HPSAA II
Volume II

HUMAN PERFORMANCE, SITUATION AWARENESS AND AUTOMATION

Edited by
Dennis A. Vincenzi • Mustapha Mouloua
Peter A. Hancock

HUMAN PERFORMANCE, SITUATION AWARENESS AND AUTOMATION:
CURRENT RESEARCH AND TRENDS

HPSAA II
Volume I

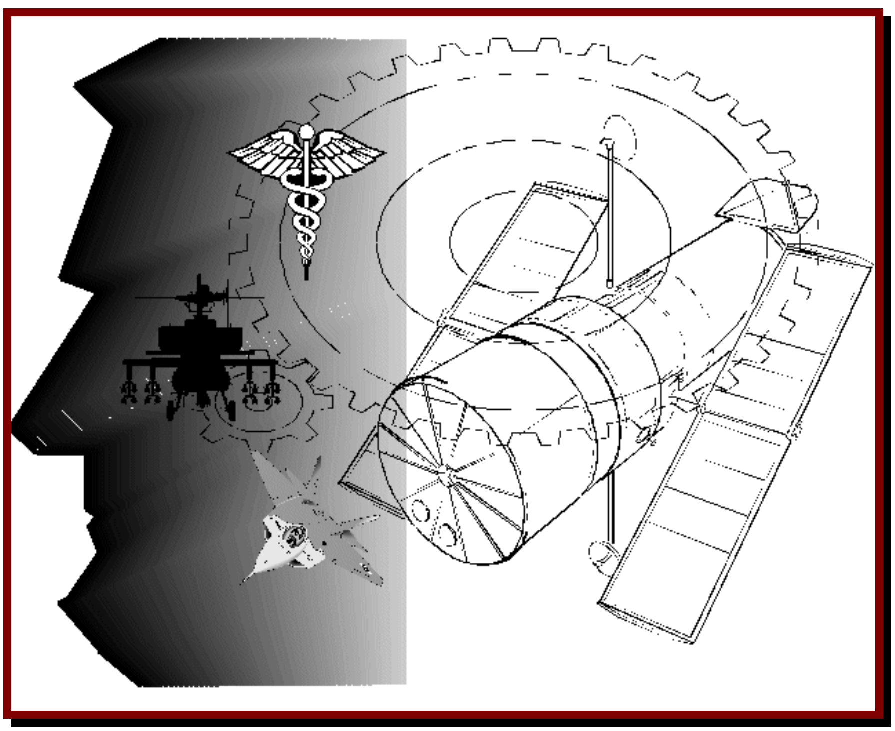

Edited by:

Dennis A. Vincenzi
Mustapha Mouloua
Peter A. Hancock

Proceedings of the Second Human Performance, Situation Awareness and Automation Conference (HPSAA II), held in Daytona Beach, FL, March 22 – 25, 2004.

Conference Chairs:

Scientific Program Committee:

COPYRIGHT

Technical Editing and Graphic Design by:

Dennis A. Vincenzi
Ryan Wasson

Department of Human Factors and Systems
Embry-Riddle Aeronautical University
600 S. Clyde Morris Blvd.
Daytona Beach, FL 32114

Human Performance, Situation Awareness and Automation: Current Research and Trends

Edited by:

Dennis A. Vincenzi, Ph.D.
Department of Human Factors and Systems
Embry-Riddle Aeronautical University
600 S. Clyde Morris Blvd.
Daytona Beach, FL 32114
dennis.vincenzi@erau.edu

Mustapha Mouloua, Ph.D.
Department of Psychology
University of Central Florida
4000 Central Florida Blvd
Orlando, FL 32816
mouloua@pegasus.cc.ucf.edu

Peter A. Hancock, Ph.D.
Department of Psychology
University of Central Florida
4000 Central Florida Blvd
Orlando, FL 32816
phancock@pegasus.cc.ucf.edu

Lawrence Erlbaum Associates
Mahwah, NJ

FOREWORD

In the era which followed the second world-war, human-machine interaction was dominated by the questions of movement control and movement accuracy as luminaries such as Craik, Crossman, McRuer, Baron and their respective colleagues fought to generate analog solutions to problems fundamentally framed as questions about analog machines. As pointed out in his admirable keynote paper (Moray, this volume), the challenge was often understood as a complex of demands but even within such multi-tasking environments, the continuous control problem was the one that thrust itself predominantly to the fore. Significant and successful efforts to address such challenges were indeed forthcoming just as the necessity of this aspect of human-machine interaction began to wane. Not merely because of the digital revolution, but also the very fact that these continuous control problems themselves were open to successful resolution left them vulnerable to the engineering community who rightfully sought to solve soluble problems.

As with all technical leaps, in solving a range of problems, such successes create a vista of others. In our science, this transition was observed in the content composition of scientific conferences. The once vibrant 'Annual Manual' Conference dissolved, not so much as with a bang (bang) as with a whimper. There being no fundamental vacuum in human endeavors, the translated problem was resurrected through the wonderful insights of individuals such as Sheridan and Moray who understood that the issue of control had been removed one level to a supervisory one, in which the actions of the operator were much more sporadic and intermittent with all the issues of dissociation, memory capacity, mental workload, and vigilance that this remove engendered. Of course, there had previously been rudimentary process control systems but now this form of interaction was thrust to the fore. For almost a decade, such work was the topic of workshops and sporadic meetings as the associated issues resolved themselves and the technical systems of this form were designed, fabricated, operated, and decimated.

It was a ground swell of understanding, largely championed by Charles Billings that found formalized expression in the first meeting in the present series, conceived and convened by Raja Parasuraman and Mustapha Mouloua. In 1994 they hosted the inaugural Conference in Rockville, Maryland with the paradoxical title, 'Human Performance in Automated Systems.' While the general perception of automation involves machine response with no human intervention, that Conference rapidly established that the human involvement, although further removed from the site of activity, nevertheless continued to exert a crucial influence over operations, even if this influence was only evident at the design stage. A strong motivation for holding this first meeting was the growing realization that automation was pervasive in virtually all walks of life, a trend which continues today and still acts as a stimulus for the series of meetings that have taken place subsequently.

From Cocoa Beach, Florida in 1996 through Norfolk, Virginia in 1998 we have seen research presentations on human interaction with automated systems in domains as diverse as transportation, medicine, manufacturing engineering, aerospace operations consumer products, and now consumer software. In 2000, the Conference on Automation joined forces with a partner group on situation awareness (SA) with a highly successful first combined meeting in Savannah, Georgia. The situation awareness group itself had a history of successful meetings, again in both Georgia and Florida. In keeping with the increasing emphasis on a move from immediate manual control to remote, periodic interaction, the ascending complexity of systems demands that one be aware of a large range of environmental and task-based stimulation in order to match what is done with what has to be done. Thus SA and automation-based interaction fall naturally together and the present Conference is the second incarnation of this most fruitful union.

The efforts evident in the present conference are captured by a number of themes. Moray in his keynote implores us not to forget the hard won knowledge of the past, protesting that the success of quantitative modeling of skill-based behavior can well form a template for successful quantitative modeling of both rule-based and knowledge-based behavior. His admonitions should receive deserved attention. There is also a small but growing recognition that affective facets of performance and detailed individual differences in operator's attitudes and capacities are issues with which we must now grapple. Primarily, we still seek unified models, theories, taxonomies and descriptions to characterize what is a dispersed domain of interaction but one in which we continue to seek and explicate core issues. There is indeed still much for us to accomplish, although where this will occur under the same banner title, or whether there will again be a technology-driven and societally-driven metamorphosis, only the future will tell. As we move inexorably further into the 21st Century, further diversification of the applications of automation will evidently continue, the revolution in genetic technology being an obvious coming example. Given the pervasive nature of this form of human-machine interaction, it is vital that we apply the lessons of the past to map a future for the symbiotic relationship between humans and the artifacts we create. It is as part of this on-going endeavor that the present volume is offered.

Peter A. Hancock, Ph.D., *University of Central Florida*
Dennis A. Vincenzi, Ph.D., *Embry-Riddle Aeronautical University*
Mustapha Mouloua, Ph.D., *University of Central Florida*

ACKNOWLEDGEMENTS

As with any successful effort, there is never one single individual, organization or agency that is totally responsible for that success. Such is the case with this conference. Many individuals contributed countless hours toward the organization and successful execution of this conference, and many organizations provided critical funding to support this conference. It is with great pleasure that I attempt to recognize their efforts and express heartfelt thanks for all their hard work.

I would like to thank the following persons for devoting their time and effort to make important contributions to the organization of the Human Performance, Situation Awareness and Automation Conference II (HPSAA II):

- Susan Vincenzi
- Ryan Wasson

Susan Vincenzi planned the conference food menus and assisted in the planning and organization of many logistical details of the conference ranging from delivery of equipment to listing and purchasing of necessary supplies for the conference. Ryan Wasson was largely responsible for the technical editing of all papers submitted for publication in the proceedings. He contributed many hours of tireless effort to make the papers and proceedings as uniform and consistent as possible. His assistance prior to and during the conference was much needed and greatly appreciated.

I would like to thank the scientific committee for their comments and suggestions during the process of organizing the meeting. Their help and guidance in this regard was invaluable. The scientific committee included Mustapha Mouloua, Jean Bresson, John Deaton, Mica Endsley, Peter Hancock, David Kaber, Anthony Majoros, Regis Mollard, Raja Parasuraman, Mark Scerbo and John Wise.

I would like to offer special thanks to the conference co-chair, Mustapha Mouloua, for his unending enthusiasm and contributions toward organization and marketing of the conference.

Finally, I would like to thank all our sponsors, without whose generous support, this conference would not have been possible. The support provided by the different sponsors ranged from manpower and supplies to major monetary contributions that funded various operational aspects of the conference that included but were not limited to funding for guest speakers, travel expenses, student support, printing of programs, and printing of proceedings to provide a permanent record for dissemination of information throughout the scientific community. The conference sponsors are:

- The National Aeronautics and Space Administration (NASA) Langley Research Center (LaRC)
- The Office of Naval Research (ONR)
- The U.S. Army Medical Research Acquisition Activity (USAMRAA)
- The Defense Advanced Research Projects Agency (DARPA)
- CHI Systems, Inc.
- Human Solutions, Inc.
- Interface Analysis Associates
- Embry-Riddle Aeronautical University
- The University of Central Florida
- Florida Institute of Technology

My sincerest thanks go out to all these individuals and organizations for their support in making this conference a great success.

Dennis A. Vincenzi, Ph.D.
Conference Chair
Embry-Riddle Aeronautical University

TABLE OF CONTENTS

CONFERENCE KEYNOTE ADDRESS

TECHNICAL PAPERS

SITUATION AWARENESS

STRESS, WORKLOAD AND FATIGUE

DECISION MAKING

DRIVER PERFORMANCE AND DISTRACTION

CONFERENCE KEYNOTE ADDRESS

Où sont les neiges d’antan?[1]

Written by:

Neville Moray
Magagnosc, France
moray2@wanadoo.fr

Presented by:

Thomas B. Sheridan
Massachusetts Institute of Technology

[1] “Where are the snows of yesteryear?”. François Villon. 1450.

ABSTRACT

Starting from an examination of two modern papers, about situation awareness and about allocation of function in human-machine systems, this paper examines the requirements for a quantitative predictive model of rule-based behaviour. It shows that much relevant empirical evidence exists in papers of 30 to 40 years ago, and that there has been convergence both of theory and data. It offers tentative guidelines towards the development of a predictive model.

INTRODUCTION: SITUATION AWARENESS AND AUTOMATION

It is well-known that elderly people are incapable of transferring new material from working to long term memory, but are adept at recalling material from the past. In the spirit of that observation I thought that I would take the one or two (fairly) recent papers, and use the remote past to illuminate the work that respected scientists and researchers are currently doing. The two modern papers that I would like to examine are Endsley, (1995), and Parasuraman, Sheridan, and Wickens, (2000).

I want to examine these papers because the results of a recent straw poll of over 100 people, inquiring as to which papers they thought were historically the most important human/factors ergonomics papers, surprised me, and made me wonder what had happened to HF models. Top of the poll was Bainbridge (1983), and equal second were Miller (1956), and Rasmussen, (1983). It is curious that there were no candidate papers from "classical" industrial ergonomics (even when deliberately requested) and very few votes for classical engineering candidate papers.

The lack of engineering papers is particularly surprising. People seem to have been thinking mainly of advanced human-machine systems when asked about their opinions of HF/ergonomics (I take it that the words mean the same thing), but there was almost no support for papers on control theory. Yet control theory is the one area where there has been quite exceptional success in theory based practical applications. It provides almost the only area in which we have developed not just good descriptive models, not just an ability to perform strong analysis of system performance, but also are able to predict, even in advance of prototypes, whether a machine will be usable by operators. Why have people forgotten its achievements so soon?

I think perhaps the vote suggests the following:

1. A shift from skill-based behaviour (SBB) to automation, with an increasing emphasis on rule-based-behaviour (RBB) and knowledge-based behaviour (KBB), (Rasmussen, 1983.).
2. A mature recognition of the importance of limits due to inherent psychological mechanisms such as memory, channel capacity, etc. (the Miller paper). (In modern work this is called "bounded rationality".)
3. Increasing interest in automation and the impact of system complexity.

The two papers which I have taken as the starting point of this talk reflect such conclusions. Perhaps control theory has lost its influence because it was traditionally concerned with SBB, and with the increasing interest in complex continuous processes with extensive automation on the one hand, and with automated and robotic discrete manufacturing on the other, there is less interest in SBB. As Tom Sheridan once remarked, it is ironic that just as we really learned how to model manual control it began to disappear. Even aviation seldom requires SBB: for an extreme example see the recent paper in Ergonomics in Design (Cummings 2004.), and of course the Airbus series of aircraft. Even so, it is remarkable that the influence of classical and optimal control theory seems to have disappeared from human factors and ergonomics literature. Why, for example, did the Annual Manual meetings cease? Is this meeting their successor? If so, where are the computational and *predictive* models? And, (what I find very worrying), why are there so few references in our literature to engineering journals, such as *Automatica, IEEE Transactions on Systems Science and Cybernetics,* and other places where engineers are describing their approaches to human-machine system design and modeling?

IN SEARCH OF PREDICTIVE MODELS

One reason for the success of control theory was the tight coupling between human and machine in manual control systems. The highly practised human became a component (even if adaptive one which could, on occasion, switch into RBB or KBB). The result was a highly quantitative and predictive set of models (McRuer and Krendel, 1959; McRuer and Jex,1967; Young, 1973). The least developed aspect was adaptive control, although there were many good papers , (e.g. Young, L. 1969; Crossman & Cooke, 1974).and a particularly fine book by Kelley (1968) that pretty well sank without trace and is almost never referenced in modern papers. If we are now all so interested in advanced, automated human-machine systems, why have has no equivalent predictive models appeared for RBB and KBB?

In what follows, when I speak of "predictive" models I am looking for more than ordinal prediction. Models should say more than, "If you make it like this it will be better than if you make it like that."; rather, one would like them to say, "This is what will happen in the next five seconds." I think that the intense interest in "situation awareness" (SA), is a reflection of a desire for such models.

Table 1 lists some of the reasons we have not seen any predictive models arise to supplement control theory.

TABLE 1.
WHY HAVE NO NEW PREDICTIVE MODELS FOLLOWED CONTROL THEORY?

- The loss of contact with engineering has led people to avoid models tightly coupled to the dynamics of the systems and environments.
- A zeitgeist developed in the 1980s that rejected mathematical models in favour of a deeper psychological understanding of the operators' cognitive processes. But almost no predictive models have resulted.
- In modern systems the coupling between human and machine is much looser. (It was loose even for very slow manual control processes and was never developed for discrete manufacturing. See papers in Edwards and Lees (1974) for discrete intervention in the case of slow process control.)
- The coupling changed from physical causality to logical and semantic relationships (Rasmussen, 1986).
- The operators' role changed from controller to supervisor and intervener.
- Operators' actions afforded by the systems became ever less analog, and less depended on physical strength and skill. (Interfaces intervened between operator actions and provided digital control of the systems.)
- As human-machine systems (HMS) become ever more complex and larger, teams, rather than individuals control them, and modeling team behaviour seems at first sight very different from modeling individual behaviour.

If we are to develop new predictive models, two things are needed: a return to a detailed interest in models that capture system dynamics, and a way to describe the coupling between operators and the new machines. Reading recent papers suggests that implicitly people recognize the need to improve coupling. Table 2 lists some work that supports this contention.

TABLE 2.
EVIDENCE OF INTEREST IN IMPROVING HUMAN MACHINE COUPLING

1. The emphasis on Situation Awareness.
2. The discussion of the design of automation by Parasuraman, Sheridan and Wickens (2000).
3. Klein's work on Recognition Primed Decisions
4. Wood's notion of "visual momentum".
5. The ideas of Rasmussen and Vicente about "ecological" displays and controls.

It is clear that a predictive model at the level of RBB and KBB will not be a purely quantitative mathematical model in the way that control theory was for SBB. In that respect Bainbridge was right to challenge control theory as a general model for process control (Bainbridge, 1981; Montmollin & De Keyser, 1985.). But while a mathematical model may not be possible, a *computational* model that includes quantitative components is surely a desirable goal, and the need for a mixed model is evident if we look closely at Endsley's account of SA.

In fact, there are several models that have tried to extend the exact modelling of SBB that was achieved by control theory to RBB and KBB, even though not necessarily while thinking in those terms:

1. SOAR
2. ACT*-R
3. PROCRU

and there are also some specialized models such as those of Hollnagel and Cacciabue (Cacciabue, Cojazzi, Hollnagel, & Mancini, 1992; Cacciabue, P.C., DeCortis, F., Drozdowicz, B., Masson, M., & Nordvik, J-P. 1992; Hollnagel, 2003);Roth, Woods & Pople, 1992; and Baron. Fehrer, Pew, and Horwitz, 1986.

The first two have been applied to modeling "real" behaviour to some extent, although SOAR and ACT*-R have seldom been used to model real time activities particularly where satisficing under time pressure rather than optimal decision making is required. (I use "ACT*-R" to cover all the recent versions.) It is particularly to be regretted that PROCRU was not further developed, because it included optimal control theory for SBB, and was reaching for the ability that SOAR and ACT*-R have to deal with semantics, logical reasoning, and other aspects of the psychological mechanisms that support complex psychological operations.

A look at Endsley's account (Endsley, 1995) of situation awareness emphasises the need for complex computational models. It is ostensibly a paper about SA, but could equally well have been published as an introductory chapter in a text on cognitive psychology: it is about *everything*, perception, attention, working memory, long-term memory, decision making, and so on. This is not really surprising, because SA is not really the name of a particular psychological function. It is a shorthand description for, "keeping track of what is going on around you in a complex, dynamic environment." Obviously, one needs to use all one's cognitive (and for that matter motivational and emotional) abilities for such a task. Equally, the aim of efficient SA is to keep the operator tightly coupled to the dynamics of the environment.

Often one thinks of automation as loosening the coupling between operators and the task, so as to lighten their loads, but it is clear that Parasuraman et al. (2000) do not think that. Indeed all the best researchers on automation, including Parasuraman et al., Woods, Hollnagel, Rasmussen, Vicente, Bainbridge, etc., point out that the major problem with automation is the desire to reduce the load on operators coupled with its inability to keep them tightly in the loop. In the classical description of Supervisory Control, (Sheridan, 1978), the flow chart of closed loop supervisory control makes this quite clear (Figure 1).

The main thrust of Parasuraman et al. (2000) is to make this coupling strong but flexible. In the supervisory control part of Figure 1 the block called "human operator transfer function" conceals a great deal. It is clear in earlier discussions of adaptive control (such as Young ,1969; and Kelley, 1968), that the block actually contains what amounts to Endsley's model of dynamic cognition, because the operator has to assess the situation, make plans, decide how to act, and if necessary intervene, as Sheridan (1978) described. Parasuraman et al.'s analysis of automation in essence disassembles the flow chart of cognition into sub-functions and proposes that automation be allocated to these different sub-functions as required to optimize the coupling of human operators to the task. On some occasions it may be most effective to substitute machine sensing and perception for that of humans, on other occasions to substitute machine planning or decisions making, and on others control or action. (And of course there is always the possibility of automating two or more sub-functions simultaneously.)

One reason to try to develop a computational predictive model is that there is a large measure of agreement about what should be in such a model, and how to link it to decisions about automation. In addition to the well-known block diagrams of Endsley and of Wickens, there are the descriptions of human machine interaction in Roth, Woods & Pople (1992), Rasmussen, (1984, 1995), and Bainbridge (1991), and less well known ones such as Hess (1987), Hollnagel, (2003), and Cacciabue et al.(1992). There is even considerable agreement between those developed by the engineering human factors community and SOAR or ACT*-R. Here, for example, is a quotation from Hess (1987):

> "(The internal model is) a volatile internal spatial/temporal representation of the environment which the human is assumed to possess and use while interacting with complex dynamic systems."

Hess goes on to suggest that in a multi-dimensional world the operator chooses a frame within which to stay, using time-driven planning, and that time scales change as O moves up or down the hierarchy of representation. It all sounds very like a synthesis of what Endsley would later describe and what Rasmussen had already discussed.

Surely the time has come to try for an overall computational synthesis. Roth et al. (1992) were well on the way to such a model, as indeed were the developers of PROCRU (Baron, Muralidharan, Lancrafty, & Zacharias, 1980; Baron et al., 1986). Sanderson and I, in the early 1990s, got as far as writing pseudo-code for Rasmussen's ladder, although we never completed it (Sanderson, 1991). (See Figure 2.)

Looking at that pseudo-code today, it appears to contain many productions rather like a first pass at a model of SA. (See particularly Productions 2, 3, and 4). Productions 6, 13, and 14 on the other hand might be relevant to deciding what part of the processing might be allocated to automation.

Let us look a bit more closely at some ways that a computational model of SA and function allocation might be tied together, and implemented as a computational model.

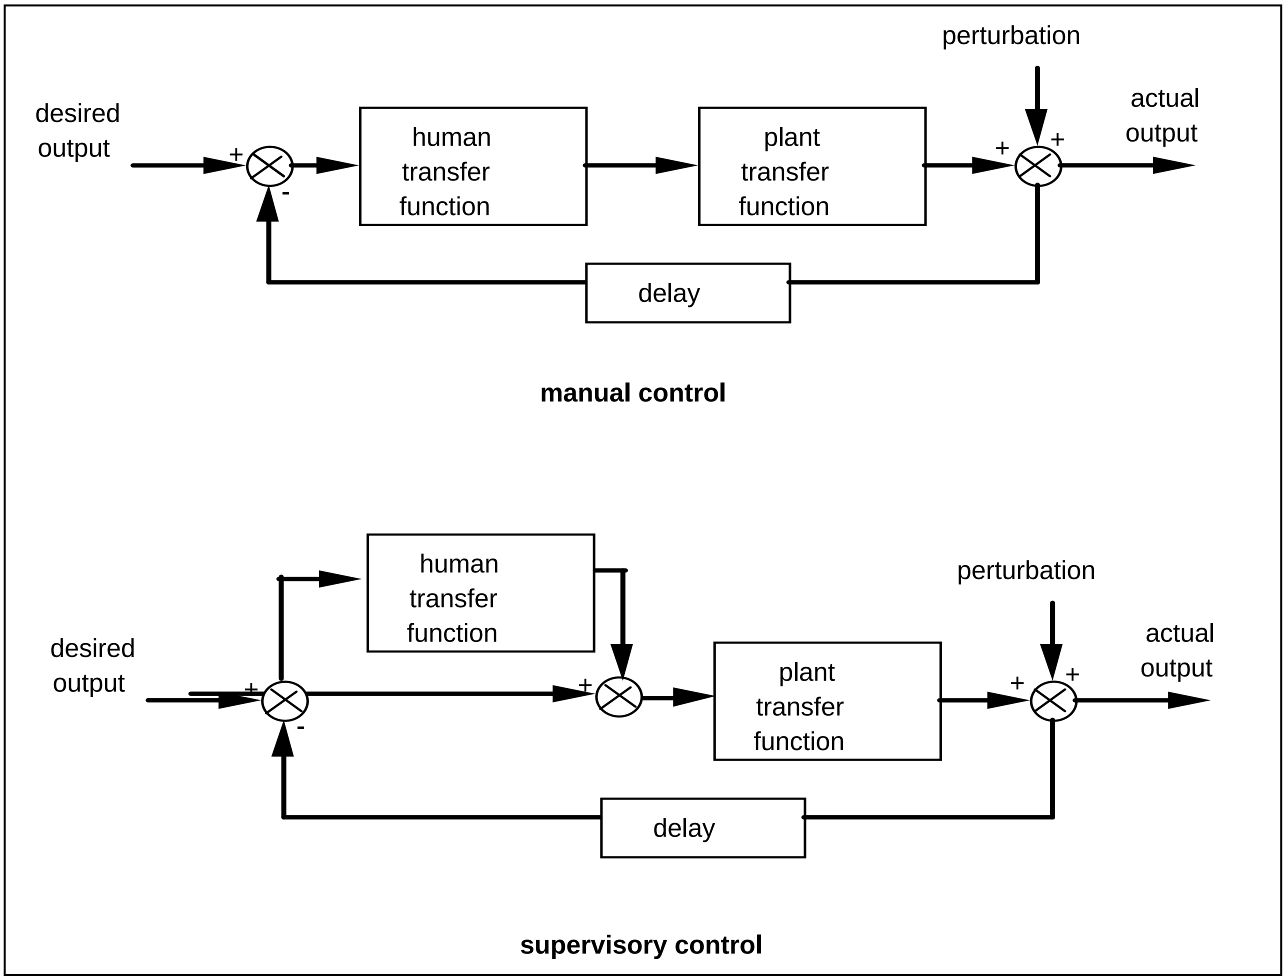

Figure 1. Manual and supervisory control

MODEL CONVERGENCE IN THE LITERATURE

Suppose that we want to decide whether, for some task, it is reasonable to expect operators to monitor the dynamic environment, or whether that task should be allocated to automation. Let us start by looking at Endsley's definition of SA. According to her it is

> "the perception of the elements in the environment within a certain volume of time and space, the comprehension of their meaning, and the projection of their status into the near future.".
>
> (Endsley, 1995.)

1. if state = Alert then ...
   p(response$_i$) = f(alert type, environmental invariance, practice, response mapping, etc.)
   ... then carry out response$_i$

2. if state = Alert and response not carried out
   then set goal = Observe information and data

3. if goal = Observe information and data then ...
   data = f(data available, scanning habits, current scheduling policy, hypotheses, incoming events, time to observe, etc.)
   ... then Data = {d_1, d_2, ... }

4. if goal = Identify state then ...
   state = f(observability, data observed, diagnosticity of data, operator knowledge, familiarity, likelihood, time, etc.)
   ... then State = s_i

5. if goal = Determine state consequences and State = s_i then ...
   state consequences = f(system mental model, performance goal, etc.)
   ... then State consequences = {sc_1, sc_2, ... , uncertain}

6. if goal = Evaluate performance criteria then ...
   performance goal = f(organizational knowledge, etc.)
   ... then Performance goal = {g_1, g_2, ... }

.
.
.

13. if goal = Choose criterion and Criterion = c_i
    then set goal = Define policy: operationalize criterion

14. if Criterion = c_i and goal = Define policy: operationalize criterion then ...
    policy = f(criterion c_i, memory of criterion-policy mappings, momentary system configuration, policy feasibility, display/technological constraints, time available, etc.)
    ...then Policy = p_i.

Figure 2. Pseudo-code for Rasmussen's Decision Ladder.

Why is this difficult for people? And why is it hard to model RBB and KBB in "a certain volume of time and space"?

It is not just the "volume" of time and space that is the problem. It is the rate at which events occur (constraints caused by limited time and the dynamics of the environment) and the quantity of information to be processed (the complexity of space, or rather, of things that exist in space). I am reminded of Simon's famous remark that it is the complexity of the environment rather than the complexity of the mind wherein complexity lies.

Effective SA requires operators to simplify the environment so that observations can be handled by limited memory, and at a rate within the capacity of limited attention. There is a considerable amount of empirical and theoretical work on how people simplify the complexity of their environments, and by so doing gain time to handle the rate of events. If we can understand how people do this, we will be in a better position to decide what aspect of their task to automate, what version of Parasuraman et al.'s (2003) array of automation to implement. We might begin by considering the spectrum of task times described by Sanderson and Fisher (1994). (See Figure 3.)

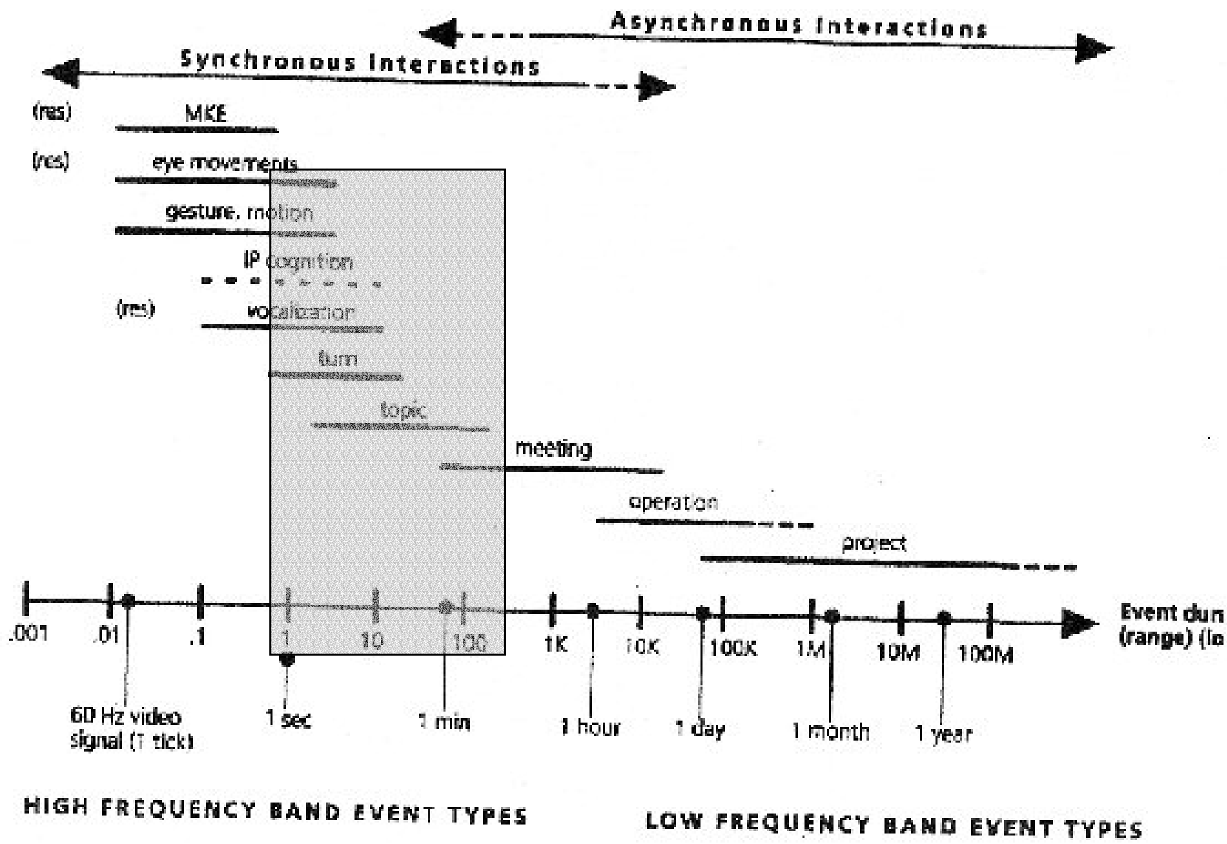

Figure 3: A spectrum of durations of event classes (Sanderson and Fisher, 1994).

I have overlain Sanderson's spectrum with a window that marks the region where I would expect people to have little difficulty in handling information in real time. If events occur faster than this, then operators will need assistance to update their SA, and if events occur very much more slowly, they will have difficulty in integrating the information to obtain an overall picture due to loss of information from working memory. Of course if the environment which they are observing is very complex, (high dimensional), even a low event rate for individual parts of it may cause too high an event rate to allow O to maintain SA. It is important to consider very large and slow process such as process control as well as high bandwidth systems such as high performance aircraft. Both ends of the spectrum are important and cause problems for the human operator.

Consider some more of the characteristics of SA as described by Endsley. SA is a dynamic description of the status of the environment and the human-machine system during a functionally important window.

> "There is evidence that an integrated picture of the current situation may be matched to prototypical situations in memory… (a theory of SA) should explain dynamic goal selection, attention to appropriate critical cues, expectancies, regarding future states of the situation, and the tie between situation awareness and typical actions."
>
> (Endsley, 1995.)

Wherever we look at Endsley's description of SA, we find that workers in other contexts have said similar things. That is most encouraging. The next three figures, 4, 5, and 6 are Figures 1, 2, and 4 from Endsley's 1995 paper.

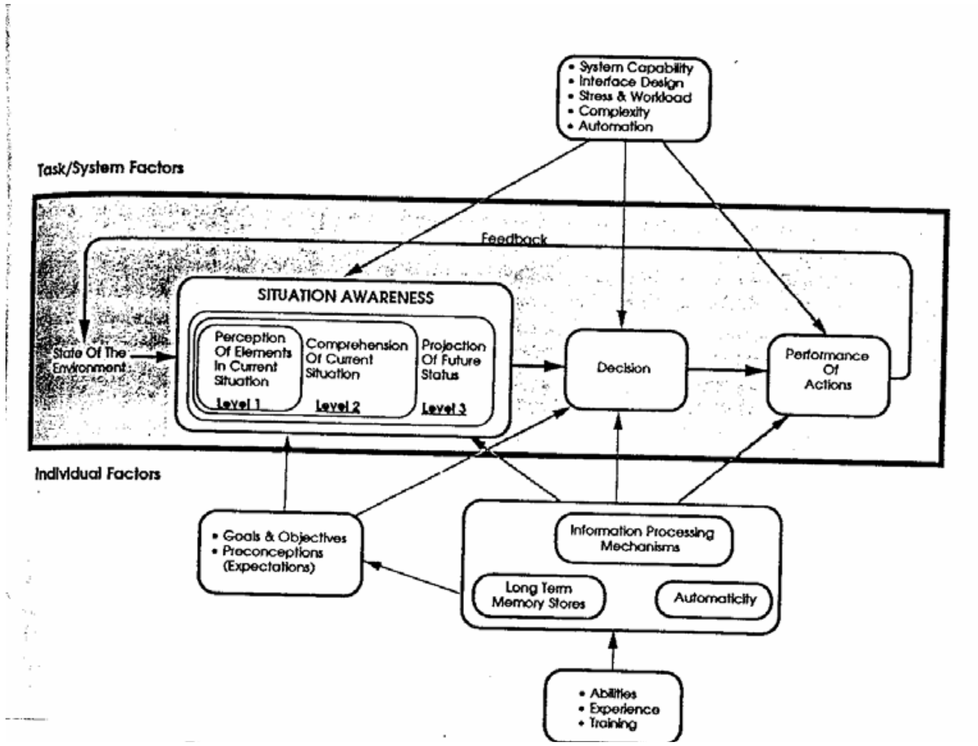

Figure 4. Endsley's model of Situation Awareness

For example, the contents of the SA box in her Figure 1 could easily map onto the hierarchy of information processing described in Rasmussen's "ladder" (Rasmussen, 1984.), and her description of "best match" categorization in SA recognition onto Reason's account of pattern matching in his discussion of the origins of human error (Reason, 1990). Her discussion of perception and the role of expectation is closely related to what one would find in the theory of signal detection. Her "space" can be related to the ecological problem spaces of Rasmussen and Vicente (1989), as can her discussion of display design. And if I may be allowed to cite some of my own work on mental models, there is a striking similarity between the contents of the SA box in her Figure 1 and my conclusions of how mental models are used to guide monitoring (Figure 7).

The most interesting thing about these similarities is that they have come about serendipitously, because the researchers involved were working on different problems. The convergence is remarkable, and Endsley's account of SA serves as an excellent convergence point.

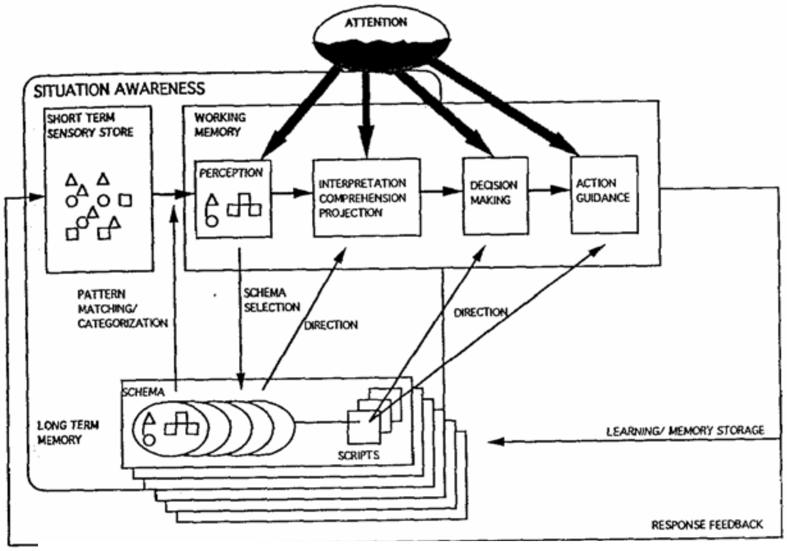

Figure 5. Endsley's model of information processing

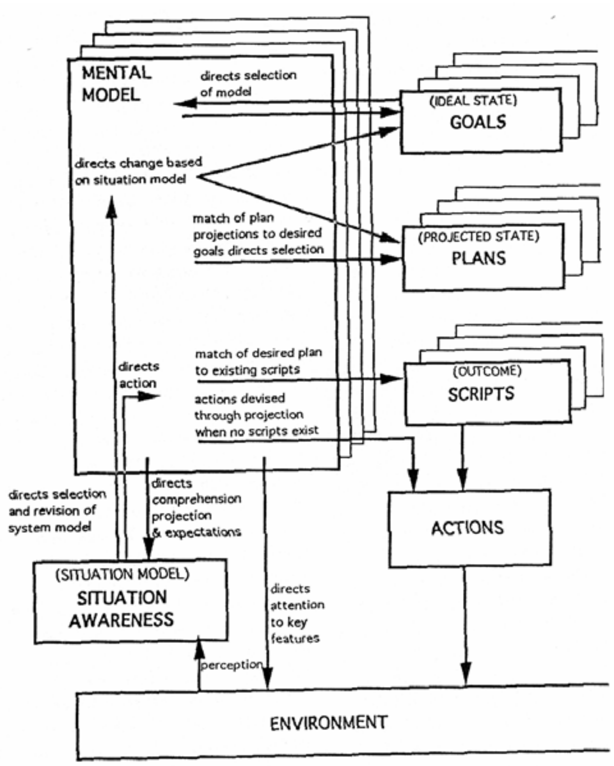

Figure 6. Endsley's flow of control in Situation Awareness

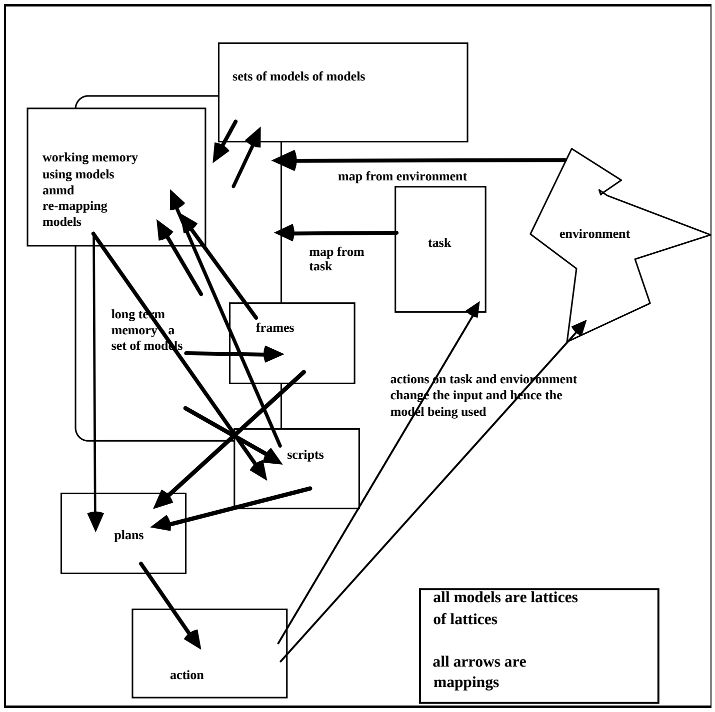

Figure 7. Moray's mental model information flow chart

HOW DO THEY DO IT? REDUCING COMPLEXITY

When we look at the immense complexity of the environments in which expert operators manage to maintain good SA, their performance is really quite remarkable. If we look back at the literature of the last 50 years we find that expert operators seem to manage remarkably well in situations that on a superficial analysis of the bandwidth and complexity of the environment would seem nearly impossible. Given that they are often probably forced to use RBB and KBB, which are known to be very demanding of cognitive skills, how are they doing it? Before looking at some of this empirical evidence, let us consider how in general complexity can be handled.

One way is decompose the environment into groupings of elements, within each group of which there is a strong mutual causal interaction and coherence, and between which there is little. We can think of this as looking at a complex body in terms of its "organs" or "molecules", rather than its "atoms". There are several suggestions as to how this might be done. Recently Yufik and Sheridan (1991); Yufik, Y. M., Sheridan, T. B., & Venda, V. F. 1992; proposed a way to build mental models with just this property. And somewhat earlier (Moray ,1986), I offered a less sophisticated account of the same topic. These papers show how the detection of coupling among variables would allow an observer to reduce the amount of processing needed to keep track of the state of the environment, at very little cost in accuracy. (We shall see later that Beishon (1974) thought that empirical evidence from process control supported this kind of behaviour in ovenmen.) The degree of re-composition of information (from "atoms" to molecules to organs, etc.) would depend on discovering the natural decomposition of the environment, and reasons to think that it is one of the main differences between a neophyte and an expert (Canas, Quesada, Antoli, & Fajarda, 2003). Furthermore, there has been extensive discussion of how to decompose a large system into its natural components even if one does not understand initially how it is put together. (See for example, the little known papers by Conant, 1976). Rasmussen (personal communication, 1993) has suggested that causal models (as distinct from mathematical descriptive models) should be based on decomposing a system into objects and its behaviour into events, rather than simply treating all variables alike as individual entities, and probably the findings of McKinney and Davis (2003) on the use of training for crisis management reflects this.

One consequences for this approach to SA and operator mental modeling would be that automation might be designed both to encourage such skills in the operators and direct their attention to appropriate levels of “molecules”, and also, equally importantly, to keep track of the finer details so that if unforeseen events or system failures occur, the automation can support a return to more detailed examination of state space when that is needed. Note that this is what Rasmussen, Vicente, Woods and other “ecological system” designers recommend, namely information at appropriate levels for the particular cognitive operations form moment to moment. By knowing how the operators mentally decompose the system, strong hints about the design of support for “visual momentum” will be found.

Let us return to the question of timing. Sanderson’s spectrum is but one way to view the rate limitation on cognition. I would like to draw your attention to two others. The first is quantitative. Forty years ago, Wohl, (1961; also 1982), reported very interesting data about the times it took people to repair radar sets. Data from Wohl (1982) are shown in Figure 8.

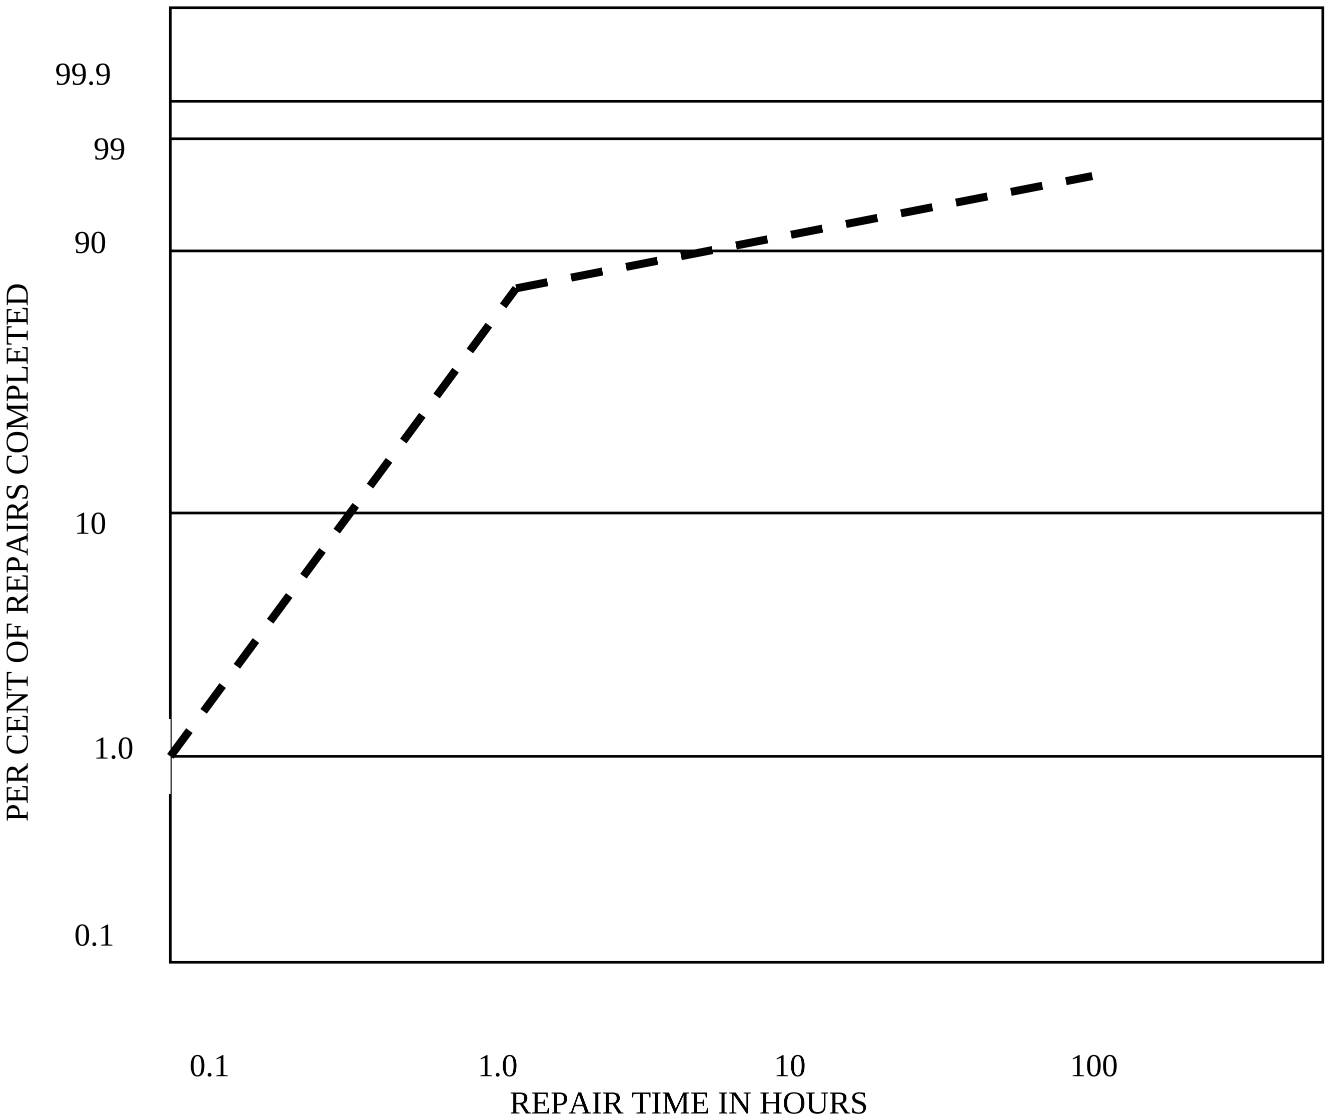

Figure 8. The proportion of radar set faults diagnosed and repaired as a function of time since work on them began. (After Wohl, 1982.)

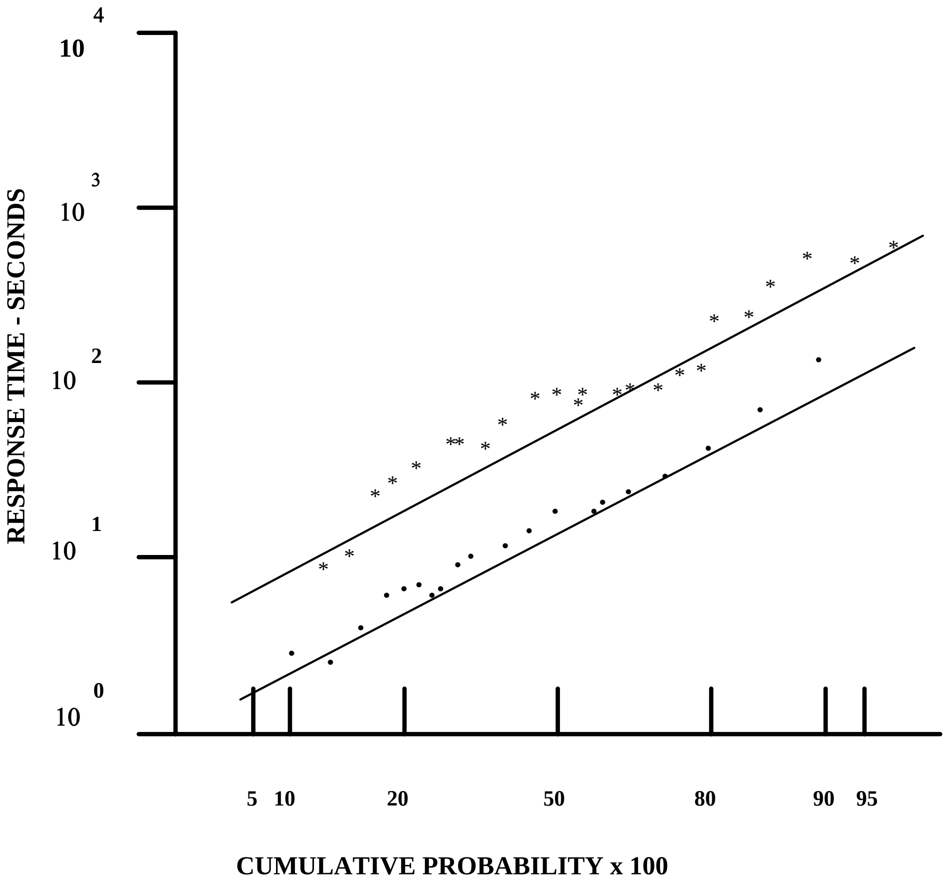

Figure 9. Cumulative probability of response related to time since incident. (After Beare et al 1984.)

Wohl discusses the shape of this plot in terms that today sound similar to the distinction between RBB and KBB. The steep initial slope, in which as time passes the repair men are more and more likely to reach a solution, seems to occur while they using a "test according to well known rules" strategy (RBB). If this finally fails, they are forced to try to think of new things to do (KBB), and from then on the slope of the curve (on a Weibull plot) is less than 45^0, meaning that it is less and less likely that they will find the solution. (A practical conclusion would be, "After an hour, throw it away and buy a new one.") What I particularly want to draw your attention to is the first part of the curve, when the operators are drawing on their deeply learned set of pattern matching rules. In this case about 90% of work on this task is completed in an hour. Consider now another, closely related curve (Beare, A. N., Dorris, R. E., Bovell, C. R., Crowe, D. S., and Kozinsky, E. J.,(1984). Here it is not a Weibull plot, but a log-normal time-reliability curve (Figure 9).

In this study, although the environment is quite different (a high fidelity nuclear power plant simulator,) these curves closely resemble the first part of Wohl's curve. They also resemble nearly two dozen others plotted by Beare et al. for a variety of diagnostic tasks in a nuclear power plant. (Note that they are plotted in an unusual way, the independent variable being on the ordinate, and the dependent variable on the abscissa, unlike Wohl's data and most others that have used this kind of plot. In Beare et al.'s version a flatter slop is a faster slope.) Much simpler systems show similar curves, as do the data in Wohl (1961), but they are flatter, i.e. approach completion faster.

Such curves could be used to calibrate the success of automation intervention. They provide an estimate of what can be expected from the general population of operators. If working with automation lowers the slope, then automation is assisting diagnosis. In Beare et al. there is also a feature that suggests that the operators are simplifying the environment. The upper curve in each case is based on field data for operators' responses to real faults in the plant: the lower curve is based on responses to simulations of the same faults. In all but one of Beare et al.'s plots the curves are linear and parallel, the curves for the simulated data being almost exactly one log unit lower than the field data, that is, ten times faster. An obvious explanation is that operators in the simulator are excluding a large range of events that are possible in the real world but are not believed to be plausible in a (limited duration but very expensive) simulator training session. In all the published curves that I know of data are pooled over different operators and different occasions. They are not based on repeated trials by one operator. That is, at present these curves represent a normal distribution of (probably Rule Based) talent over the population of operators. If anyone is looking for a simple but interesting Ph.D. project, they could look at the extent to which plots such as these are similar to those from single operators performing a task many times. (The implications, as reports like to say, "is left as an exercise for the reader".)

These plots are another way of describing the boundaries of rate of information processing in complex environments, that is, defining the envelope within which SA can be effectively performed. I would like to consider one more way of representing rate limits and that is by the use of Minkowski diagrams, one of which is shown in Figure 10.

The original Minkowski diagram was invented (Minkowski, 1909) to represent the implications of the fact that nothing can travel faster than light, and what is represented is spacetime, not space and time. The "world line" represents the history of an entity. At the intersection P(0,0) ("here and now") there are some parts of the universe whose state is hidden from the entity because there has not been time for information traveling at the speed of light to reach it. Similarly there are regions of the future that the entity cannot causally affect because there will not be time for any control signal to reach them, limited again by the speed of light. Here I will use a "pseudo-Minkowski" space to discuss the constraints on keeping SA and intervention up to date. (Peter Hancock (personal communication) has independently suggested using these diagrams.)

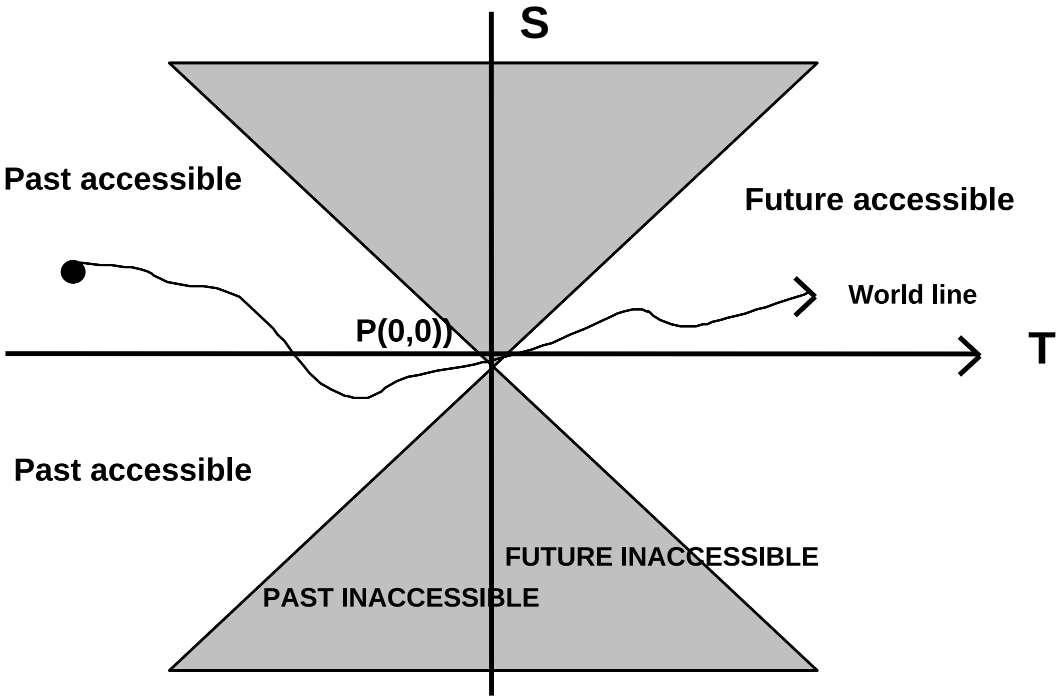

Figure 10. Minkowski Space-Time Diagram

Figure 10 shows such a diagram adapted to our purposes. At the "here and now" is our operator, O. The vertical direction is a dimension representing a set of facts about the system state. The horizontal axis is real time. O's knowledge of what has happened in the environment is limited by the time it takes to access information about it. For example, the slope of the accessible/inaccessible boundary on the left of "Now" will depend on how fast attention can be switched, how fast eyes can be moved, how fast the interface can deliver state information from sensors, and so on. To the right of "Now" is O's ability to affect the world causally in the near future. The slope in general will be different from that prior to now, because the time it takes to perform actions, which depends on the biodynamics of muscles, on the speed with which sensors act, on how long it takes to call up screens with icons for actions, etc., will be different from the rate of sensing the environment.

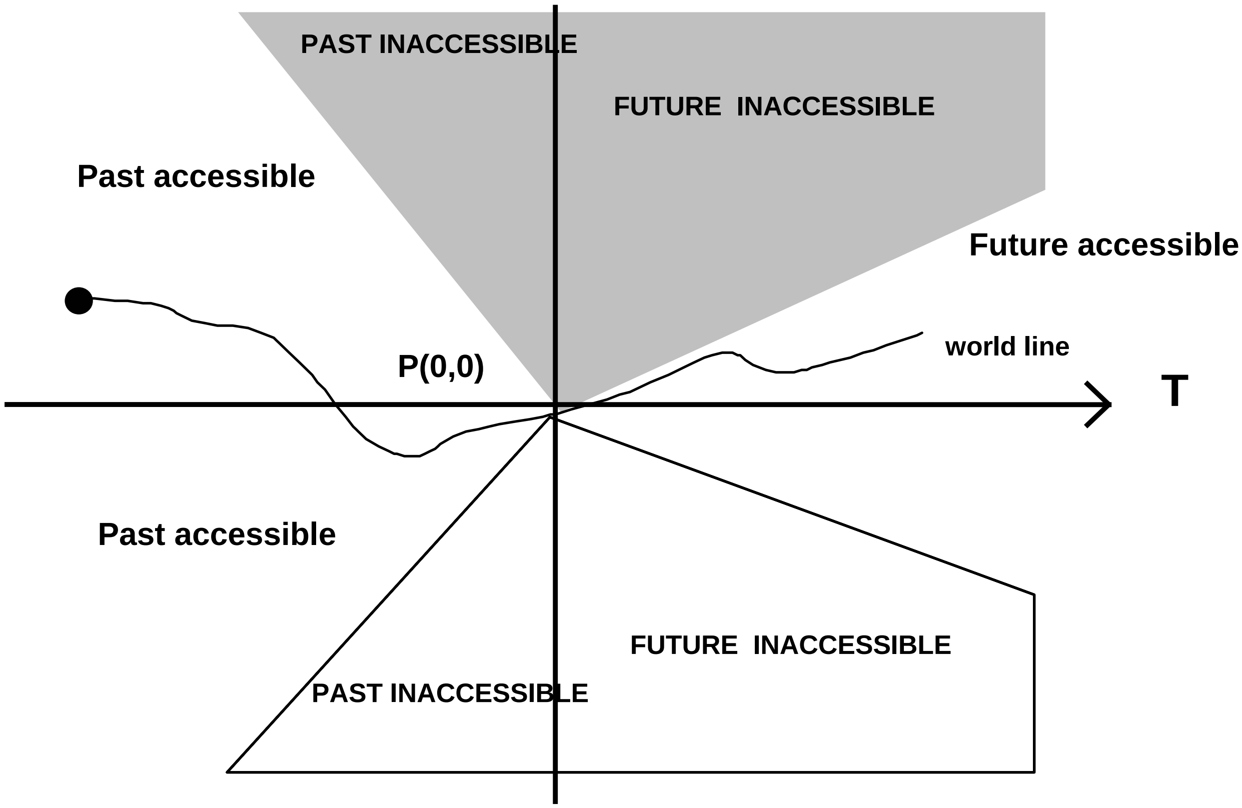

Figure 11. A pseudo Minkowski space for assessing human-machine systems

The problem for O in maintaining effective SA is clear in such a diagram. The properties of the human-machine system and the properties of the environment outwith that system determine the slopes of the Minkowski space, and hence what is in principle accessible, either for updating O's information about the state of the world, or for affecting its future. The task for automation is also clear: it is to make the slopes of the accessible/inaccessible boundaries as steep as possible, thus reducing the inaccessible regions. If we were able to determine the actual structure of the real information space, and the slopes of the boundaries, we could, for example, decide whether to put effort into the input or the output side of the system (the sensing or acting automation in Parasuraman et al.'s formulation). In Figure 11 it is clear that above all, automation is needed on the output side.

Just for fun, in Figure 12 we can represent the effect of predictor displays or fast time models, whether mental or automated. They are valuable because they let us know the future *in the inaccessible regions.*

I find this representation interesting and stimulating, but I have to admit a small weakness. I think that one could define the slopes of the boundaries for specific aspects of a task, using what we know about eye movements, Card et al.'s (1983) Model Human Processor, and a direct examination of a particular interface. But I have completely failed, despite some fairly vigorous efforts, to discover a general metric for the space as a whole, or even really to be satisfied that I understand what the notion of vertical "distance" really is as a property of the human-machine system and the environment outwith it.. Again, I am happy to leave that problem to the reader.

That is probably enough speculation about theory and the future. I now want to remind you of how much is already known about SA and how operators become coupled to their environment at the RBB and KBB levels. When we have reviewed the empirical data, we will be able to see what a tightly coupled RBB system would be like, and hence how a computational predictive model may be developed.

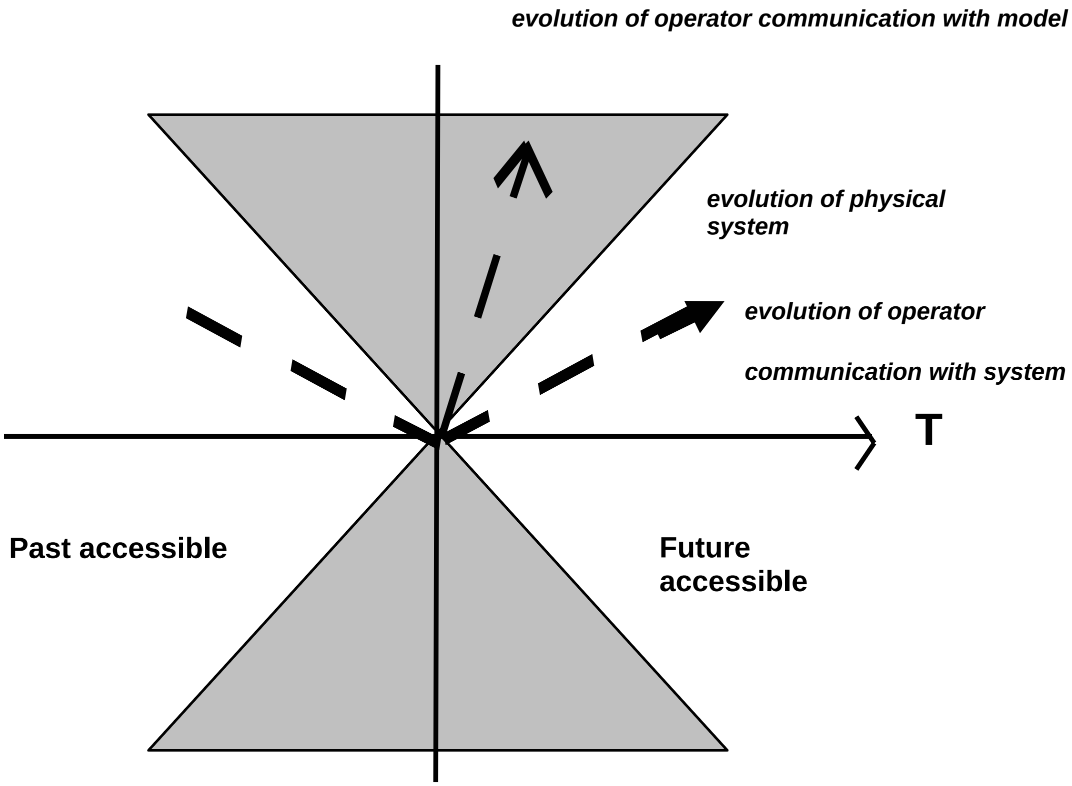

Figure 12. Minkowski Space-time Diagram with operator knowledge and model of process

ON TO THE PAST: SOME EARLY EMPIRICAL WORK

As many of you will know, Europeans view the recent American enthusiasm for the study of naturalistic decision making with a slightly ironic eye, since there is a long-standing tradition of field work in human industrial decision making in Europe that is at least fifty years old. I want to draw on some of that early work to support the claims by Endsley and others concerning the characteristics of SA.

As long ago as 1960, Crossman was describing supervisory control, and the effects of automation in steel, chemical and petroleum industries. In his paper on "Automation and Skill" there are diagrams that anticipate Sheridan (1978) and show the flow of information in manual and "automatic" control. Here are some quotations from his work (Crossman, 1960).

> "Analytic study suggests that a specific control skill comprises five components:
>
> 1) *Sensing* - the ability to detect the signs and indications such as "noises, smells and appearance, which indicate how the plant is running.
> 2) *Perceiving*- the ability to interpret these signs and the instrument readings in relation to one another, and to infer what is happening.
> 3) *Prediction* – of what is likely to happen in a given situation of the controls are left alone.
> 4) *Familiarity with the controls* – knowing what means can be use to influence the process, what their effects are, and how they interact with others.
> 5) *Decision* – the ability to select the control action most likely to achieve the desired result in the given circumstances or to avert unfavourable developments when they threaten."

The last item, decisions, can be carried out in several ways.

i. “The operator may follow a ‘rule of thumb’ – doing what has always been done in a given situation, or what worked last time; but this allows little flexibility.

ii. He may use a ‘mental model’ or idea of the process, on which he can try out the different possible control actions in his imagination and pick the best bet. A good operator seems to ‘feel’ his way into the process, becoming *intuitively* aware of what is going on and what to do about it.

iii. The operator may use a logical approach and consciously reason out the meaning of things, analyze the situation, and come to a *rational* decision.

On the whole, discussions with operators have suggested that the first, or ‘rule of thumb’, methods is common among the less good operators, and the second, or ‘intuitive’, method is often characteristic of the better ones.

In another of his papers he discussed strategies of monitoring and sampling behaviour, both for the demands made by the dynamics of the process being controlled, and also in regard to the use by O of mental models (Crossman, Cooke and Beishon, 1974.) In particular he drew attention to the fact that the Sampling Theorem was an effective model for a limit on sampling, but that it needed to be modified by cognitive factors, to produce what Crossman called “the effective bandwidth” of the task:

> “The more normal use of human vision evidently is to create and maintain an internal image, map or model of the environment, from which information can be extracted to determine future action.”

“Conclusions from field and laboratory studies

1. The operator’s basic minimum rate of sampling in the two process tasks studied was determined from the system bandwidth as predicted by Sender’s application of the Shannon-Wiener sampling theorem, provided that account is taken of the allowed error tolerance by calculating an ‘effective bandwidth’.
2. However, a much more detailed analysis of factors n contributing to the operator’s uncertainty, its rate of growth over time, and the cost attached to sampling is needed to give an even moderately accurate estimates of sampling rate in the various circumstances encountered when they rose above the minimum.
3. The problem of sampling could not be divorced from the more general problem of control, which in turn raised questions of the required accuracy, cost of error, operator’s knowledge of system structure, degree, type, and predictability of disturbance, and effects of response lag.
4. While forgetting was not positively identified as a cause of increased sampling rate, the data were consistent with this possibility.
5. The data suggest the following empirical generalizations about sampling behaviour which agree fairly well with predictions from the uncertainty analysis given above.

a. When a variable is at its desired value and the system is correctly adjusted so that there is no residual drift due to small errors of control setting, the ‘background’ sampling rate is determined by the highest frequency component of random disturbance that has an *amplitude great enough to cause excursions exceeding the allowed tolerance*. This is the ‘effective bandwidth’ which the system presents to input noise.

b. When a variable is within the specified range but the system is not quite correctly adjusted so that it tends to drift off, sampling rate is determined by the rate of drift, and rises *when the variable is near either of the limits of its tolerance range.*

c. When a variable is outside its tolerance band and the operator is making large stepwise control changes in an attempt to correct it, *a sample is taken after each control change at a time when the response is expected to have reached some 80% of its final value* << about 2 time constants>> … (this rule) only applies when the operator is uncertain of the precise effect of control changes.

d. Sampling rate rises when whenever general observation of the system or its surroundings shows that anything unusual may be happening, even though it is not known to be relevant to the particular variable sampled.
e. Operators may estimate the rate of change of a variable either by prolonged observation during one sample, or by remembering its value at one sample and comparing it with the next. In general they do not attempt to estimate higher derivatives.

He also commented on the need to mix quantitative and qualitative analysis, anticipating the ideas of De Keyser (1981) on what she called "informal information":

> "The analyst should note that not only quantities expressed in numerical form are variables, but also factors assessed subjectively such as state of a firebed or the turbulence of a liquid. The following five factors normally govern the sampling rate of a given variable:
>
> 1. Its (true) bandwidth – roughly speaking the speed with which the system can permit the variable to change.
> 2. The amplitude of any random disturbance or *noise* which may affect it.
> 3. Its *tolerance* the latitude for variation about a desired value which is permissible without incurring a penalty.
> 4. *Predictability* …operators can often forecast future system state changes from known patterns of behaviour or from auxiliary information.
> 5. *Control calibration* …sampling rate is increased whenever the operator makes control adjustments because a variable's response to changes in a relevant control setting is imperfectly predictable."

Perhaps the most remarkable of the early work is that of Iosif (1968a,b; 1969a,b) performed in Rumania, using the crews of manually controlled fossil-fuel power stations, refineries and simulators. This work is based on *nearly 700 hours of observations of skilled operators* and how they scanned the instrument panels during process control to maintain knowledge of the system state, in other words how they maintained situation awareness. His conclusions are as follows. The monitoring strategy is a function both of subjective factors and of the physical characteristics of the process (the environment). The monitoring strategy depends on the operator having a deep understanding of the technical process, and on the attitudes adopted, and the latter particular include the degree of "prudence" of the operators (Os), (Iosif specifically mentions confidence and Os' decision criterion,) and on the degree to which Os anticipates disturbances. Iosif actually derives a statistic "ε" which is a measure of "prudence". This statistic depends non-linearly on the frequency of disturbances, and depends also both on the technical characteristics of the plant and on individual differences, and seems to relate also to how long it takes Os to notice that a disturbance has occurred. Anticipation of disturbances is very important and is also a function of the frequency of disturbances. (This ability of Os to anticipate disturbances, and then to use the instrumentation not to inform them that a disturbance has occurred but to confirm their anticipation is again very reminiscent of some of the observations by De Keyser (1981, 1987) some 15 years later in Belgium.) What is important is the way Os use mental causal models to anticipate the evolution of the process. This means that even for rare disturbances, they tend to be looking in the right place when the disturbance actually occurs. In Endsley's terms, SA permits prediction, and monitoring is not a stochastic process, but one driven by a recognition of system state and a prediction of its future value.

Iosif's description of his operators is strikingly similar to the optimal control model, in which observations are made to establish the state vector of the controlled process, related to a model of the process in the Kalman filter, and from moment to moment choose either the past history embodied in the model, or the recent observations, to minimise the error of their estimate. Iosif claims that Os make strategically appropriate observations because they have a mental model of the frequencies of disturbances, and in addition use current observations to understand the evolution of the process. They also make use of observations on one variable to predict the values of causally related (tightly coupled) variables which they then do not need to observe.

> "The ability to anticipate is the result of a deep understanding of these correlations, leading to a prompt detection of variations, observations of values that signal certain effects."

This of course is another way of reducing the complexity of the system: as I suggested earlier Os are using the values of observations on one part of a "molecule" to predict other parts of the same molecule without observing them, as I suggested in Moray (1976)). Iosif also notes, as did De Keyser (op. cit.) that operator use information from other members of the team, from telephone conversations, etc., to supplement the "official" information displayed on the console.

Iosif also performed some studies in a process control simulator. He found that the vast majority of fault diagnoses required only the initial hypothesis or a hypothesis revised on the basis of operator "deep causal knowledge". He was able to understand about 80% of the diagnoses made. In about 70% of diagnoses the first hypothesis was correct, or deep causal knowledge led to a rapid revision of the hypothesis. Only about 6% used guesswork and only 1.5% failed to diagnose the disturbance.

Compare this with Endsley (1995):

> "There is evidence that an integrated picture of the current situation may be matched to prototypical situations in memory."
>
> "…(SA) should explain dynamic goal selection, attention to appropriate critical cues, expectancies regarding future states of the situation, and the time between situation awareness and typical actions"
>
> "…(SA) is the perception of elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future."

Iosif's data could almost have been designed to support Endsley's description, although anticipating the notion of SA by 30 years. I commend his papers to your attention. They deserve to be far better known. Some of his results are shown in Figures 13 – 17.

And as we have seen, many of Crossman's observations from the 1960s support Endsley's proposals (Crossman, 1960; 1964).

There is a great deal of empirical data that should be regarded as the fundamental data for any discussions of these problems. An outstanding source book for much of this work is *The Human Operator In Process Control* (Edwards and Lees, 1974). To give you an idea of the richness of these data, Lees (1974) in that book cites 135 references specifically to the behaviour of the human operator in process control, and in another book by Edwards and Lees they list more than 2000 studies of process control!

Another place of where Endsley's discussion of SA resembles earlier work is her discussion of the role of expectation in perception . This could be fitted into to the framework of the Theory of Signal Detection (TSD), and her discussion of the way in which Os use a "best match" categorization of observations on the environment is very similar to James Reason's "pattern matching" hypothesis in his discussion of human error. Her discussion of display design (p. 51.) and her notion of the "space" to which SA is relevant would fit very well with Vicente's discussions of ecological displays, and in general her ideas of system design anticipate those of Parasuraman et al. (2000) on automation. The SA box in her Figure 1 could contain something very similar to a Rasmussen "Decision Ladder". It is even worth perhaps going back to the work on perceptual learning and concept formation in the 1950s, such as that of Gibson and Gibson (1955) to understand how Os assemble the perceptual information for SA, although of course the early work will need to be re-interpreted for our current context.

Let me make plain that I am not criticizing Endsley for lack of originality. Far from it. As she herself says, she is trying to synthesize previous work, and her paper is masterly in that respect. Indeed since I first read it I have thought that it is the only paper that does justice to the complexity of SA. I want to emphasise how much we know, and suggest that if there is so much agreement among major researchers, surely we could move from analysis to a predictive model, and hence to more powerful support for decisions about automation and allocation of function.

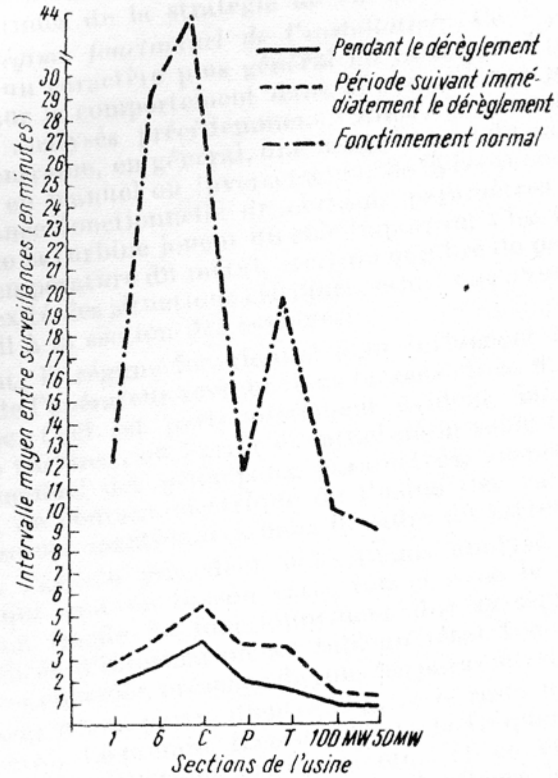

Graphique 7. — Fréquence d'observation des situations du processus technologique à un moment donné.

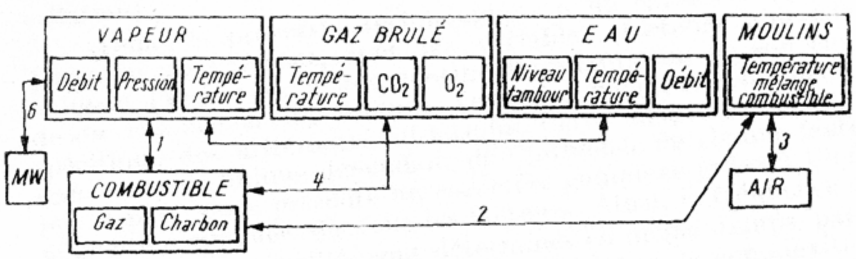

Figure 13. Intervals between successive monitoring as a function of the subsystem and whether an abnormal disturbance has occurred. (Iosif, 1968)

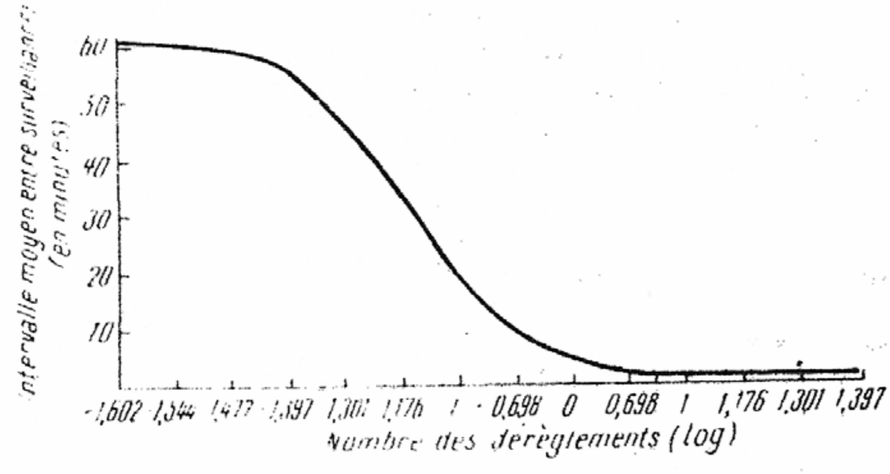

Figure 14 Interval between observations as a function of log(number of disturbances. (Iosif, 1968)

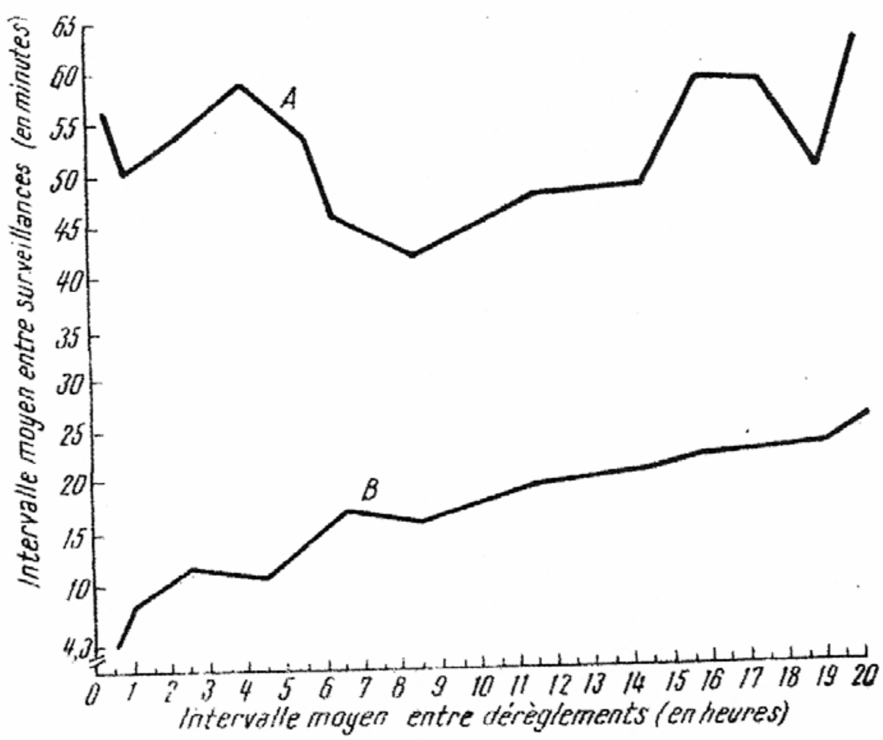

Figure 15 Mean intervals (in minutes) between observations as a function of intervals (in hours) between disturbances. Curve A an unimportant variable. Curve B a critical variable. (Iosif, 1968)

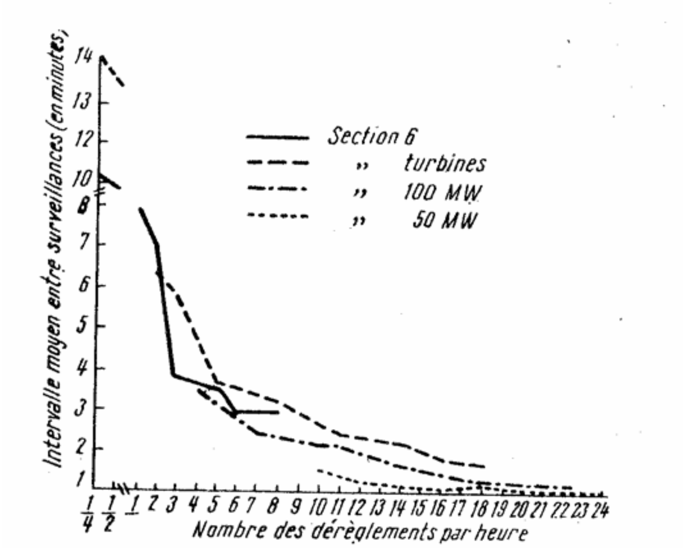

Figure 16 Mean intervals (in minutes) between observations as a function of mean intervals (in hours) of disturbances for different plant (Iosif, 1968).

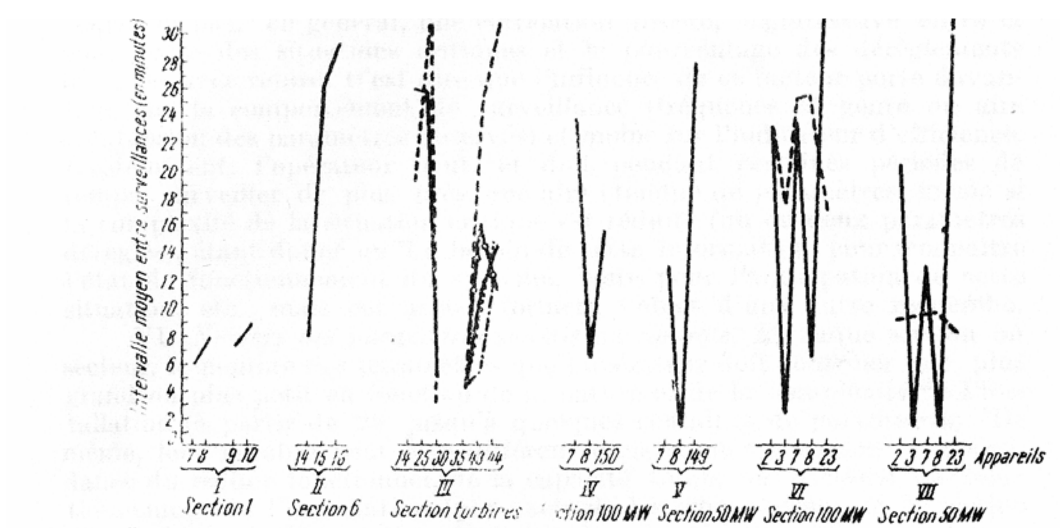

Figure 17. Mean interval (in minutes) between observations on groups of variables that are strongly correlated. (Iosif, 1968.)

I think there are two central functions in SA, from which others follow, (1) a monitoring strategy, and (2) recognizing the state of the system based on the resulting observations. What do we know about these functions that could be used to develop a computational model?

MODELS OF STRATEGIC MONITORING.

In the case of monitoring there is almost an *embarras de richesse* Moray (1986) reviewed nearly a dozen models for monitoring, and since then others have appeared. From what Iosif and Endsley have said about SA monitoring, it is clear that simple stochastic model such as that of Senders (1964) will not be appropriate. We require at least a model that responds to an observation that the value of the variable is near to a limit by increasing the probability that another sample will be taken soon after. There are several such models, such as those by Senders, Elkind, Grignetti and Smallwood (1965); Carbonell (1966), Sheridan (1970), and Moray, Richards and Low (1980). There are abundant empirical data going back to Crossman et al..(1974) and even to the work of Fitts 55 (!) years ago (Jones, Milton and Fitts, 1949). However, we need a model that contains more of a semantic component. Iosif found that people made use of the structure of their mental causal models and of the correlation among the values of groups of variables, and Rasmussen and Batstone (1989) suggest that people actively probe the boundaries of system performance, deliberately allowing systems to reach constraint boundaries to learn how to handle the dynamics of the system:

> "Proposition 16. Designers should carefully consider that humans are boundary seeking and that their reliability depends on the opportunity to touch the boundary to loss of control and to learn to cope with recovery.". . .
>
> Proposition 20. Before you consider to support (sic) operators by solving their problems through predictive situation analysis, make sure they have available the information about:
>
> 1) actual system state
> 2) design basis and intended function
> 3) boundaries around acceptable performance and problem
> 4) available resources for action.
>
> Consider that the information is structurally related and that integrated symbolic displays can be based on primary data."
>
> (Report to World Bank,1989)

Endsley also speaks of pilots and other operators pushing at the boundaries of performance.

Sometimes reliance on correlations among variables may lead to cognitive tunnel vision (which incidentally was reported over 50 years ago by Russell Davis (1948) in the context of prolonged attention in pilots). Endsley makes the point that in automated systems SA (the knowledge of current system state) decreases. This is almost a matter of logic rather than an empirical fact: if the environment is less observable at an appropriate bandwidth, then its state will be less accurately identified – unless the automation is arranged so that O can predict unsampled values on the basis of correlation and causal modeling.

One model that could be looked at more closely uses queuing theory, (Senders and Posner, 1976). But still to my way of thinking it is too close to a purely stochastic model. Certainly we need something that is sensitive to system bandwidth. When Endsley speaks of SA being concerned with a certain volume of time and space, we can see the space as being the number of sources of information (either in terms of the instruments on a control panel, or as the number of locations in the environment in which a significant event may occur when flying at attack aircraft). The volume of time is of course directly related to the bandwidth of the system, (or rather, as Crossman says, the "effective bandwidth" (Crossman et al., 1974). That in turn depends on where in Sanderson's spectrum our task is located, which in turn (for a task such as flying an aircraft) depends on what maneuvers the pilot chooses to make, and how fast he or she flies. For a very high frequency of events SA must be updated very rapidly. For very low frequencies updating at infrequent intervals is sufficient providing that the correct variables are sampled and memory adequate. Iosif was dealing with a low frequency system, where some significant events occurred at intervals of several hours, and it is desirable to extend the concepts out into the ultra-low-frequency (ULF) region of project management, thus generalizing a model into the realm of "organizational accidents" (Reason, 1997).

A very interesting sampling model, which appears never to have been developed beyond its original definition, is that due to Milgram (1983). Its great advantage is that in addition to covering high dimensional multivariate displays Milgram claims that it can be related to the perception of patterns, citing Wewerinke (1980).

Based as it is in optimal control theory, Milgram's model is attractive for two further reasons. The first is that the covariance matrix at the heart of the model is a natural representation of coupling between variables, and hence of the dependence on correlations to reduce complexity; and secondly because it naturally leads to a consideration of what action to take in response to an observation (and hence to the updating of the system state estimate). This is important because as we have seen, Crossman and his co-workers found that scan patterns were affected by the discriminability of signals (signal/noise ratio, as Senders, 1964 predicted) and also by the actions taken by operators (as Iosif also found). Following the adjustment of a control which forced a variable to a new value, Crossman found that the next look at the value of adjusted variable tended to be when it had reached about 80% of its final value. Crossman noted that the bandwidth of a process control system, which seems at first sight to be toward the right hand end of Sanderson's spectrum, is realty higher: what he called "the effective bandwidth" was the highest frequency of random fluctuation that had a peak amplitude that could exceed the tolerance band of the process. This again relates to Iosif's discussion of the effect of frequency of disturbances. Moray and Synnock and Sims (1973) on the other hand, in a task where the bandwidth depended on the velocity generated chosen by O (roughly, a laboratory equivalent of driving at a self-selected rate along a twisty road) found that observers sampled at fixed intervals, apparently as a function of forgetting (see also Moray, Richards and Low, 1980 for a further discussion of forgetting-driven sampling; and Senders (1983) on monitoring while driving. and Moray and Inagaki, 2002 for a recent attempt to formulate a general model.

To sum up this section, we need a model of strategic sampling that can encompass not just bandwidth and tolerance threshold, but is sensitive to effective bandwidth, memory liability, correlations (that represent causal coupling) among variables, relative values of different variables, control actions and their effects, and can handle *patterns* of variable values as its effective input. Milgram's model seems to me very promising. It could handle both SBB and RBB if we knew how operators use the patterns of information they construct from sampling. Like all the sampling models it is quantitative. What is remarkable, when one looks back to at the literature of the 1960s and 1970s is the very large number of quantitative models that have been proposed. I must plead guilty myself for joining in the lemming-like tendency to rush down theoretical slopes into the conceptual sea without sufficient attention to what has already been achieved (see, e.g., and Moray and Inagaki,). There is really no excuse for not stopping, drawing breath, and reviewing what they all have in common.

HOW IS KNOWLEDGE REPRESENTED IN SA?

We now come to the second topic which I want to look at in the light of old data, namely the nature of operators' knowledge. Obviously Os acquire the values of variables, be they environmental or on control panels, as a series of sequential perceptions. But we have seen good reason from the empirical and theoretical literature to believe that they are not stored in that way. Forty five years ago Miller (1956) spoke of "chunking" in short term memory, and, as we saw, Yufik & Sheridan (1991), Yufik et al. (1992) , and Moray (1986) have all discussed possible ways that the complex knowledge might be simplified in long-term memory. At the same time there is widespread agreement among people who have studied workers in complex processes that reasoning (KBB) is rare. Rather, Os most commonly use RBB, that is to say, "IF (state X) THEN (action Y)" rules. (Rasmussen, 1984; Woods, 1984, 1991; Wohl, 1982;). When we look at the old empirical literature there is strong support for this. We will examine two studies.

The first is by Beishon, (1969) who studied the control of bakers' ovens in an industrial bakery. The task was complex. What we might call embryonic cakes, loaves and buns entered an oven in batches, but as a continuous process, and had to be baked for different times, and at different temperatures, without waiting to cool down or heat up the oven between batches. O had to use an extensive knowledge of different kinds of dough, the effect of the size of the loaves, etc., temperatures, composition of the mix of products, etc., and observations from moment to moment of the oven temperature, to determine the time of baking and the changes in the heater settings. Beishon's conclusion was that his Os did not work these out, but used a large mental look-up table. That is, the signals for control formed a complex pattern, what we could call the left hand side of the "IF…THEN… " rule. As with Crossman, I will quote extensively from his report: