

Analysing and modelling train driver performance

Ronald W McLeod¹, Guy H. Walker², Neville Moray¹ and Ann Mills³

¹Nickleby HFE Ltd, Hillington Park Innovation Center, Glasgow, G52 4RU

²Rail Safety and Standards Board, Euston, London, NW1 2DX

³Brunel University, Ergonomics Research group, Uxbridge, UB8 3PH

Keywords: Train driver performance; Automatic Warning System; Models of human performance; human reliability

Abstract

Arguments for the importance of contextual factors in understanding human performance have been made extremely persuasive in the context of the process control industries. This paper puts these arguments into the context of the train driving task, drawing on an extensive analysis of driver performance with the Automatic Warning System (AWS). The paper summarises a number of constructs from applied psychological research which are thought to be important in understanding train driver performance. A "Situational Model" is offered as a framework for investigating driver performance. The model emphasises the importance of understanding the state of driver cognition at a specific time ("Now") in a specific situation and a specific context.

Introduction

The work reported in this paper was carried out as part of a study for the Rail Safety and Standards Board (RSSB). The aim was to understand and assess the risks of driver unreliability associated with extended uses of the Automatic Warning System (AWS) on the UK rail network. The study is summarized in references 1 and 2.

Despite what, to all everyday experience, appears to be an extremely powerful attention-getting device, combined with a highly visible visual reminder, there is a significant, and recurring risk of AWS failing to prevent experienced train drivers passing signals at danger. According to Vaughan (3), *"No-one ever thought that the driver would ignore the warning. It was utterly taken for granted that a driver would always take notice of the signal. (p.12)*. In connection with the crash at Watford Junction in 1996, Hall (4) wrote that; *"It seems inconceivable that a driver can acknowledge receiving two warnings and yet take no action to apply the brake to stop the train at the red signal, yet it happens". (p.97)*

Although relatively straightforward conceptually, the ways in which AWS is used in the UK introduce a number of complexities for the driver. These reflect the

complexity and variability of UK signalling, as well as the various ways AWS is used. These include;

- Δ AWS does not adequately discriminate between the possible states giving rise to the alarm. The system therefore depends on the driver having an adequate appreciation of the existing situation to be able to correctly interpret the alarm.
- Δ There are a variety of situations where the sunflower can refer not to the immediate past signal, but to a signal some time prior to that.
- Δ The time frame of AWS activity (i.e. the time period over which any AWS signal is 'active' in terms of conveying information about the track ahead) can vary from a few seconds to possibly many minutes.

Train driver performance is dominated by cognitive and perceptual factors. Much of the existing human factors research in the area draws on relatively simple models of human decision making and information processing. Traditional, hierarchically-based forms of task analysis are at the core of much of this research. From the perspective of cognitive psychology and cognitive engineering, these approaches are held to be fundamentally flawed. Specifically, they are not able to capture the contextual and situated nature of human performance in general, and train driver performance in particular.

Vicente (5) offers a compelling critique of traditional approaches to Task Analysis. Similarly, Hollnagel (6) provides a detailed critique of human reliability estimation techniques from the perspective of cognitive engineering. The arguments made by both Vicente and Hollnagel are extremely well made, and may be considered definitive. There are doubts about whether the techniques they propose overcome these limitations (Cognitive Work Analysis (CWA) in the case of Vicente and the Cognitive Reliability and Error Analysis Method (CREAM) in the case of Hollnagel), or whether they will translate from a process control context to train driving. Nevertheless, we believe that many of the limitations and weaknesses of traditional forms of task analysis and human reliability estimation they identify are equally valid for understanding train driver performance.

Other than in the sense of identifying factors likely to influence or shape performance little if any of the existing research has tried to understand the mechanisms by which contextual and situational factors influence driver performance. The relative simplicity of the psychological basis of the existing work, and the persistence of SPAD incidents, suggests that more comprehensive insights are needed.

Recent thinking in the psychology of real-world human performance

There is now a large, established and well respected body of thinking and research that focuses on understanding human cognition and performance in the real world contexts in which it occurs. Recognising the psychological complexity of the AWS problem, a number of ideas from this wider knowledge-base were reviewed. These go beyond the immediate literature on human factors aspects of AWS. The aim was to introduce a number of areas of thinking about cognitive performance in real world tasks that seemed relevant to understanding the risks associated with train drivers use of AWS.

The following sections very briefly introduce some of these areas. A more detailed review, with full reference list, is available in (7).

Strategic behaviour: The importance of strategy and strategic behaviour in human performance has been recognized, at least in the research literature, for many years. In particular, strategy is essential in helping the human maintain performance in situations of stress, very high or very low workload or when subject to other influences. The concept has however rarely been used explicitly to try to understand or explain human behaviour and performance in real-world situations. Recent research by Merat, et al. (8) into train driver eye movements provides contemporary insights into strategic behaviour in train drivers.

Situation awareness: Situation awareness is possibly the most widely cited construct in understanding human performance in complex real-time systems. Not surprisingly, it is also considered a key psychological construct in safe train driving.

Loss of situation awareness leads to what Sarter and Woods (9) have described as automation “surprises”. In connection with the Purley rail crash in 1989, Vaughan (3) describes “*Driver Morgan of the Littlehampton train [...] completely at a loss to understand what had happened*” (p. 91). Vaughan notes that driver Morgan must have acknowledged and cancelled the two previous AWS warnings leading up to the red signal, but failed to initiate an appropriate braking maneuver until it was too late.

A number of factors can potentially interact to determine the train driver’s understanding or belief about the current situation and therefore how an AWS alarm is interpreted. These include; the nature of the alarm (bell or horn); visibility of signals and magnets on the track ahead; the driver’s interpretation of the nature of the preceding alarm; expectations about the current location (from route knowledge, as well as WONS, PONS and Late Notices); the train speed to be achieved, and by what time or by what location it is to be achieved.

Changes in any of the factors can potentially lead to a change in the driver’s interpretation of a particular AWS alarm. In combination, they have the potential to cause the driver to misinterpret what an AWS alarm refers to, and what action to take.

Situated behaviour: The term 'situated behaviour' refers to a broad group of areas of psychological thinking. They all seek to take the understanding of human performance out of the laboratory context, and to ground it in the real-world environment in which it occurs. The common themes are that the context and particularly the situation at the time, have an extremely powerful influence on how the human performs in real world tasks.

The work of Rasmussen, Pedersen and Goodstein (10), Vicente (5) and many others on the analysis of cognitive work¹ takes the view that it is not possible to understand human performance without locating it within organizational and technological constraints. They demonstrate why traditional normative analysis approaches are not well suited to understanding complex socio-technical systems because they under-

¹ Note that this approach is concerned with the analysis of cognitive work: it is not cognitive task analysis.

estimate (if they capture them at all), the context dependent aspects of human performance.

Distributed cognition: Ecological-based research such as Ed Hutchins's ground-breaking 'Cognition in the Wild' (11) and much of the varied research into Distributed Cognition emphasises the critical role that the situation and environment play in cognitive performance. The Distributed Cognition concept recognises that cognitive performance draws on, and is supported by, many artefacts of the organisation and surrounding environment.

It is, for example, clear that the placement of a prominent AWS magnet in the middle of the track can help to cue the driver in advance of a warning sounding. The driver is also likely to take advantage of many other environmental cues. These might include TPWS antennas, hot axle box detectors, other incidental AWS sounds (such as a relay's clicking) and a myriad of other track/trackside artefacts.

Cognitive control modes and subjectively available time: Hollnagel (6) has developed a model of cognitive performance (COCOM) that recognises and seeks to account for the critical role of context in cognition. His ideas of cognitive control and subjectively available time seem potentially important in understanding train driver behaviour. Route knowledge supports anticipation and future-orientated behaviour. It allows the driver to think ahead, and helps control the allocation of cognitive and perceptual resources based on expectations about the future. It also helps the driver in spotting and interpreting cues and other information. In the COCOM framework, the effect of route knowledge might be to increase subjectively available time, thereby allowing the driver greater cognitive control over performance. Among the more important aspects of tactical behaviour is the (often unconscious) scheduling of visual attention, which switches the eyes to and from the different sources of information

Ecological optics: Following the lead of JJ Gibson (e.g. 12), there is a very large body of knowledge about the ways in which humans, as well as many other species, make use of information directly available from movement of the eyes of an observer in a visual field. A sizeable body of this work has concentrated on how information directly (i.e. without requiring cognition) available from what is known as the 'optic flow' is used to control the timing and co-ordination of movement through the world. The information is derived directly from pattern and motion perception, and guides action without the need for thought, calculation, or conscious decision-making. This model of the visual control of movement has been shown to apply to many areas, including walking, running, driving cars, and flight (both by humans and birds).

From the perspective of ecological perception, the train driver is in a relatively unique and paradoxical situation. Seated at the front of a fast moving vehicle, the driver is on the one hand enclosed in an extremely powerful and compelling optic flow. But, apart from very few situations (such as approaching buffers), the information that their senses would normally draw on to control movement is of no relevance to their driving task. The driver has no control over directional movement, and control of speed is entirely mediated by cognitive processes such as perception of signs, and route knowledge. Whether and how the driver balances the (cognitive) use of the

speedometer and the (direct) use of the optic flow in estimating train speed is not known.

The implications - if any - of this paradoxical situation for driver performance and reliability are not known, and, so far as we are aware, have never been investigated explicitly. It is however possible that it could, for example, contribute to the 'driving without attention' phenomenon (discussed below).

Attention: There are at least two quite distinct meanings of the word 'attention' that are relevant to the study of train driver performance. Both have been intensively investigated for many years in laboratory and practical situations.

The first is often called 'vigilance', and is typical of 'watch-keeping' tasks, where an observer tries to detect signals that arrive at infrequent intervals during which nothing much happens. Often the vigilance decrement can be avoided by an occasional message to the observer that requires an answer. The experimental Pro-Active AWS system (18) provided a means to achieve this in requiring, effectively, an 'answer' from the driver as to the signal aspect that was approaching. The vigilance decrement is also dependent on the extent of trust that the driver has in AWS. The occurrence of a string of frequent signals that do not require a response may also cause the vigilance decrement to increase.

The second meaning of attention is in "dynamic selective attention". In real situations, particularly in the case of vision (though also in hearing), attention tends to be directed to one source of information at a time. To some extent attention can be shared between two or more tasks, although performance will usually be degraded.

It is interesting that although evidence given to the Southall Inquiry maintained that drivers would always be looking at signals before the AWS sounded, in their study of driver eye movements Meret et al. (8) found that on a significant proportion of signals the contrary was true. As predicted by quantitative attention models, attention was elsewhere and was attracted by the sound of the AWS.

Driving without attention: There is considerable interest in psychological research and among safety professionals in the concepts of 'driving without attention' (DWAM), and 'looking-without-seeing'. These notions may have relevance to the AWS context, although no specific literature related to train driving has been identified.

Allen (13) wrote that, "*In modern multi-aspect signalling there is a tendency to shorten the signal sections. [...] When traffic is dense through multi-aspect territory, a train will often come upon a quick succession of double or single yellow aspects. At every one the driver will get an AWS warning, which he has to cancel. In such conditions, the risk of frequent canceling action degenerating into a reflex action that loses its full meaning for the driver is a real one*" (p. 332).

In their extensive review of human factors in road traffic safety, Dewar and Olsen (14) consider the possible reasons for the large number of driving accidents associated with car driver inattention, "highway hypnosis" and "DWAM". They refer to the

state as being associated with monotonous, uneventful driving, where a lack of novelty promotes automatic responses.

Expectancy: One of the key benefits of route knowledge is that it allows the driver to prepare in advance for the route ahead. With good route knowledge, the driver is able to plan ahead, to prepare for events before they occur, and to quickly understand and interpret information. That is, it generates expectancy. In the great majority of situations, expectancy and the ability to plan ahead is effective; indeed, without it, it would not be possible for humans to perform many of the complex tasks they manage. On occasions, however, expectancy can lead to pre-planned behaviours being inappropriate, or can lead to incorrect interpretation of information.

Trust in automation: A growing body of research has been investigating the role that trust and confidence in technology plays in influencing the way people interact with technology-based systems.

If a system is regarded as untrustworthy then it will tend not to be used correctly. Likewise, if an unreliable system is trusted then it will tend to be used even though it malfunctions. Despite the importance of trust in technology, it remains relatively ill-defined and only partly understood psychologically.

The dynamics of trust in technology have been dealt with by Lee and Moray (15) and Parasuraman and Riley (19) among others in connection with supervisory control and automated systems. They define the dynamics of trust around predictability, dependability and faith. In simplistic terms, trust is established by consistent and desirable behaviour. This relies on the observability of system behaviour, as well as the individual actively sampling it. As trust is gained, the focus shifts away from observing specific behaviours towards assessing the global disposition of the device (it is said to be “dependable”). Over time, trust serves to reduce the effort expended in checking the system to ensure it is performing as expected.

In our questionnaire survey of 277 drivers conducted as part of the AWS study (1), a substantial proportion reported issues suggesting drivers know they cannot always trust an AWS warning.

Simple heuristics and recognition-primed decision making: Perhaps most intriguing, is recent work by Gigerenzer and Todd on what they term 'Simple Heuristics' (16). Gigerenzer and Todd take the view that complex human performance in the real world cannot possibly involve the rational cognitive strategies assumed by much experimental psychology. Understanding decisions in the real world requires different, more psychologically plausible explanations. The work on simple heuristics focuses on what they term "fast and frugal heuristics - simple rules for making decisions with realistic mental resources".

The notion of simple heuristics is similar to the well-established concept of recognition-primed decision making (see for example, Klein 1989, 2003). Instead of emphasising heuristics and rules, Klein's model of decision making is bound in the ability of humans to pattern match. Repeated experiences are unconsciously linked together to form a pattern. Links between features in past and current situations enable

typical features of situations to be extracted, and typical responses generated (Klein, 1989). This model is similar to the well-known GEMS model of human error (17).

Whether or not 'simple heuristics' and/or 'recognition-primed decisions' are important in determining train driver performance remains to be seen. It might seem surprising if, in some form, they were not.

Summary of psychological thinking

So how does consideration of these wide variety of psychological constructs help in understanding train driver performance? Two general observations seem relevant.

The first is that, however desperately the human factors professional might seek tools and analysis methods that can be applied systematically across application domains, the fact is that the basis of human performance is exceedingly complex. Generic information processing 'models' of the human operator can certainly provide extremely useful generalizations. They have value in systems engineering, and in seeking to predict asymptotic limits of human performance in complex systems. Even relatively simple models can have great value as engineering tools in providing structure and drawing attention to human capabilities.

However, as a means of understanding why specific individuals, on a particular day in a particular set of circumstances behaved (or failed to behave) in a particular way, or predicting how an individual might behave in a given set of unexpected circumstances, such models have very little to offer.

The second general observation is that if such an understanding of what happened in an incident, or what is likely to happen, is important, it is essential to identify and understand the characteristics of the context and situation at the time the behaviour of interest occurred or is expected to occur.

A Situational Model of driver performance with AWS

A principle conclusion of the study, then was that driver performance with AWS can only be understood in terms of the context and situation at the time the system is intended to influence driver behaviour. From this perspective, existing approaches to modeling driver performance appear relatively superficial. Developing a comprehensive model of the psychology of driver behaviour would be a significant task. However, based on the material reviewed, and drawing on insight from discussion with drivers, a tentative framework that might underpin the assessment of risk associated with Extended AWS was developed.

The objective was to provide a framework to help understand the possible state of cognition of a particular driver between the point at which the AWS signal is detected and the driver's subsequent behaviour. The emphasis on time reflects the fact that many of the circumstances and situational factors involved are inherently time-dependent, or time-limited. These include:

- Δ The role of time in controlling the safe passage of a train (such as the time for which signals are visible, and the time involved in accelerating or braking)

- Δ Events or features which are only true at particular locations on a track
- Δ Expectations or beliefs a driver might hold because of the way in which events have evolved over time prior to an event (due to the spatial relationship between signals and speed restrictions, say, or because of the relative speed and distance between trains further down the track)

Figure 1 illustrates a “Situational Model” of driver behaviour. A number of points of introduction are needed to put this framework into context:

- Δ The Situational Model has been developed as a framework to direct attention towards those situational factors which seem to be important in understanding driver behaviour.
- Δ The model emphasises the state of cognition at a specific time (“Now”) in a specific situation and a specific context. The emphasis is therefore on the immediate situation facing the driver, and the immediately preceding history (what is shown on the framework as “The Now”). Longer term knowledge and experience is clearly important. However, the emphasis is on understanding the behaviour of a particular individual at a particular moment.

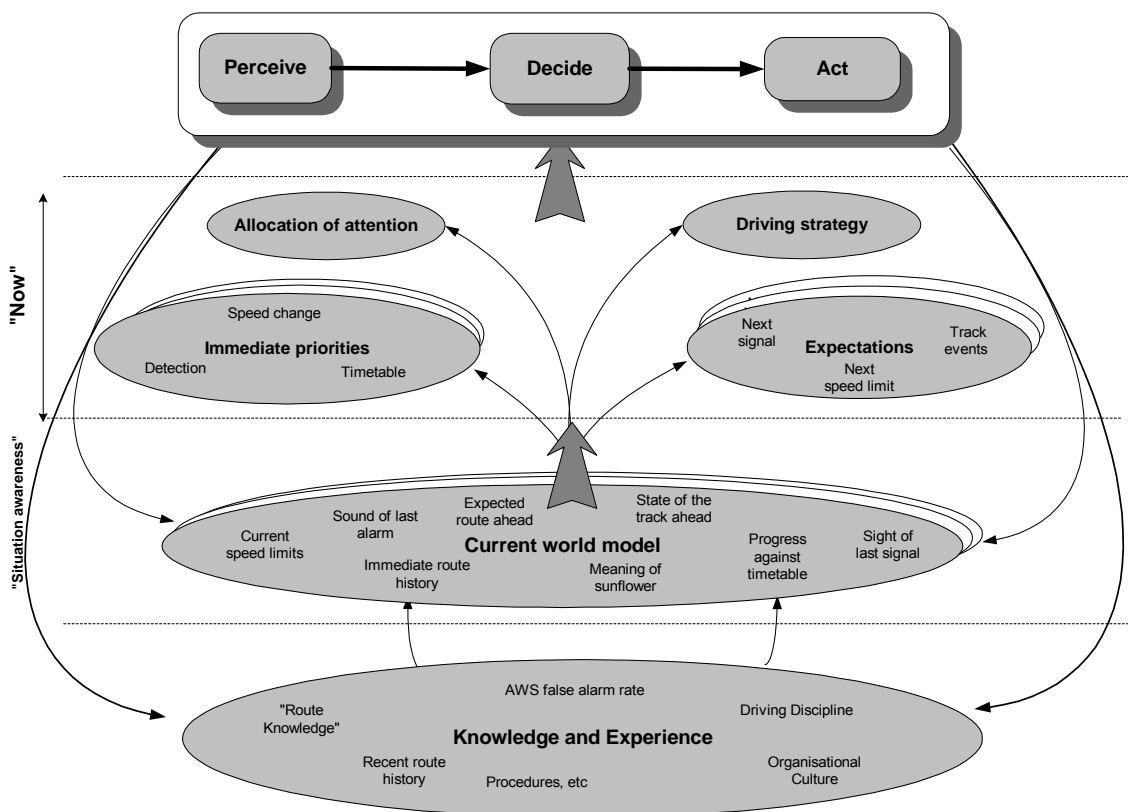


Figure 1: Situational model of driver performance in interacting with AWS

- Δ The situational model is aligned against a standard information processing model (Perceive-Decide-Act) of human performance. This seeks to contextualise the model of the human as an information processor in terms of the situational factors that determine actual performance in a real-world, real-time, environment. This emphasises that any analysis of driver performance using an Information Processing model needs to take proper account of situational issues.
- Δ The model can be no more than a ‘snapshot’. Details of what matters at any moment, and how factors interact will vary over time, particularly as the ‘state’ of the driver changes. For example, different states of attention and cognitive control could be particularly powerful in mediating how elements of the framework might interact at any time. Momentary distractions might also disrupt interactions between elements of the framework.

Briefly, the framework incorporates the following assumptions (each of these is represented as an element on the figure);

- Δ driver performance is based on a foundation of knowledge and experience derived over a relatively long time from training and general experience. As well as explicit knowledge of the rules, procedures, and disciplines governing professional driving, this knowledge base includes recent experience of the route being driven (e.g. the next signal has been at caution on every approach over the past fortnight). It will also include such things as the typical false alarm rate associated with AWS signals.
- Δ the driver has a mental representation of the current state of the world as it affects the current driving activity. This includes elements such as the current speed limit, state of the track (are track side workers in the vicinity) and significant features in the route over a psychologically relevant timescale ahead. It will also hold awareness of the immediate history (the last six caution signals have cleared before the train reached them because of a queue of trains ahead). Much of this world model will be held sub-consciously.
- Δ two elements are particularly important in determining the driver’s cognitive state at any specific moment (“Now”). These are the **immediate priority** and the **expectation** of what the world will be like in the next few moments. The immediate priority might be to reduce speed, or to ensure every individual in a track-side party can be seen. Expectation might include the location of the next signal, the point at which a speed limit will come into force, the expectation that an unsuppressed AWS magnet is about to cause an alarm that can be ignored, or the expectancy that the signal will change it’s state while its still in view.
- Δ both immediate priorities and expectations are continually changing and updating as the train progresses along the route.
- Δ The driver directs attentional resources based on current priorities and expectations. Responses to events in the outside world are determined by a driving strategy which can vary from moment-to-moment, within reasonably broad constraints. The nature of the driving strategy, and especially how they vary, will depend extent on the drivers training, experience and confidence.

We have found that the situational model of driver behaviour provided an extremely useful framework for assessing the risk of driver unreliability with AWS. Consideration of the model led to eighteen situational factors being identified which might influence the drivers state of cognition at the time AWS warnings are encountered (i.e. “Now” on the model). These eighteen factors provided the basis for an assessment of the likely risk associated with twenty scenarios involving encounters with AWS (see 1 for a summary of the assessment).

Summary

The study reported in this paper sought to understand the nature of driver reliability with the Automatic Warning System, and to estimate the likely change in driver reliability if the use of AWS is further extended. Considerations from the psychological research base identified a reasonably large number of mechanisms and constructs that might be important in understanding moment-to-moment driver performance. Most important among these appear to be:

- Δ the importance of expectancy and route knowledge
- Δ the drivers mental representation of the situation and how the understanding of immediate priorities and objectives are maintained and controlled
- Δ trust in the system, specifically, the extent to which drivers’ experience might cause them to expect a certain level of ‘false’ alarms or incorrect signals.

Both the situational model of driver behaviour, and the method used to assess driver reliability with AWS might have wider application in understanding safety issues associated with driver performance.

References

1. McLeod, R. W., Walker, G. Moray, N., and Love, G (2003) Project Summary Report: Driver reliability with Extended AWS Nickleby HFE Ltd B/C271/FD.5
2. McLeod, R. W., Walker, G., Mills, A (2003) Assessing the Human Factors risks in extending the use of AWS 1st European Conference on Rail Human Factors
3. Vaughan, A. (2000). Tracks to Disaster. Shepperton, Surrey: Ian Allen.
4. Hall, S. (1999). Hidden dangers: railway safety in the era of privatisation. Shepperton, Surrey: Ian Allen.
5. Vicente, K.J. (1999). Cognitive Work Analysis: Towards Safe, Productive and Healthy Computer Based Work. Marwah, N.I: Lawrence Erlbaum.
6. Hollnagel, E. (1998). Cognitive Reliability and Error Analysis Method. London: Elsevier.
7. McLeod, R. W., Walker, G. Moray, N., (2003) Extended AWS Study: Review of the knowledge base Nickleby HFE Ltd B/C271/FD.8
8. Merat, N., Mills, A., Bradshaw, M., Everatt, J., & Groeger, J. (2002). Allocation of attention among train drivers. In P. T. McCabe (Ed.), Contemporary Ergonomics 2002. London: Taylor and Francis.

9. Sarter, N. B., & Woods, D. D. (1997). Team play with a powerful and independent agent: Operational experiences and automation surprises on the Airbus A-320. *Human Factors*, 39, (4), 553-569.
10. Rasmussen, J. Pedersen, A-M, & Goodstein, L. (1995). *Cognitive Engineering: Concepts and applications*. New York: Wiley.
11. Hutchins, E. (1995). *Cognition In The Wild*. Cae: MIT Press.
12. Gibson, J. J. (1979) *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.
13. Allen, G. F. (1965). *British rail after Beeching*. London: Ian Allen.
14. Dewar, R. E., & Olsen, P. (2002). *Human Factors in Traffic Safety*. Criterion Press.
15. Lee, J. & Moray, N. (1992). Trust, control strategies and allocation of function in human-machine systems. *Ergonomics*, 35, 10, 1243-1270.
16. Gigerenzer, G. & Todd, P. M. (1999). *Simple Heuristics That Make Us Smart*. Oxford: Oxford University Press
17. Reason, J. (1990). *Human Error*. Cambridge: Cambridge University Press.
18. Dore', D. M. (1998). *The Pro-Active Automatic Warning System (PAWS)*. Report to the New Measures Working Group.
19. Parasuraman,, R., and Riley, V. 1997. Humans and automation: use, misuse, disuse, abuse. *Human Factors*, 39(2), 230-253. IS