ROUTLEDGE REVIVALS

Coping with Computers in the Cockpit

Edited by

Sidney Dekker and Erik Hollnagel

ROUTLEDGE REVIVALS

Coping with Computers in the Cockpit

Edited by
SIDNEY DEKKER AND ERIK HOLLNAGEL
Linköping University, Sweden

Routledge
Taylor & Francis Group
LONDON AND NEW YORK

First published 1999 by Ashgate Publishing

Reissued 2018 by Routledge
2 Park Square, Milton Park, Abingdon, Oxon OX14 4RN
711 Third Avenue, New York, NY 10017, USA

Routledge is an imprint of the Taylor & Francis Group, an informa business

A Library of Congress record exists under LC control number: 99065450

ISBN 13: 978-1-138-60851-1 (hbk)
ISBN 13: 978-0-429-46060-9 (ebk)

Contents

PROLEGOMENA

BEING THERE: AUTOMATION AND INTERACTION DESIGN

Explains how and why the age-old question of function allocation is no longer adequate to describe the automation technology of today and tomorrow. Provides new sorts of guidance to human-automation design.

Presents the theoretical and practical challenges of designing feedback to make automation a teamplayer, given the realities of actual practice in aviation and related domains.

Explains from the manufacturer's standpoint how certification of flight management systems fails to catch both very basic and more subtle human-computer interaction flaws. Recommends changes to current practice.

Lays out what the regulators would like to see in terms of systems design and certification on the basis of human factors criteria. Takes into account changing JAR requirements in the near future.

Shows how an airline can re-organize its training of crews for cockpit resource management in automated flight decks.

Discusses the organizational and technical issues that surround the creation of an automation policy within an airline. Lays out a prototypical automation policy that strikes the balance between being too general (no specific guidance) and too flight deck specific.

1 Computers in the Cockpit: Practical Problems Cloaked as Progress

SIDNEY DEKKER AND ERIK HOLLNAGEL

Linköping University, Sweden

Introduction

Another book on aviation automation? Well, perhaps this is not a book on aviation automation *per se*. It is a book, rather, on how the entire aviation industry is coping with automation. Or more precisely, on how it is coping with the human consequences of automation, which it has fielded over the last two decades. The aviation domain, and the cockpit in particular, is frequently seen to be on the forefront of technological and human-machine interface developments. And indeed, in some sense, progress in the cockpit has been enormous. But from another angle, innovations presented as progress have brought along a large number of unanticipated practical problems - practical problems that today form the inherited by-products of once-vaunted automation technologies. Practical problems cloaked as progress, in other words.

Not only individual pilots have to learn how to operate automated aircraft. The entire aviation industry is learning how to deal with the profound implications that automation carries for the operation, design, regulation and certification of passenger aircraft. The industry is struggling to find ways to meaningfully educate and train operators for their new and different work in the automated environment. It is reconsidering who to select for these jobs and how. And now that current cockpit designs are firmly in place and its problems better-accounted for, it is regrouping to begin to regulate and certify cockpit equipment on the basis human factors criteria. This while manufacturers are voicing continued concern over the lack of concrete and specific ideas for better feedback design in the next generation of flightdeck automation.

One result of being ahead of the pack is that an industry encounters and, hopefully, solves a host of new problems and thereby generates an experience that can be helpful to others. It is therefore quite ironic that many other industries are in fact (re)-discovering similar automation related problems for themselves as they stumble ahead on technology-driven paths. For example, ship bridges are seeing more and more moded automation technology become responsible for navigation and many other on-board tasks. Standardised design of interfaces or system logic does not appear to exist and formal operator training is neither required nor well-organised. The result is that ships have begun to show the same pattern of human-machine breakdowns and automation surprises that were discovered in aviation years ago (see for example the grounding of the *Royal Majesty*, NTSB, 1996). Hence the need for this book: not only is it relevant to exchange experiences and swap lessons across one industry - aviation - it is also critical to show how one industry has to cope with the consequences of its own automation to industries that are poised to adopt similar systems in their operational environments.

Practical problems galore

To some extent research efforts and operational experience are beginning to pay off. In itself, this book is an outflow of the increasing realisation that automation is a mixed blessing. It reflects operational, educational and regulatory countermeasures that were for example inspired by the 1996 U.S. Federal Aviation Administration report on human-automation interfaces onboard modern airliners (FAA, 1996). Closer to the ground, many organisations that deal with complex automation acknowledge that changes in technology can be a double-edged sword. One defence procurement agency for example, says that they must strike a balance between simpler equipment and highly automated equipment. The reason they cite is that the former imposes greater manpower burdens but the latter can create excessive demands on operator skills and training. Such lessons learned indicate that old myths about automation (for instance that it reduces investments in human expertise) are becoming unstuck.

Nevertheless, almost all sectors of the aviation industry are still struggling in one way or another to adapt to the emerging realities of automation technology - to which this entire book is testimony. The training of pilots from the *ab initio* (zero-hour) level upward, for instance, has come to the fore as a key issue relative to automated flight decks (Nash, 1998; Lehman, 1998). Does requisite time in single piston aircraft of light wing loading have anything to do with becoming a jet transport pilot in a world of near sonic, satellite-guided computer-managed flight at 35,000 feet? These questions emerge during a time when European operators and regulators are attempting to harmonise training and licensing standards across the continent and while North-American operators are gradually losing a major source of pilots (the military), with collegiate aviation programs working to fill the gap (NRC, 1997). Questions about preparing pilots for their new supervisory roles do not stop at the *ab initio* level. The debate about optimal training strategies pervades the airline induction (multi-crew, operational procedures) and type-rating stages as well. A new pilot's first encounter with automation is often delayed to late in his or her training. This means it may fall together with the introduction to multi-crew and jet-transport flying, creating excessive learning demands. Telling pilots later on to be careful and not to fall into certain automation traps (a common ingredient in classroom teaching as well as computer-based training - CBT) does little to prevent them from falling into the traps anyway. The end result is that much of the real and exploratory learning about automation is pushed into line-flying.

Automation also erodes the traditional distinction between technical and non-technical skills. This tradition assumes that interactions with the machine can be separated from crew co-ordination. But in fact almost every automated mishap indicates that the two are fundamentally interrelated. Breakdowns occur at the intersection between crew co-ordination and automation operation. Crew resource management training is often thought to be one answer and is by now mandatory. It is also regulated to include some attention to automation. But all too often CRM is left as a non-technical afterthought on top of a parcel of technical skills that pilots are already supposed to have. Air carriers are coming to realise that such crew resource management training will never attain relevance or operational leverage.

Another issue that affects broad sections of the aviation industry is the certification of flight decks (and specifically flight management systems) on the basis of human factors criteria (Harris, 1997; Courteney, 1998). One question is whether we should certify the process (e.g. judging the extent and quality of human factors integration in the design and development process) or the end-product. Meanwhile, manufacturers are working to reconcile the growing demand for user-friendly or human-centred technologies with the real and serious constraints that operate on their design processes. For example, they need to design one platform for multiple cultures or operating environments. But at the same time they are restricted by economic pressures and other limited resource horizons (see e.g. Schwartz, 1998). Another issue concerns standardisation and the reduction of mode complexity onboard modern flight decks. Not all modes are used by all pilots or carriers. This is due to variations in operations and preferences. Still all these modes are available and can contribute to complexity and surprises for operators in certain situations (Woods & Sarter, 1998). One indication of the disarray in this area is that modes which achieve the same purpose have different names on different flight decks (Billings, 1997).

Air traffic control represents another large area in the aviation industry where new technology and automation are purported to help with a variety of human performance problems and efficiency bottlenecks (e.g. Cooper, 1994). But the development of new air traffic management infrastructures is often based on ill-explored assumptions about human performance. For example, a common thought is that human controllers perform best when left to manage only the exceptional situations that either computers or airspace users themselves cannot handle (RTCA, 1995). This notion directly contradicts earlier findings from supervisory control studies (e.g. the 1976 Hoogovens' experience) where far-away operators were pushed into profound dilemmas of when and how to intervene in an ongoing process.

Technology alone cannot solve the problems that technology created

In all of these fields and areas of the aviation system we are easily fooled. Traditional promises of technology continue to sound luring and seem to offer progress towards yet greater safety and efficiency. For example, enhanced ground proximity warning systems will all but eradicate the controlled flight into terrain accident. We become focused on local technological solutions for system-wide, intricate human-machine problems. It is often very tempting to apply a technological solution that targets only a single contributor in the latest highly complex accident. In fact, it is harder take co-ordinated directions that offer real progress in human-centred or task-centred automation than to go with the technological; the latest box in the cockpit that can putatively solve for once and for all the elusive problems of human reliability.

Many of our endeavours remain fundamentally technology centred. Ironically, even in dealing with the consequences of automation that we have already, we emphasise pushing the technological frontier. We frame the debate of how to cope with computers in the cockpit in the technical language of the day. For example, can we not introduce more PC-based instrument flight training to increase training effectiveness while reducing costs? Should we put Head-Up-Displays on all flight decks to improve pilot awareness in bad weather approaches? How can we effectively teach crew resource management skills through computer-based training tools? With every technical question asked (and putatively answered), a vast new realm of cognitive issues and problems is both created and left unexplored. The result, the final design, may beleaguer and surprise the end-user, the practitioner. In turn the practitioners' befuddlement and surprise will be unexpected and puzzling to us. Why did they not like the state-of-the-art technology we offered them? The circle of miscommunication between developer and user is complete.

One reason for this circle, for this lack of progress, is often seen to lie in the difficulties of technology transfer - that is, the transfer of research findings into usable or applicable ideas for system development and system improvement. This book is one attempt to help bridge this gap. It provides yet another forum that brings together industry and scientific research.

Investing in human expertise and automation development

The book echoes two intertwined themes. The first theme explores how and where we should invest in human expertise in order to cope with computers in the cockpit today and tomorrow. It examines how practitioners can deal with the current generation of automated systems, given that these are likely to stay in cockpits for decades to come. It examines how to prepare practitioners for their fundamentally new work of resource management, supervision, delegation and monitoring. For example, various chapters converge on what forms cockpit resource management training could take in an era where flying has become virtually equated with cockpit resource management (managing both human and automated resources to carry out a flight). There are chapters that target more specific phases in a pilot's training career, for instance the *ab initio* phase and the transition training phase. Yet another chapter makes recommendations on how an air carrier can proceduralise the use of automation in terms of how different levels of automation affect crewmember duties, without getting bogged down in details that are too prescriptive or too fleet-specific.

The second theme explores what investments we must make in the development of automated systems. The industry would like to steer the development of additional cockpit equipment and air traffic management systems in human-centred directions - but how? Current processes of development and manufacturing sometimes seem to fail to check for even the most basic human-computer interaction flaws in for example flight management systems coming off the line today. Two chapters examine whether and how certification and increased regulation could help by setting up standards to certify new or additional cockpit equipment on the basis of human factors criteria. Although these chapters represent current European state of the art in this respect, much work needs to be done and much more agreement needs to be reached, for example on whether to pursue quantitative measures of human error and human performance in system assessment. Another chapter lays out how we could finally break away from the traditional but unhelpful, dichotomous notion of function allocation in our debates about automation. As this automation is becoming more and more powerful, allocation of *a priori* decomposed functions misses the point. Also, such increasingly powerful automation needs to show feedback about its behaviour, not just its state or currently active mode - an issue targeted in a chapter on automation visualisation. Finally, one chapter looks into the problem of extracting empirical data about the future. As technology development goes ahead, in aviation and in many other fields of human endeavour, it becomes ever more important to be able to evaluate the human factors consequences of novel technological solutions before huge resources are committed to a particular system design. This chapter explains how to generate empirical data relating to human performance in systems that do not yet exist.

Real progress

As automation has brought along many practical problems under the banner of continued progress, the aviation industry is struggling to cope with the human-machine legacy of two decades of largely technology-driven automation. The lessons learned so far and the lessons still to be learned, carry information not only for aviation, but for a number of industries that are opening their doors to similar problems. Real progress across multiple industries, not the kind that cloaks the sequential introduction of practical problems into different worlds of practice, can only be achieved through acknowledging the similarity in challenges that we have created for ourselves. Hopefully this book offers some leads.

Acknowledgements

The initiative for *Coping with Computers in the Cockpit* was developed in the Swedish Centre for Human Factors in Aviation at the Linköping Institute of Technology with the encouragement and support from *Luftfartsinspektionen*, the Swedish Flight Safety Department. The editors owe particular gratitude to the Director of Flight Safety, Arne Axelsson, and the Flight Safety Department's technical monitors of civil aviation human factors work at Linköping University, Kaj Skärstrand and Bo Johansson.

2 Automation and its Impact on Human Cognition

SIDNEY DEKKER AND DAVID WOODS*

Centre for Human Factors in Aviation, Linköping Institute of Technology, Sweden
** The Ohio State University, USA*

Introduction

We introduce new technology because we think it helps people perform better. For example, we expect technology - and especially automation technology - to reduce people's workload, improve situation awareness, and decrease the opportunity for human error. Indeed, we value automation for its impact on human cognition. Through aiding the operator's awareness and decision making, new technology can increase system safety and improve the economy or accuracy of operations.

Looking a little closer, it becomes clear that we express the promises of automation technology almost always in quantitative terms. For example, *less* human workload will result if we replace a portion of the human's work with machine activity. And when we give the human less to do - when we shrink the bandwidth of human interference with system operations - we leave *fewer* opportunities for human error. Indeed, this is the traditional idea: that the replacement of human activity with machine activity has no larger consequences on the overall human-machine ensemble. The only thing that is affected is some kind of outcome measure. This outcome measure may be error count, or workload, or economy, and - indeed - all of them are quantifiable, and all of them somehow get better when we introduce automation technology.

Some of these quantitative effects have been realised, but only in a narrow empirical sense. For example, in highly automated systems there is less workload during certain times. There is also less opportunity - or no more opportunity - to make certain kinds of errors. No one is going to stick the wrong punch card into a flight management system: flight management systems don't work on punch cards (although the equivalent of downloading pre-created flight plans from an airline's operational base into an individual FMS does in fact exist).

But our pre-occupation with the quantifiable is distracting. The real and overall impact of technological innovations is qualitative - not quantitative. Humans and technology are not two separate and interchangeable components of a human-machine system. We cannot simply put in a little more of one and leave out a little more of the other. Instead, changes in one have fundamental consequences for the other. Changes in what we make the machine or the human do also have fundamental consequences for the interaction between them, for how they have to behave in relation to one another. The thought that humans and machines are substitutable or interchangeable without creating any larger impact on the human-machine system other than on some quantitative output measure, has turned out to be a myth. We call this the substitution myth (see also the chapter by Hollnagel in this volume).

Automation technology has had a profound impact on the way people in aviation and other systems do their work. And on what kind of work they do in the first place. Indeed, automation technology has fundamentally changed people's tasks, roles and responsibilities. For example, automation has lifted human responsibilities up into the realm of supervisory control. Here activities like delegation and monitoring are crucial. Human assessments and decisions increasingly have to be about the future. Such new work means new knowledge; new expertise and skill requirements. People who remain at work in automated systems have to know and be good at new and different things. Comments from practitioners across the aviation industry reflect these new realities. "I have never been so busy in my life, and someday this automation is going to bite me", says one pilot. "I diverted my attention from flying the aircraft to attend to the intricacies of reprogramming the computer", says another.

Automation technology has also created novel and unprecedented opportunities for human error and opened doors to new forms of system breakdown. For instance, automated airliners have electronic cocoons of protection wrapped around their failure-prone human pilots. These prevent pilots who are flying manually from going outside the normal flight envelope; from going outside the cocoon specified by the engineers. For example, pilots are not able to stall these aircraft, i.e. go beyond a critical angle of attack. But together - jointly - automation and pilots actually can take an aircraft beyond this angle and outside the envelope. Mishaps such as the crash at Nagoya (1994) show a pilot and autopilot fighting each other for control over the aircraft, which in this case lead to extreme pitch excursions far outside the flight envelope.

Particular patterns, persistent problems

The problems that surround practitioners in their interaction with new technology are more than a series of individual glitches. Of course, we can easily label the incidents to "human error" or attribute the handling difficulties to a "learning curve" which is inevitably associated with the introduction of a new generation of technology, whether aircraft (Benoist, 1998) or something else.

But based in part on a series of investigations of practitioner interaction with highly advanced aviation control environments (Sarter & Woods, 1992; 1994b; 1995; Dekker & Woods, in press), we can conclude that the difficulties that practitioners encounter indicate deeper patterns and phenomena. They all concern human-machine interaction. Aviation is not unique in experiencing these problems. Similar kinds of human-machine breakdowns occur in critical-care medicine (Moll van Charante *et al.*, 1993; Obradovich & Woods, 1996; Cook & Woods, 1996), the nuclear field and railways (Hollnagel, 1997) and the maritime world (NTSB, 1996). Various efforts have compiled these and related research results (Woods *et al.*, 1994; Billings, 1997; Abbott *et al.*, 1996). These works constitute markers in our progressive understanding of the effects that automation has on human performance and cognition. They have also begun to point to strategies that can help us deal with the perceived automation learning curve, or the perceived human error problem.

But problems and false promises persist. In this chapter we will first present two typical automation mishaps. They are from different domains but their anatomy has too much in common to ignore. Their similarity affirms that aviation holds no majority stake in automation problems. It also illustrates how the patterns of human-automation breakdown are of a certain kind - similar from one incident to the next, and from one domain to the next. After describing these incidents we will cover, in turn, what these problems and incidents mean for our investments in human expertise and then how knowledge of them should influence the development of additional or new automated equipment. We will try to map our current knowledge in both areas - setting the stage for the rest of the book which takes us further into what we have learned or still can learn from coping with automation.

Typical mishaps with automated systems

One June 10, 1995, a Panamanian passenger ship named *Royal Majesty* left St. Georges in Bermuda. On board were 1509 passengers and crewmembers who had Boston as destination - 677 miles away, of which more than 500 would be over open ocean. Innovations in technology have led to the use of advanced automated systems on modern maritime vessels. Shortly after departure, the ship's navigator set the ship's autopilot in the navigation (NAV) mode. In this mode, the autopilot automatically corrects for the effects of set and drift caused by the sea, wind and current in order to keep the vessel within a preset distance of its programmed track. Not long after departure, when the *Royal Majesty* dropped off the St. Georges harbour pilot, the navigator compared the position data displayed by the GPS (satellite-based) and the Loran (ground/radio-based) positioning systems. He found that the two sets of data indicated positions within about a mile of each other - the expected accuracy in that part of the world. From there on, the *Royal Majesty* followed its programmed track (336 degrees), as indicated on the automatic radar plotting aid. The navigator plotted hourly fixes on charts of the area using position data from the GPS. Loran was used only as a back-up system, and when checked early on, it revealed positions about 1 mile Southeast of the GPS position - nothing unusual.

About 34 hours after departure, the *Royal Majesty* ran aground near Nantucket Island. A quick check revealed that it was about 17 miles off course and that Nantucket Island was actually rather close by. The accident investigation found that the cable leading from the GPS receiver to its antenna had come loose and that the GPS unit (the sole source of navigation input to the autopilot) had defaulted to dead-reckoning (DR) mode about half an hour after departure. In DR mode, there was no more correction for drift. A northeasterly wind had blown the *Royal Majesty* further and further west.

After the fact, the grounding incident looks mysterious to outsiders who have complete knowledge of the actual state of affairs (Woods *et al.*, 1994). The benefit of hindsight allow reviewers to comment things such as:

- "How could they have missed X (the DR mode indication), it was *the* critical piece of information?"
- "Why didn't they double-check X and Y (GPS against Loran data), it could have avoided the mishap!"
- "Why didn't they understand that X (tripping the cable) would lead to Y (default to DR mode), given the inputs, past instructions and internal logic of the system?"

The wake of the *Royal Majesty's* grounding shows a maritime industry trying to make sense of the consequences of advanced integrated ship bridge systems on human performance. In this new operating environment, for example "the crew's failure to detect the ship's errant navigation for more than 34 hours raises serious concerns about the performance of the watch officers and the master... The watch officer is relegated to passively monitoring the status and performance of the automated systems. As a result of passive monitoring, the crewmembers of the *Royal Majesty* missed numerous opportunities to recognise that the GPS was transmitting in DR mode and that the ship had deviated from its intended track... As the grounding of the *Royal Majesty* shows, shipboard automated systems such as the integrated bridge system and the GPS, can have a profound influence on the watchstander's performance" (NTSB, 1997, pp. 30, 34 and 35).

Almost identical exclamations about human performance followed the 1995 crash of a Boeing 757, near Cali Colombia. After accepting a runway change, the crew programmed the automation to fly the aircraft to a beacon at the end of the new runway. Due to internal logic, however, the flight management computer interpreted the instruction as a different waypoint than the one the crew intended, and it made the autopilot commence a left turn to fly to that waypoint instead. The crew became very busy setting the aircraft up for the new arrival, while the aircraft strayed into the mountains off to the side of the valley in which Cali lies. Almost under total automatic control, it hit a mountain a few minutes after the wrongly interpreted computer instruction. From the position of a retrospective outsider, it is hard to understand how the Cali crew could miss so many critical cues during their gentle but quick progression towards disaster. Indeed, the official investigation blamed various controversial crew decisions for the crash. Aeronautica Civil determines that among the probable causes of this accident were:

- The flightcrew's failure to adequately plan and execute the approach to runway 19 at Cali and their inadequate use of automation;
- Failure of the flightcrew to discontinue the approach into Cali, despite numerous cues alerting them of the inadvisability of continuing the approach;
- Failure of the flightcrew to revert to basic radio navigation at the time when the flight management system-assisted navigation became confusing and demanded an excessive workload in a critical phase of flight (Aeronautica Civil, 1996).

Common reactions to failure

Cali and the *Royal Majesty* have much in common. One thing they share (and share with most other mishaps in high-technology human-machine systems) is the reactions they spawn. Reactions to failure can often indicate the extent of our misunderstandings about human performance in complex worlds, and our underestimation of how new technology has changed their operating environments and the work that has to go on within them.

When we go through some of the reactions ("the crew failed to monitor the vertical flight path...adequately plan and execute the approach...discontinue the approach", etc. or "none of the officers determined that the GPS had switched to DR mode ... their monitoring was deficient ...

they continued to miss opportunities to avoid the grounding"), we get the impression that the respective crews had motivational shortcomings. If only they had tried a little harder, then they would have picked up the data critical to their situation and integrated those in their assessments and decisions - so that the disaster could have been averted. Data critical to resolving the situation was available, so what is puzzling to the retrospective outsider is how these data could possibly have been missed. As the director of safety of an aircraft manufacturer put it: "You can incorporate all the human engineering you want in an aircraft. It's not going to work if the human does not want to read what is presented to him, and verify that he hasn't made an error" (quoted in Woods *et al.*, 1994).

Data availability

The common reactions to failure are predicated on a particular assumption about human performance. This is the assumption of data equi-availability - the idea that if in hindsight data can be shown to have been physically available to practitioners, it should have been obvious or picked up by them.

While it is easy to say that a critical piece of data should have been picked up, the feedback properties of the automated systems that bring these data forth often make it very difficult. In fact, automated systems have made it really hard for practitioners to pick up subtle changes in mode or status. According to one pilot, "unless you stare at it, changes can creep in". Ship bridge systems are generally no better. The GPS unit aboard the *Royal Majesty* was mounted across a chart table and only a two tiny letters (DR) notified the user of the currently active mode. The default to DR mode was automatic: it required no user concurrence. Future behaviour of the ship as governed by the DR mode (in terms of a projected track) was shown nowhere. In other words, no representation on the ship bridge showed in one picture where the automation was taking the ship, even though the automation was in charge of the ship's heading for a day and a half.

These properties of automation have created an interesting situation, described by Nadine Sarter in her doctoral work (see Sarter, 1994). Automated systems have become *stronger* (they can carry out long sequences of action without the user providing any inputs): they have increased authority and autonomy ("Here, you take the boat from Bermuda to Boston"). But the same system remains *silent* about what or how well it is doing, and it is difficult to direct ("How do I get it out of this mode?"). The paradox is that the strength of automation increases co-ordination demands, while the properties of the automation's feedback and interface (showing a tiny "DR"; not asking for user concurrence; not annunciating the mode change in any other way) make such co-ordination extremely difficult.

The reason why co-ordination demands go up is that today's powerful automation is no longer a subsystem that can easily be switched on or off. Operators no longer treat automation as a separate component in the larger operational system. Instead they approach automation as an animate partner in systems operations. "What's it doing now? Why is it doing this?" The kind of automation that steers a large ship towards Boston, or directs an aircraft across the Pacific, no longer just *is*. It *behaves*. "How did it get into this mode? How do I stop it from doing this?" These questions indicate how automation has become an agent, capable of pursuing its own goals (getting this boat from here to Boston) by using knowledge about itself and about the world.

Introducing a new agent has enormous effects on human cognition. An agent is in some sense a new team member. This team member can do certain things on its own, but has to be informed and supervised at the same time. It must also communicate about its work and progress. This means that automation imposes co-ordination demands: a new team member means more co-ordination - precisely the kind of co-ordination that today's automated systems are rather bad at (given that they are silent and difficult to direct (Sarter, 1994)).

The dissociation between availability and observability

The silent strength of automated systems produces a dissociation between data availability and data observability. There is a large difference between data that can be shown to have been available in hindsight ("DR" was available, how could they have missed it?"), and data that was actually observed, used and integrated by the crew given their ongoing tasks and attention demands. Observability is a technical term. It refers to the cognitive work that users need to do to extract meaning from available data. Observability refers to processes, the cognitive work, involved in extracting useful information. It results from the interplay between a human user knowing when to look for what information at what point in time and a system that structures data to support attention guidance. How much cognitive work does the user need to do to make sense out of "DR"? How much guidance does "DR" provide? It is a large step from this unspecific, tiny annunciation to the understanding that your boat isn't heading for Boston after all. To understand what "DR" means, not only in general but also in this particular situation, can take significant mental resources. If the mode annunciation is seen in the first place, knowledge has to be called from memory (what was this "DR" again?) and translated into the expected behaviour of the ship. The knowledge of where the boat is headed is not in the world, not observable. It is up to the head of the user to figure out or remember what DR means and what it does to the boat. This is in a sense an example of Don Norman's "conspiracy against memory". "DR" might have been available, but technically it was not observable.

The dissociation between data availability and data observability is especially noticeable in highly dynamic situations. In these cases, novel combinations of factors can push incident evolutions beyond the routine. Practitioners themselves have little control over the pace of process activities and evidence about unfolding conditions gets generated over time. Gaps and uncertainties are common and practitioners have to make assessments and decisions on the basis of partial and ambiguous evidence. The Cali sequence of events represents precisely such an event-driven, busy period. Saying that something should have been obvious in this situation reveals our own ignorance of the demands and activities of people in complex, dynamic domains, where they must juggle multiple interleaving tasks and sift through uncertain and changing evidence. The fact that certain data was available somewhere in the world during these times does not mean that it was relevant to the multiple tasks at hand, that it was expected, that it was understandable, or that it was in a location or format that made it compelling to look at in the first place.

The critical test for level of observability is when annunciations help practitioners notice what they did not expect to see. Or when it helps them notice more than what they were specifically looking for. If a display only shows users what they expect to see or ask for, then it is merely making data available. The measure of true support comes when the representation helps users see or find what they were not explicitly looking for. Increasing machines' autonomy, authority and complexity creates the need for increased observability. The automation characteristics require new forms of feedback, emphasising an integrated dynamic picture of the current situation, automation activities, and how these may evolve in the future. Striving for these larger representational goals also helps designers achieve a balance between underinforming people about automation activities and overwhelming them with details about every minor action.

Common patterns of breakdown

Cali and the *Royal Majesty* represent a common pattern of breakdown between humans and automation. The incident signature is that perfectly functioning machines are flown or sailed into the ground. There is nothing mechanically wrong with these systems. But through a series of persistent and deepening misassessments and miscommunications between human operators and the automation, they evolve towards failure. This kind of accident sequence can be called the "going sour" accident (see Cook, Woods & McDonald, 1991). In this progression, a small event (e.g. an uncommanded turn to the left; a tripped GPS cable) triggers a situation from which it is in principle possible to recover. But through a series of commissions and omissions, misassessments and miscommunications, the human-automation ensemble gradually manages the situation into a serious and risky incident or even accident. In effect, the situation is managed into hazard.

Accidents such as Cali can hardly be classified as controlled flight into terrain (CFIT) anymore. Instead, "typed flight into terrain", or "managed flight into terrain" would better reflect the critical human-machine interactions preceding the mishap. The automated system is handling the aircraft, which has relegated the crew to a supervisory and directory role rather than a controlling one. As the *Royal Majesty* shows, aviation does not hold a patent on the going sour sequence. The gradually managed "radar-assisted collisions" of the seventies (see Perrow, 1984 for some excellent examples) have now made room for "programmed groundings" such as the one near Nantucket.

The going sour progression is consistent with research findings on complex system failure. Very small - even trivial - events (such as a GPS antenna cable kicked loose) can start a progression towards breakdown (see Perrow, 1984). Many other factors, individually insufficient, are jointly necessary to push a system closer to the edge. Simultaneously, system defences need to breached (e.g. the erosion of double-checking Loran and GPS position data) to allow full-scale system breakdown (Reason, 1990). Having authority and autonomy over safety-critical tasks, highly automated systems can sponsor in their own ways this pattern of complex system failure. The going sour scenario is an important accident category which represents a significant portion of the residual risks in aviation.

Much of the management toward breakdown has to do with the fact that automation often does not help during busy periods. In fact, it gets in the way. When there is already a lot to do, automation will give the user even more to do. It will ask for inputs, it may spring surprises. This problem occurs because of a fundamental relationship. The greater the trouble in the underlying system or the higher the tempo of operations, the greater the information processing activities required to cope with the trouble or pace of activities. The more unusual the situation, the higher the tempo of operations, the more demands there will be for monitoring, for attentional control, for information and for communication among team members. This includes human-machine communication and co-ordination. The upshot is that the burden of interacting with the automated system tends to be concentrated at the very times when the practitioner can least afford new tasks, new memory demands or attentional diversions.

Clumsy automation is a label coined by Wiener (1989) to describe such poor co-ordination between the human and machine. The benefits of new technology accrue during workload troughs: when there was already virtually nothing to do, technology will give the user even less to do. But the costs or burdens imposed by the technology (the additional tasks, new knowledge, forcing the user to adopt new cognitive strategies, new communication bur-