiding data overload by being able to recognize out of a large, changing data set the findings to be explained (including the role of expectations and event patterns in recognizing anomalies),
- The demands for control of attention in a multi-threaded process,
- The knowledge that revising assessments is the key to expert performance,
- Broadening the set of hypotheses considered, even as information and assessments change,
- How replanning is connected to diagnostic assessment.

To summarize: *Anomaly Response*

Anomaly response is the general form of work for JCSs, and the chapter provides a functional synthesis of the basic demands in this class of work derived from converging methods observations across multiple fields of practice. Synthesizing and abstracting patterns in anomaly response provides a set of recurring story lines about resilience, coordination, and affordance. The discovery of these patterns provides a case study on understanding JCSs at work.

Additional resources: Murray & Cox's (1989) history of mission control during the Apollo project is a compact resource that captures how control centers are adapted to the demands of anomaly response. Additional studies of how mission control works are available, such as Patterson et al. (1999), Patterson & Woods (2001), Garrett & Caldwell (2002), Mark (2002), and Shalin (2005).