

- Woods, D.D. (1986). Paradigms for decision support. In, E. Hollnagel (Ed.), *Intelligent decision support in process environments* (NATA ASI Series, Vol. F21, 153-173). Heidelberg: Springer-Verlag.
- Woods, D.D. (1995). The alarm problem and directed attention in dynamic fault management. *Ergonomics*, 38, 2371-2393.
- Yang, L.C. and Kuchar, J.K. (1997). Prototype conflict alerting logic for free flight. *Journal of guidance, control and dynamics*, 20, 768-773.

FORMAL PAPERS

## The disembodiment of data in the analysis of human factors accidents

Sidney W. A. Dekker  
Centre for Human Factors in Aviation, IKP  
Linköping Institute of Technology, Sweden

### Abstract

This article examines the theoretical underpinnings of human factors investigations in aviation. Many investigations today disembody their human factors data, either through micro-matching (behavioural fragments with after-the-fact worlds) or cherry-picking (fragments that seem to point to a common psychological condition). The reasons and mechanisms of disembodiment stem in part from conventional and technological restrictions on the gathering and analysis of human factors data. But they also have to do with our reactions to failure, which easily make investigations retrospective, proximal, counterfactual and judgmental. In addition, methodical guidance on how to map context-specific details of a complex behavioural sequence onto a conceptual description is hardly available. Instead, investigators typically rely on inarticulate folk models that make broad, unverifiable psychological assertions about the observed particulars. This article presents possible progress in the form of steps investigators can take to reconstruct the unfolding mindset of the people they are investigating, in parallel and tight connection with how the world was evolving around these people at the time.

## Introduction

Developments in digital and other recording technologies have enlarged our ability to capture human performance data—in aviation as well as other industries where people interact with technology in complex, dynamic environments. In commercial aviation, the electronic footprint that any incident or accident leaves can be huge. Recent initiatives such as the formation of a Future Flight Data Committee in the US, and regulatory proposals to lengthen recording times on cockpit voice recorders (Flight International, 1999) testify to the value we attach to accessible chronicles of human behaviour (e.g. Baker, 1999). But capturing human factors data addresses only one side of the problem. Our ability to make sense of these data, to reconstruct how humans contributed to an unfolding sequence of events, may not have kept pace with our growing technical ability to register traces of their behaviour. In other words, the growing dominance of human factors in incidents and accidents (Boeing, 1994) is not matched by our ability to analyse or understand the human contribution for what it is worth (McIntyre, 1994).

Data used in human factors accident analysis often comes from a recording of human voices and perhaps other sounds (ruffling charts, turning knobs), which can be coupled to a greater or lesser extent to contemporaneous system or process behaviour. A voice trace, however, represents only a partial data record. Human behaviour in rich, unfolding settings is much more than the voice trace it leaves behind. The voice trace always points beyond itself, to a world that was unfolding around the practitioners at the time, to tasks, goals, perceptions, intentions, thoughts and actions that have since evaporated. But aircraft accident investigations are formally restricted in how they can couple the cockpit voice recording to the world that was unfolding around the practitioners (e.g. instrument indications, automation mode settings). ICAO Annex 13 prescribes how only those data that can be factually established may be analysed in the search for cause. This provision often leaves the cockpit voice recording as only factual, decontextualized and impoverished footprint of human performance. Making connections between the voice trace and the circumstances and people in which it was grounded quickly falls outside the pale of official analysis and into the realm of what many would call inference or speculation.

Apart from the provisions of Annex 13, this problem is complicated by the fact that current flight data recorders often do not capture many automation-related traces: precisely those data that are of immediate importance to understanding the problem-solving environment in which most pilots today carry out their work. For example, FDR's in many highly automated aircraft do not record which ground-based navigation beacons were selected by the pilots, what automation mode control panel selections on airspeed, heading, altitude and vertical speed were made, or what was shown on the both pilots' moving map displays. As pilot work has shifted to the management and supervision of a suite of automated resources (Billings, 1997; Dekker and Hollnagel, 1999), and problems leading to accidents

increasingly start in human-machine interactions, this represents a large gap in our ability to access the reasons for particular human assessments and actions in cockpits.

## Situated cognition—but decontextualized investigations

The inability to make clear connections between behaviour and world straight-jackets any study of the human contribution to a cognitively noisy, evolving sequence of events. Basic findings from cognitive science and related fields keep stressing how human performance is fundamentally embedded in, and systematically connected to, the situation in which it takes place (Neisser, 1976; Gibson, 1979; Winograd and Flores, 1987; Varela *et al.*, 1995; Clark, 1997). Human actions and assessments can be described meaningfully only in reference to the world in which they are made (Winograd, 1987; Suchman, 1987); they cannot be understood without intimately linking them to details of the context that produced and accompanied them (Orasanu and Connolly, 1993; Woods *et al.*, 1994; Hutchins, 1995; Klein, 1998).

Despite this balance of scientific opinion (e.g. Willems and Raush, 1969), critical voice traces from aircraft mishaps are basically left to be examined outside the context that produced them; and outside the context in which they once carried meaning. As a result, investigators easily confuse their own reality with the one that surrounded the practitioners in question. There appear to be two broad patterns in this confusion—two ways in which human factors data are taken out of context and given meaning in relation to after-the-fact worlds. The section that follows discusses these two patterns of disembodiment. The remainder of the article explores the reasons and mechanisms behind the disembodiment of data, including retrospection and the social construction of cause, and human factors' susceptibility to folk modelling. The article then points to opportunities for progress.

## Two forms of disembodiment: Micro-matching and cherry picking

*Micro-matching: holding individual performance fragments against the background of a world we now know to be true*

Faced with complex and temporally extended behavioural situations, a dominant tactic people use is to restrict the situation under consideration by concentrating on only one isolated fragment at a time (Woods, 1993). Thus, an incident or accident sequence is typically parsed up into individual controversial decision points that can then be scrutinised for accuracy or reasonableness in isolation. This can for example be the 'decision' of a crew to continue their approach into conditions investigators now know produced a microburst (NTSB, 1995). Rather than understanding these controversial fragments in relation to the circumstances

that brought them forth, and in relation to the stream of preceding as well as succeeding behaviours which surrounded them, the performance fragment is held up against a world investigators now know to be true. The problem is this after-the-fact-world may have very little relevance to the actual world that produced the behaviour under investigation. The behaviour is contrasted against the investigator's reality, not the reality surrounding the behaviour in question. There are various ways in which such after-the-fact-worlds can be brought into being; three are mentioned here.

First, individual fragments of behaviour are frequently contrasted with written guidance, which can be found to have been applicable in hindsight. Compared with such written guidance, actual performance is often found wanting; it does not live up to procedures or regulations. For example, 'One of the pilots...executed (a computer entry) without having verified that it was the correct selection and without having first obtained approval of the other pilot, contrary to AA's procedures.' Investigations invest considerably in organisational archaeology so that they can construct the regulatory or procedural framework within which the operations took place. Inconsistencies between existing procedures or regulations and actual behaviour are easy to expose when organisational records are excavated after-the-fact and rules uncovered that would have fit this or that particular situation. This is not, however, very informative. There is virtually always a mismatch between actual behaviour and written guidance that can be located in hindsight (Suchman, 1987; Woods *et al.*, 1994). Pointing that there is a mismatch sheds little light on the *why* of the behaviour in question. And for that matter, mismatches between procedures and practice are not unique to mishaps (Degani and Wiener, 1991).

Second, to construct the world against which investigators hold individual performance fragments, can exist of elements in the situation that were not picked up by the practitioners (Endsley, 1999), but that, in hindsight, proved critical. Take the turn towards the mountains on the left that was made just before an accident near Cali, Colombia in 1995 (Aeronautica Civil, 1996). What should the crew have seen in order to notice the turn? They had plenty of indications, according to the manufacturer of their aircraft:

'Indications that the airplane was in a left turn would have included the following: the EHSI (Electronic Horizontal Situation Indicator) Map Display (if selected) with a curved path leading away from the intended direction of flight; the EHSI VOR display, with the CDI (Course Deviation Indicator) displaced to the right, indicating the airplane was left of the direct Cali VOR course, the EADI indicating approximately 16 degrees of bank, and all heading indicators moving to the right. Additionally the crew may have tuned Roza in the ADF and may have had bearing pointer information to Roza NDB on the RMDI' (Boeing, 1996, p. 13).

This is a standard response after mishaps: point to the data that would have revealed the true nature of the situation (Woods, 1995). Knowledge of the 'critical' data comes only with the omniscience of hindsight, but if data can be shown to have been physically available, it is assumed that it should have been

picked up by the practitioners in the situation. The problem is that pointing out that it should have does not explain why it was not, or why it was interpreted differently back then (Weick, 1995). There is a dissociation between data availability and data observability (Woods, 1995)—between what can be shown to have been physically available and what would have been observable given the multiple interleaving tasks, goals, attentional focus, interests, and—as Vaughan (1996) shows—culture of the practitioner.

Third, there are less obvious or documented standards. These are often invoked when a controversial fragment (e.g. a decision to accept a runway change (Aeronautica Civil, 1996; or the decision to go around or not, NTSB, 1995) knows no clear pre-ordained guidance but relies on local, situated judgement. For these cases there are always 'standards of good practice' which are based on convention and putatively practised across an entire industry. One such standard in aviation is 'good airmanship', which, if nothing else can, will explain the variance in behaviour that had not yet been accounted for.

Investigators micro-match controversial fragments of behaviour with standards that seem applicable from their after-the-fact position. The problem is that these after-the-fact-worlds may have very little relevance to the circumstances of the accident sequence, and that the investigator has substituted his own world for the one that surrounded the practitioners in question.

#### *Cherry picking: grouping similar performance fragments under a label identified in hindsight*

Grouping individual fragments of human performance that *prima facie* represent some common condition, is the second pattern in which meaning is imposed on available data from the outside and from hindsight. Consider the following example, where diverse fragments of performance—that are not temporally co-located but spread out over half an hour—are lumped together to build a case for haste as explanation of the (in hindsight) bad decisions taken by the crew:

'Investigators were able to identify a series of errors that initiated with the flightcrew's acceptance of the controller's offer to land on runway 19...The CVR indicates that the decision to accept the offer to land on runway 19 was made jointly by the captain and the first officer in a 4-second exchange that began at 2136:38. The captain asked: 'would you like to shoot the one nine straight in?' The first officer responded, 'Yeah, we'll have to scramble to get down. We can do it.' This interchange followed an earlier discussion in which the captain indicated to the first officer his desire to hurry the arrival into Cali, following the delay on departure from Miami, in an apparent to minimise the effect of the delay on the flight attendants' rest requirements. For example, at 2126:01, he asked the first officer to 'keep the speed up in the descent'... The evidence of the hurried nature of the tasks performed and the inadequate review of critical information between the time of the

flightcrew's acceptance of the offer to land on runway 19 and the flight's crossing the initial approach fix, ULQ, indicates that insufficient time was available to fully or effectively carry out these actions. Consequently, several necessary steps were performed improperly or not at all' (Aeronautica Civil, 1996, p. 29).

It is easy to pick through the voice record of an accident sequence and find fragments that all seem to point to a common condition. The investigator treats the voice record as if it were a public quarry to select stones from, and the accident explanation the building he needs to erect. The problem is that each fragment is meaningless outside the context that produced it: each fragment has its own story, background, and reasons for being, and when it was produced it may have had nothing to do with the other fragments it is now grouped with. Also, behaviour takes place in between the fragments. These intermediary episodes contain changes and evolutions in perceptions and assessments that separate the excised fragments not only in time, but also in meaning.

Thus, the condition that binds *prima facie* similar performance fragments arises not from the circumstances that brought each of the fragments forth; it is not a feature of those circumstances. It is an artefact of the investigator. In the case described above, 'hurry' is a condition identified in hindsight, one that plausibly couples the start of the flight (almost 2 hours behind schedule) with its fatal ending (on a mountainside rather than an airport). 'Hurry' is a retrospectively invoked *leitmotif* that guides the search for evidence about itself. This leaves an investigation not with findings, but with tautologies.

## Investigations as social reconstructions

### *Retrospective*

Investigations aim to explain the past. Yet they are conducted in the present, and thus inevitably influenced by it. One safe bet is that investigators know more about the mishap than the people who were caught up in it. Investigators and other retrospective observers know the true nature of the situation surrounding the people at the time (where they were versus where they thought they were, what mode their system was in versus what mode they thought it was in, and so forth). Investigators also know the outcome and inevitably evaluate people's assessments and actions in the light of it. The effects of this hindsight bias have been well-documented and even reproduced under controlled circumstances (Fischhoff, 1975). One effect is that 'people who know the outcome of a complex prior history of tangled, indeterminate events remember that history as being much more determinant, leading 'inevitably' to the outcome they already knew' (Weick, 1995, p28). Hindsight allows us to change past indeterminacy and complexity into order, structure, and oversimplified causality (Reason, 1990). The problem is that structure is itself an artefact of hindsight—only hindsight can turn ill-structured, confusing and blurred events into

clear decision paths with junctures which shows where practitioners went the wrong way. Second, to explain failure, we seek failure. We seek the incorrect actions, the flawed analyses, and the inaccurate perceptions, even if these were not thought to be influential or obvious at the time (Starbuck and Milliken, 1988). This search for people's failures is the result of the hindsight bias: knowledge of outcome influences how we see a process. If we know the outcome was bad, we can no longer objectively look at the behaviour leading up to it—it must also have been bad (Fischhoff, 1975; Woods *et al.*, 1994).

### *Proximal, counterfactual and judgmental*

It is not just knowledge of outcome and wider circumstance that colours the interpretation of past behavioural data. Investigations are governed by implicit goals that go beyond merely 'understanding' what went wrong and preventing recurrence. Mishaps are surprising relative to prevailing beliefs and assumptions about the system in which they happen (Wagenaar and Groeneweg, 1987). Thus, investigations are inevitably affected by the concern to reconcile a disruptive event with existing views and beliefs about the system. In this process of reconciliation, something has to give—either the event, or the existing beliefs about the system.

In the immediate aftermath of failure, people and institutions may be willing to question their underlying assumptions about the system they use or operate. Perhaps things are not as safe as previously thought; perhaps the system contains error-producing conditions that could have spawned this kind of failure earlier, or worse, could do it again. This openness does not typically last long. As the shock of an accident subsides, parts of the system mobilise to contain systemic self-doubt and change the fundamental surprise into a merely local hiccup that temporarily ruffled an otherwise smooth operation. The reassurance is that the system is basically safe—it's only some people or other parts in it that are unreliable. In the end, it is not often that an existing view of a system gives in to the reality of failure. Instead the event, or the players in it, are changed to fit existing assumptions and beliefs about the system, rather than the other way around. Expensive lessons about the system as a whole, and the subtle vulnerabilities it contains, can go completely unlearned (e.g. Rasmussen and Batstone, 1989; Wilkinson, 1994).

Human factors investigations must be understood against the backdrop of the 'fundamental surprise error' (Lanir, 1986) and examined for the role they play in it. The inability to deal with the fundamental surprise of a failure shines through investigations that are:

- proximal: they concentrate on local perpetrators: those closest to 'causing' and potentially preventing the failure are seen as the sole engine of action;
- counterfactual: they state what these practitioners could have done to prevent the outcome;
- judgmental: they state how practitioners 'failed' to take these actions or notice data that they should have noticed.

An investigation's emphasis on the proximal ensures that the mishap remains the result of a few uncharacteristically ill-performing individuals who are not representative of the system or the larger practitioner population in it. It leaves existing beliefs about the basic safety of the system intact.

The pilots of a large military helicopter that crashed on a hillside in Scotland in 1994 were found guilty of gross negligence. The pilots did not survive—29 people died in total—so their side of the story could never be heard. The official inquiry had no problems with 'destroying the reputation of two good men', as a fellow pilot put it. Potentially fundamental vulnerabilities (such as 160 reported cases of Uncommanded Flying Control Movement or UFCM in computerised helicopters alone since 1994) were not looked into seriously (Sunday Times, 25 June 2000).

Counterfactuals enumerate the possible pathways that practitioners could have taken to prevent the mishap (the pathways that look so obvious in hindsight). For example:

'... the airplane could have overcome the windshear encounter if the pitch attitude of 15 degrees nose-up had been maintained, the thrust had been set to 1.93, and the landing gear had been retracted on schedule' (NTSB, 1995, p. 119).

Much investigative energy is invested in proving what could have happened if certain minute or utopian conditions had been met, and the result of these efforts is often accepted as explanatory substitute. Counterfactuals may be fruitful in exploring potential countermeasures against the same kind of failure. But counterfactuals explain nothing. Laying out what did not happen does not explain what happened. Counterfactuals also have no empirical justification, as there are no events in the world that beg their kind of illumination.

The step from counterfactual to judgement is a small one, and often taken by investigations. Judgements occur when counterfactual pathways are held up as the ones that *should* have been taken, and they slip into investigations surprisingly easily. Take as an example the probable causes of the previously alluded 1994 aircraft accident at Charlotte, N.C. Among the probable causes were '(2) the flightcrew's failure to recognize a windshear situation in a timely manner; (3) the flightcrew's failure to establish and maintain the proper airplane attitude and thrust setting necessary to escape from windshear' (NTSB, 1995, p. 120). Instead of explaining events from the point of view of the pilots inside of them, the investigation lays out counterfactual pathways and stresses that they should have been taken—which is obvious from hindsight—by asserting that not doing so constituted a failure. None of this explains the practitioners' actual behaviour.

Counterfactual and judgmental language, typical of aircraft accident investigations, justify their proximal emphasis. If only these people had done something different (which was so obvious! How could they have missed it!), then the accident would not have happened. The reason for failure has been located. It truncates the need for deeper probing into the systemic conditions that perhaps laid the groundwork for controversial proximal assessments and actions in the first place.

### *The social construction of cause*

The fundamental surprise error means that events get picked apart and re-inscribed to fit existing assumptions about the system in which it happened, rather than the other way around. Investigations contribute to the error through the disembodiment of data, their proximal focus and counterfactual and judgmental language. As a result, different and mutually exclusive accident investigations can emerge that each emphasise and rely on their own sets of 'cherries' from a single sequence of events (see table 1 for an example)—leaving the larger system or industry none the wiser.

**Table 1** Different statements of cause about the same accident. Compiled from Aeronautica Civil (1996) and Flight Safety Foundation (1997).

The Colombian authorities say:	The airline says:
The Cali approach controller did not contribute to the cause of the accident	Approach control clearances were not in accordance with ICAO standards The approach controller's inadequate English and inattention during a critical phase of the approach were causal
... and among the causes are:	... and among the causes are:
Inadequate use of automation	Inadequate navigational database
The lack of pilot situation awareness regarding terrain and navigation aids	Lack of radar coverage over Cali
Failure to revert to radio navigation when FMS created confusion and workload	Manufacturer's/vendor's over-confidence in FMS technology and resultant influence passed onto pilots regarding FMS capabilities
The flight crew's failure to adequately plan and execute the approach to runway 19	Failure of those responsible to ensure that FMS database matched industry advisory
Ongoing efforts to expedite the approach to avoid delays	The crew's task overload caused by unexpected change in assigned runway

The reason that investigations have so many causes to choose from is that the systems that are vulnerable to failure are so well-protected against it. A lot needs to go wrong for a system to be pushed over the edge of breakdown (e.g. Reason, 1997). When tracing a failure chain, causal webs quickly spread out, like cracks in a window. The pattern of contributions to any complex system failure is dense, and 'primary' or 'root' causes that are found in it are arbitrary choices or constructs by the one doing the looking. The example, above shows that the selection or construction of cause is determined in large part by organisational or sociological factors that lie outside the actual events. It also shows that accuracy or comprehensiveness are not criteria for a successful explanation, but plausibility is—plausibility from the point of view of those who have to accommodate the surprise that the failure represents for *them* and *their* organisation. Accident investigations have this purpose to fulfil, even if they become selective oversimplifications because of it. Even if—in the words of Weick (1995)—they make lousy history.

### Folk models

Another reason for the disembodiment of data is the paucity of agreed-upon methods for matching the behavioural particulars of an accident sequence with articulated, theory-driven models of human performance. In this sense, human factors data analysis lags behind the investigation of, say, structural crash evidence where recovered bits and pieces can be contrasted against quantitative models of component performance or tests of identical components under similar circumstances (McIntyre, 1994). Conclusions and causal statements flowing forth from such investigations have a measure of reliability—they could in principle be replicated. In contrast, human factors investigators are left to draw inferences and produce ad-hoc assertions that bear some relationship with an ill-defined psychological or sociological phenomenon. For example:

- 'Deficient situation awareness is evident' (Aeronautica Civil, 1996, p. 34)
- 'The CRM (Crew Resource Management) of the crew was deficient' (Aeronautica Civil, 1996, p. 47)
- There was a 'casual atmosphere' (NTSB, 1994, p. 106)
- 'They had lost situation awareness and effective CRM' (Aeronautica Civil, 1996, p. 48).

Such 'explanations' frequently make it into conclusions and statements of cause. For example, 'Aeronautica Civil determines that the probable causes of this accident were:...(3) the lack of situational awareness of the flightcrew...' (1996, p. 57).

What does not help is that the human performance models investigators use, or have available (e.g. 'loss of situation awareness or loss of CRM') are often theory-begging folk models that do little more than parroting popular

contemporary consensus between experts and non-experts on the nature of an everyday phenomenon, for example getting lost or confused (Hollnagel, 1998). In response to the increasing dominance of human factors in accidents over the past decades, investigators (as well as researchers) have introduced, loaned or overgeneralised concepts that try to capture critical features of individual or coordinative crew behaviour (for example workload, complacency, stress, situation awareness). While easily mistaken for deeper insight into human factors issues, these concepts can create more confusion than clarity. A credible and detailed mapping between the context-dependent (and measurable) particulars of an accident sequence and pertinent conceptual models is often lacking. The jump from accident details to conceptual conclusions is typically a single and large one, supported by the kind of cherry-picking described above—which immunises it against critique or verification. This is similar to the issue faced by psychological field studies, where the translation from context to concept needs to go through various steps or levels in order for those without access to the original situation to trace and appreciate the conceptual description (Hollnagel, Pederson and Rasmussen, 1981; Woods, 1993).


The problem is that folk models typically lack the human performance measures or probes that would be necessary to reach down into the context-specific details, because they postulate no underlying psychological theory that could deliver any (Hollnagel, 1998). For example, when one pilot asks the other 'where are we?' (Aeronautica Civil, 1996, p. 33), this may be a clear instance of a loss of situation awareness to a lay observer. But there is nothing inherent to models of situation awareness (e.g. Endsley, 1999) that dictates that this would be so: the model proposes no performance measurement based on an underlying psychological theory, i.e. that asking a question involving direction indicates a loss of situation awareness.

Human factors can offer theory-based, well-articulated conceptual models of human performance. But, similar to the issue that confronts field experimentation (Hoffman and Woods, 2000), problems occur because there is little guidance to help investigators establish the mapping between the particulars of context-bound behaviour in an accident episode on the one hand and models or descriptions of human performance that offer a sufficient level of conceptual detail for those kinds of settings on the other (see Woods, 1993 and Klein, 1998 for some achievement in this respect). The lack of guidance is surprising in an era where human factors are judged to be the dominant contributor to system failure (e.g. NTSB, 1994; Boeing, 1994), yet some progress is being made (see Dekker, in press).

### Progress on human factors investigations

How do we prevent the disembodiment of data? In order to understand the actual meaning that data had at the time and place it was produced, investigators need to

step into the past themselves (Vaughan, 1996). According to Tuchman (1981): 'Every scripture is entitled to be read in the light of the circumstances that brought it forth. To understand the choices open to people of another time, one must limit oneself to what they knew; see the past in its own clothes, as it were, not in ours' (p. 75). When left in the context that produced and surrounded it, human behaviour is inherently meaningful. The challenge for an investigator is to reconstruct the unfolding situation as it looked to the people whose controversial actions and assessments are under investigation. Taking the perspective of people inside the situation means adopting the local rationality principle: these people were doing reasonable things given their point of view, their available knowledge, their objectives and their limited resources (De Keyser and Woods, 1990; Woods, Johannesen, Cook and Sarter, 1994).



**Figure 1** The challenge for an investigator is to reconstruct the situation as it unfolded from the point of view of the people inside of it—free from the biases of the perspective of retrospective outsider.

The local rationality principle implies that investigations should not try to find where people went wrong, but rather understand how people's assessments and actions made sense at the time, given the circumstances that surrounded them.

This requires the investigator to build a picture of:

- how a process and other circumstances unfolded around people;
- how people's assessments and actions evolved in parallel with their changing situation;
- how features of people's tools and tasks and organisational and operational environment influenced their assessments and actions.

In other words, the reconstruction of a practitioner's unfolding mindset begins not with the mind. It begins with the situation in which that mind found itself. Note how this confirms the basic findings about situated cognition: if we understand the circumstances in which human cognition took place, we will begin to understand those cognitive activities. Such reconstruction also diminishes an

investigator's dependency on large psychological (folk) labels, instead creating a more direct coupling between (possible) perceptions and actions. Here are some concrete steps that investigators can take (see for more detail: Dekker, in press):

#### *Mark the beginning and end of a sequence of events*

This may seem an obvious step to take in any analysis—bound the event under investigation by marking the start and the finish. Yet many investigations do not explicitly say where in a sequence of events their work really begins and where it ends. One reason is the inherent difficulty in deciding what counts as the beginning (especially the beginning—the end of a sequence of events often speaks for itself). But investigations can take as their beginning the first assessment, decision or action by people close to the mishap—the one that, according to the investigator, set the sequence of events in motion. Such a decision may be the pilot's acceptance of a runway change that led to trouble later on. Such a proximal assessment or action can be seen as a trigger for the unfolding events that follow. Of course the trigger itself has reasons, a background, that extends beyond the mishap sequence—both in time and in place. The whole point of taking a proximal assessment or action as starting point is not to ignore these backgrounds, but to identify concrete points to begin an investigation into them.

#### *Lay out the junctures in a sequence of events*

This step 'digitises' an analogue history of events to make further analysis possible. In the meandering sequence of events towards an outcome, an investigator can locate junctures ('critical decision points, see Klein, 1998). These are points where for example:

- The sequence of events took a turn towards the outcome.
- The sequence of events momentarily veered away from the outcome.
- The sequence of events could have taken a turn away from the outcome altogether but did not.


What counts as a juncture? As a rule, what people did and what their processes did is tightly interconnected—the two rarely develop independently from one another. Where the process makes its contributions (e.g. an automation mode change) people can get different insights, come to different conclusions or move towards particular decisions. Which in turn may influence how the process is managed. This means that discovering changes in one may lead onto a juncture in the other. Junctures in a sequence of events towards failure can be identified by cross-examining people's decisions, cognitive resets (DeKeyser and Woods, 1990), shifts in behaviour or strategy, actions to influence the process, and changes in the process itself.

Junctures are starting points for investigating the backgrounds, reasons and histories behind them. Where did decisions come from? What pushed them one



way rather than the other? In other words, these junctures form the organising thread, for reconstructing the situation that surrounded the people whose assessments and actions are being investigated. Such deeper investigation of a juncture can also make it go away. For example, people seem to decide, in the face of evidence to the contrary, to not change their course of action; to continue with their plan as it is. With hindsight, one may see that people had opportunities to recover from their misunderstanding of the situation, but missed the cues, or misinterpreted them. These 'decisions' to continue, these opportunities to revise, are junctures only in hindsight. To the people caught up in the sequence of events there was not any compelling reason to re-assess their situation or decide against anything. Or else they would have. They were doing what they were doing because they thought they were right; given their understanding of the situation; their pressures.

The challenge becomes to understand how this was not a juncture to the people in question; how their 'decision' to continue was nothing more than continuous behaviour—reinforced by their current understanding of the situation, confirmed by the cues they were focusing on, and reaffirmed by their expectations of how things would develop in the near future.



**Figure 2** Connecting critical process parameters to the sequence of assessments and actions

#### *Reconstruct the situation at each juncture*

How was the world unfolding around people? In an effort to couple cognition to the situation in which it took place, investigators can begin with laying out the critical process parameters and how they were changing over time, both as a result of human influences and of the process moving along. The values of these

parameters were likely available to people in all kinds of ways—dials, displays, knobs that pointed certain ways, sounds, mode annunciators, alarms, warnings.

Presenting critical parameters is not new to investigations. Many accident report appendices contain read-outs from data recorders, which show the graphs of known and relevant process parameters. But building these pictures is often where investigations stop today. Tentative references about connections between known parameters and people's assessments and actions are sometimes made, but never in a systematic, or graphic way. The point of step three is to marry all the junctures identified above with the unfolding process—to begin to see the two in parallel, as an inextricable, causal *dance-a-deux*. The point is to build a picture that shows these connections (see figure 2).

As said before, there is a difference between data availability and data observability. Out of the critical parameters what did people actually notice? Where did they look? How did they interpret their evolving situation? The next step can shed some light on this.

#### *Identify tasks and goals*

People do not wander through situations aimlessly, simply receiving inputs and producing outcomes as they go along. They are there to get a job done, to accomplish tasks, to pursue goals. If there is anything that determines where people look and how they interpret what they see, it is the goals that they have at the time, and the tasks they are trying to accomplish. Finding what tasks people were working on does not need to be difficult. It often connects directly to how process parameters were unfolding around them.

- What is canonical, or normal at this time in the operation? Tasks relate in systematic ways to stages in a process;
- What was happening in the managed process? Changes here obviously connect to the task people were carrying out;
- What were other people in the operating environment doing? People who work together on common goals often divide the necessary tasks among them in predictable or complementary ways.

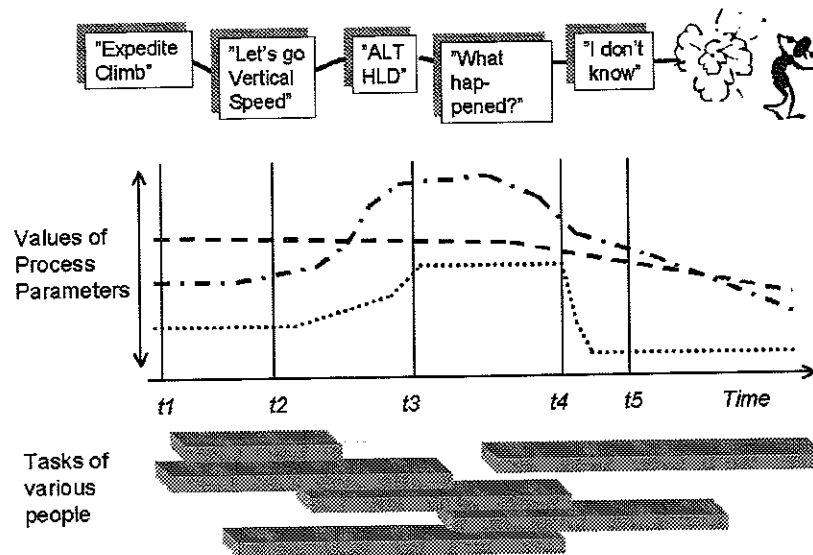
Tasks and goals pull red threads through the previously digitised assessments and actions; they connect junctures in coherent and meaningful ways. What people saw or did at any one point is determined in part by past assessments of the situation, and expectations of how it would develop in the future.

#### *Identify other influences on assessments and actions*

Human behaviour is determined by many more factors than process parameters. As a rule, however, other influences are less visible and more difficult to recover from the evidence of a mishap. Take organisational pressures to choose schedule over safety, for example. Such pressures exist and exert a powerful influence on



the many little local trade-offs people make. Yet especially in the aftermath of failure, these factors easily get rationalised away as being irrelevant or insignificant. As in: real professionals should not be susceptible to those kinds of pressures. Such reactions, however, reveal a profound shortcoming in the understanding of human error. Efforts have been made (e.g. Reason, 1997; Dekker, in press) to direct investigators and managers to those features of an operation and organisation that contribute systematically, but subtly, to the assessments and actions that people make on the line.



**Figure 3** Laying out the tasks that people tried to accomplish during the sequence of events

### Conclusion

The treatment and analysis of human factors data in aircraft accident investigations is still a young and somewhat uncertain activity. Micro-matching (fragments of behaviour with after-the-fact worlds) and cherry-picking (fragments that seem to fit a common psychological condition) as investigative practices divorce human factors data from the context in which they were produced. Methodical guidance on how to map context-specific details of a complex behavioural sequence onto a conceptual description instead, is hardly available, and investigators typically rely on inarticulate folk models that make broad, unverifiable psychological assertions about the observed particulars. Also,

investigations inevitably have to deal with the surprising nature of failures, easily making them retrospective, proximal, counterfactual and judgmental.

Conventional and technological restrictions also play a role in the disembodiment of data. Further refining of flight recorder parameters, especially on automated flight decks, will help capture those traces necessary to understand the human contribution to most failures today. And a reconsideration of ICAO Annex 13 in light of the increasing popularity of human factors—making that human behaviour and the world in which it was embedded can formally be coupled together during an investigation—is another good start. In the words of McIntyre (1994, p. 18):

'Human factors data should be subject to the same rules of evidence that are applied to other mishap data. Findings must relate to factors which were significant in the mishap scenario and be integrated with other factual evidence.'

### References

- Aeronautica Civil (1996). *Aircraft accident report: Controlled flight into terrain, American Airlines flight 965, Boeing 757-223, N651AA near Cali, Colombia, December 20, 1995*. Bogota, Colombia: Aeronautica Civil.
- Baker, S. (1999). Aviation incident and accident investigation. In, D. J. Garland, J. A. Wise, and V. D. Hopkin (Eds.). *Handbook of aviation human factors*, (pp. 631-642). Mahwah, NJ: Erlbaum.
- Billings, C. E. (1997). *Aviation automation: The search for a human-centered approach*. Mahwah, NJ: Erlbaum.
- Boeing Commercial Airplane Group (1994). *Statistical summary of commercial jet aircraft accidents: Worldwide operations 1959-1993* (Boeing Airplane Safety Engineering B-210B). Seattle, WA: Boeing.
- Boeing Commercial Airplane Group (1996). *Boeing submission to the American Airlines Flight 965 Accident Investigation Board*. Seattle, WA: Boeing.
- Clark, A. (1997). *Being there: Putting brain, body and world together again*. Cambridge, MA: MIT Press.
- Degani, A. and Wiener, E. L. (1991). Philosophy, policies and procedures: The three P's of flightdeck operations. *Paper presented at the Sixth International Symposium on Aviation Psychology, Columbus, OH, April*.
- De Keyser, V. and Woods, D. D. (1990). Fixation errors: Failures to revise situation assessment in dynamic and risky systems. In, A. G. Colombo and A. Saiz de Bustamante (Eds.). *System reliability assessment* (pp. 231-251). The Netherlands: Kluwer Academic.
- Dekker, S. W. A. (in press). *The field guide to human error*. Bedford, UK: Cranfield University Press.
- Dekker, S. W. A. and Hollnagel E. (1999). *Coping with computers in the cockpit*. Aldershot, UK: Ashgate.

- Endsley, M. R. (1999). Situation Awareness in aviation systems. In, D. J. Garland, J. A. Wise and V. D. Hopkin (Eds.), *Handbook of aviation human factors* (pp. 257-276). Mahwah, NJ: Erlbaum.
- Fischhoff, B. (1975). Hindsight is not foresight: The effect of outcome knowledge on judgement under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance*, 1, 288-299.
- Flight Safety Foundation (1997). Preparing for a last-minute runway change, Boeing 757 flight crew loses situational awareness, resulting in collision with terrain. *FSF Accident Prevention, July-August* (pp. 1-23). Washington, DC: FSF.
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston, MA: Houghton-Mifflin.
- Hoffman, R. R. and Woods, D. D. (2000). Studying cognitive systems in context. *Human factors*, 42, 1-7.
- Hollnagel, E., Pederson, O. M. and Rasmussen, J. (1981). *Notes on human performance analysis* (Tech. Rep. Riso-M-2285). Denmark: Riso National Laboratory.
- Hollnagel, E. (1998). Measurements and models, models and measurements: You can't have one without the other. In, *Proceedings of NATO AGARD conference*, Edinburgh, UK.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.
- Klein, G. (1998). *Sources of power: How people make decisions*. Cambridge, MA: MIT Press.
- Langewiesche, W. (1998). *Inside the sky*. New York, NY: Random House.
- Lanir, Z. (1986). *Fundamental surprise*. Eugene, OR: Decision Research.
- McIntyre, J. A. (1994). Perspectives on human factors: The ISASI perspective. *Forum*, 27, 18.
- Moshansky, V. P. (1992). *Commission of inquiry into the Air Ontario accident at Dryden, Ontario* (Final report, Vols. 1-4). Ottawa, ON: Minister of Supply and Services, Canada.
- National Transportation Safety Board (1994). *Safety study, A review of flightcrew-involved major accidents of U.S. air carriers, 1978 through 1990* (NTSB/SS-94/01). Washington, DC: NTSB.
- National Transportation Safety Board (1995). *Aircraft accident report: Flight into terrain during missed approach, USAir flight 1016, DC-9-31, N954VJ, Charlotte Douglas International Airport, Charlotte, North Carolina, July 2, 1994* (NTSB/AAR-95/03). Washington, DC: NTSB.
- Neisser, U. (1976). *Cognition and reality*. San Francisco, CA: Freeman.
- Orasanu, J., and Connolly, T. (1993). The reinvention of decision making. In, G.A. Klein, J. Orasanu, R. Calderwood and C. E. Zsombok (Eds.), *Decision making in action: Models and methods* (pp. 3-20). Norwood, NJ: Ablex.
- Perrow, C. (1984). *Normal accidents*. New York, NY: Basic Books.
- Rasmussen, J. and Batstone, R. (1989). *Why do complex organizational systems fail?* Environment Working Paper No. 20. Washington, DC: World Bank.
- Reason, J. (1990). *Human error*. Cambridge, UK: Cambridge University Press.
- Reason, J. (1997). *Managing the risks of organizational accidents*. Aldershot, UK: Ashgate.
- Starbuck, W. H. and Milliken, F. J. (1988). Challenger: Fine-tuning the odds until something breaks. *Journal of management studies*, 25, 319-340.
- Suchman, L. A. (1987). *Plans and situated actions: The problem of human-machine communication*. Cambridge, UK: Cambridge University Press.
- Tuchman, B. W. (1981). *Practicing history: Selected essays*. New York, NY: Norton.
- Varela, F. J., Thompson, E. and Rosch, E. (1995). *The embodied mind: Cognitive science and human experience*. Cambridge, MA: MIT Press.
- Vaughan, D. (1996). *The Challenger launch decision*. Chicago, IL: University of Chicago Press.
- Wagenaar, W. A. and Groeneweg, J. (1987). Accidents at sea: Multiple causes and impossible consequences. *International Journal of Man-Machine Studies*, 27, 587-589.
- Weick, K. (1995). *Sensemaking in organizations*. London: Sage.
- Willems, E. P. and Raush, H. L. (Eds.) (1969). *Naturalistic viewpoints in psychological research*. New York, NY: Holt, Rinehart and Winston.
- Winograd, T. and Flores, F. (1987). *Understanding computers and cognition*. Reading, MA: Addison-Wesley.
- Winograd, T. (1987). *Three responses to situation theory* (Technical report CSLI-87-106). Stanford, CA: Center for the study of language and information, Stanford University.
- Woods, D. D. (1993). Process-tracing methods for the study of cognition outside of the experimental laboratory. In, G.A. Klein, J. Orasanu, R. Calderwood and C. E. Zsombok (Eds.), *Decision making in action: Models and methods* (pp. 228-251). Norwood, NJ: Ablex.
- Woods, D. D., Johannesen, L. J., Cook, R. I. and Sarter, N. B. (1994). *Behind human error: Cognitive systems, computers and hindsight*. Dayton, OH: CSERIAC.

## Acknowledgement

Work for this article was supported by a grant from the Swedish Flight Safety Directorate.